

Validation of 3D Radiative Transfer in Coastal-Ocean Water Systems
as Modeled by DIRSIG

by

Jacqueline Amy Speir

B.S McGill University, 1999

M.S. CUNY John Jay College of Criminal Justice, 2004

A dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Chester F. Carlson Center for Imaging Science
Rochester Institute of Technology

21 July 2010

Signature of the Author _____

Accepted by _____
Coordinator, Ph.D. Degree Program Date

CHESTER F. CARLSON CENTER FOR IMAGING SCIENCE
ROCHESTER INSTITUTE OF TECHNOLOGY
ROCHESTER, NEW YORK

CERTIFICATE OF APPROVAL

Ph.D. DEGREE DISSERTATION

The Ph.D. Degree Dissertation of Jacqueline Amy Speir
has been examined and approved by the
dissertation committee as satisfactory for the
dissertation required for the
Ph.D. degree in Imaging Science

Dr. John Schott, Dissertation Advisor

Dr. Carl Salvaggio

Dr. Emmett Ientilucci

Dr. Joseph DeLorenzo

Date

DISSERTATION RELEASE PERMISSION
ROCHESTER INSTITUTE OF TECHNOLOGY
CHESTER F. CARLSON CENTER FOR IMAGING SCIENCE

Title of Dissertation:

**Validation of 3D Radiative Transfer in Coastal-Ocean Water Systems as
Modeled by DIRSIG**

I, Jacqueline Amy Speir, hereby grant permission to Wallace Memorial Library of R.I.T. to reproduce my thesis in whole or in part. Any reproduction will not be for commercial use or profit.

Signature _____ Date _____

Validation of 3D Radiative Transfer in Coastal-Ocean Water Systems as Modeled by DIRSIG

by

Jacqueline Amy Speir

Submitted to the
Chester F. Carlson Center for Imaging Science
in partial fulfillment of the requirements
for the Doctor of Philosophy Degree
at the Rochester Institute of Technology

Abstract

The radiative transfer equation (RTE) is a mathematical description of radiative gains and losses experienced by a propagating electromagnetic wave in a participating medium. Except for an isotropic lossless vacuum, all other volumes have the potential to scatter, absorb and emit radiant energy. Of these possible events, the global scattering term is the greatest obstacle between a radiative transfer problem and its solution. Historically, the RTE has been solved using a host of analytical approximations and numerical methods. Typical solution models exploit plane-parallel assumptions where it is assumed that optical properties may vary vertically with depth, but have an infinite horizontal extent. For more complicated scenarios that include pronounced 3D variability, a Monte Carlo statistical approach to the radiative transfer solution is often utilized. This statistical approach has been integrated within the Digital Imaging and Remote Sensing Image Generation (DIRSIG) model in the form of photon mapping. Photon mapping provides a probabilistic solution to the in-scattered radiance problem, by employing a two-pass technique that first populates a photon map based on a Monte Carlo solution to the global scattering term, and then later uses this map to reconstruct the in-scattered radiance distribution during a traditional raytracing pass. As with any computational solution, the actual implementation of the technique requires assumptions, simplifications and integration within a cohesive rendering model. Moreover, the realistic simulation of any environment requires several other radiometric solutions that are not directly related to the photon mapped in-scattered radiance. This research attempts to validate raytraced and photon mapped contributions to sensor reaching radiance that can be expected in typical littoral environments, including boundary interface, medium and submerged or floating object effects. This is accomplished by comparing DIRSIG modeled results to those predicted analytically, by comparison to other numerical models, and by comparison to observed field

phenomenology. When appropriate, first-order estimates of a computational solution's ability to render a given phenomenon are provided, including any variance or bias that may result as a function of the user-specified solution configuration.

Acknowledgements

I would like to thank my advisor, Dr. John Schott, for his constructive criticism, comments and suggestions regarding this body of work. On a personal note, I couldn't have asked for a more productive and enjoyable working environment and relationship.

I would like to sincerely thank each committee member for their respective contributions to this research. I particularly appreciate Dr. Carl Salvaggio's guidance on the topics of C++ and computational radiometry, Dr. Emmett Ientilucci's considerable effort in making corrections to this manuscript, and Dr. Joseph DeLorenzo's encouragement and flexibility regarding the entire process.

In addition to the committee members, there are several individuals I would like to acknowledge. First and foremost, I would like to thank Dr. Adam Goodenough. I will be eternally grateful for the role Adam played in allowing me to accomplish this research. My progress was considerably dependent on Adam's mentorship and technical expertise, both of which he gave so freely.

I would also like to thank Dr. David Messinger for his continual support and encouragement; several aspects of my research experience at RIT would not have been possible without his help. I want to thank Scott Brown for his words of encouragement. He had a knack of finding a way to offer encouragement when I needed it most (and usually before I even realized that I needed it). I would like to thank Dr. Michael Gartley for his help with the BTDF and concepts related to ESFs and PSFs. I would also like to thank Cindy Schultz; her assistance was invaluable to my survival at RIT over the past four years. Finally, I would like to thank Paul Mezzanini and all the Research Computing staff for their assistance with large memory DIRSIG simulations. Their help was invaluable.

Last but certainly not least, I want to thank my family. I am extremely lucky to have such a supportive network filled with people that always find a way to amaze me. They consider encouragement to be more than just words, and without fail, they can be counted on to actually get involved when I need them most...thank you Mom, Dad, Shelly and Jack.

This work is dedicated to my family.

Contents

1	Overview & Motivation	12
1.1	Work Statement	14
1.1.1	Overview of Proposed Methodology	17
1.1.2	Scope	19
1.2	Outline of Subsequent Chapters	21
2	Water	23
2.1	The Scattering Phase Function (Indicatrix) and Absorption Coefficient . . .	24
3	The Radiative Transfer Equation	29
3.1	Variable & Clarification	29
3.2	Overview of the Radiative Transfer Equation	30
3.3	Overview of Analytical Approximations to Radiative Transfer	32
3.4	Overview of Photon Mapping	33
3.5	The “ <i>Big Equation</i> ”	34
4	Simple Boundary Interface Effects	35
4.1	Proposed Methodology - Fresnel Reflectance	35
4.1.1	Uniform Illumination	35
4.1.2	Results	39
4.1.3	Solar and Skylight	39
4.1.4	Results	44
4.2	Proposed Methodology - Fresnel Transmittance & Attenuation	48
4.2.1	Fresnel Transmittance & Snell’s Cone	48
4.2.2	Results	52
4.2.3	Proposed Methodology - Fresnel Transmittance, MediumPM & the Generic Radiometry Solver	54
4.2.4	Results	57
4.3	Proposed Methodology - Refraction	59
4.3.1	Predicted Results	59

4.3.2	Results	65
4.4	Transition From Simple Interface Problems to More Complicated Boundaries	67
5	Waves & Caustics	68
5.1	Overview	68
5.2	Sub-Pixel Capillary Waves	69
5.2.1	Bidirectional Reflectance Distribution Function (BRDF)	69
5.2.2	Results	75
5.2.3	Bidirectional Transmittance Distribution Function (BTDF)	76
5.2.4	Results	81
5.3	Sinusoidal Gravity Waves - Caustics	83
5.3.1	Results	85
5.4	Sinusoidal Gravity Waves - Combination of the Interface and Scattering	90
5.5	Transition to Bottom Effects	94
6	Bottom Effects, Skyfraction & Background Reflection	95
6.1	Overview	95
6.2	Comparison of GRS, SRS and Analytical Predictions for Skyfraction & Background Reflection in the Absence of a Water Volume	96
6.2.1	The Generic Radiometry Solver	96
6.2.2	GRS Configuration	101
6.2.3	SurfacePM versus the Generic Radiometry Solver	106
6.2.4	Analytical Predictions	109
6.2.5	Results	115
6.3	Comparison of GRS, SRS and Analytical Predictions for Skyfraction in the Presence of an Air-Water Interface	119
6.3.1	Refraction & Skyfraction	119
6.3.2	Results	128
6.4	Comparison of GRS, SRS and Analytical Predictions for Skyfraction & Attenuation in the Presence of a Water Volume	131
6.4.1	Attenuation as a Function of Absorption	131
6.4.2	Analytical Prediction	131
6.4.3	Results	134
6.4.4	Attenuation as a Function of Absorption & Scattering - Using the Full Scattering Coefficient	135
6.4.5	Results	138
6.4.6	In-Scattering	147
6.4.7	The Effective Scattering Coefficient	151
6.4.8	Results: Attenuation - SRS versus GRS	152
6.4.9	Results: In-Scattering - SRS versus GRS	153

7	Absolute Sensor Reaching Radiance (& Irradiance) - Direct Reduced Plus In-Scattered	155
7.1	The Two-Stream Approximation	156
7.1.1	Results	162
7.2	Canonical Problem # 6	163
7.2.1	Results	163
7.3	Field Profiles	167
7.3.1	Results	172
8	Ship-Shadow and Self-Shading	174
8.1	Proposed Methodology - Simulated Ship-Shadow	180
8.2	Validation of a 3D Radiative Problem: Ship-Shadow	187
9	Image Quality	189
9.1	Overview of Density Estimation	190
9.1.1	Univariate Density Estimation	190
9.1.2	Mean Square Error	192
9.1.3	Multivariate Density Estimation	193
9.1.4	Radiance Reconstruction as a Density Estimation Problem	195
9.2	User-Defined Photon Mapping Variables	197
9.3	Evaluation of Image Quality	200
9.4	Expected Trends	202
9.5	Scenes with Considerable Spectral Sampling	206
9.6	Scenes with Large Spatial Extent	207
9.7	Noise and Asymmetry in the PSF	208
9.8	Theoretical Proof of Concept	208
9.8.1	Simulation Trade Space	209
9.9	Alternative Uses, Outcomes & Considerations	211
9.10	Integration within DIRSIG: The Virtual Step-Function	212
9.11	Results: SNR	213
9.11.1	Summary of Factors that Influence SNR	229
9.12	Results: PSFs	231
9.12.1	Summary of PSF Results	235
9.13	Spectral Considerations	238
9.13.1	Results: Spectral SNR	238
9.13.2	Results: Spectral PSF	248
9.13.3	Summary of Spectral Considerations	248
9.14	Image Quality Overview	252
9.15	Collection of Panchromatic Images	252
9.16	Collection of Spectral Images	277

10 Visual Demonstrations	282
10.1 Coastal Scene	282
10.2 Results	285
10.3 Image Summary	292
10.4 Miscellaneous Visual Demonstrations	302
10.4.1 Surface & Volumetric Shadows	302
10.4.2 Caustics	302
10.5 Chapter Summary	304
11 Final Comments	305
Appendices	312
A Two-Stream Solution & Related Variables	312
A.1 The Average Cosine of the Radiance Field	312
A.2 Uniform Illumination, Isotropic Scattering & a Non-Reflecting Bottom Boundary	315
A.2.1 Uniform Illumination, Isotropic Scattering & a Lambertian-Reflecting Bottom Boundary	319
A.2.2 Anisotropic Two-Stream Approximation	322
B Trends in Variance	324

Chapter 1

Overview & Motivation

The radiative transfer equation (RTE) is a mathematical description of radiative gains and losses experienced by a propagating electromagnetic wave in a participating medium. Except for an isotropic lossless vacuum, all other volumes have the potential to scatter, absorb and emit radiant energy. Of these possible events, the global scattering term is the greatest obstacle between a radiative transfer problem and its solution. Historically, the RTE has been solved using a host of analytical approximations and numerical methods. The solution models usually exploit plane-parallel assumptions where it is assumed that the optical properties of the volume may vary vertically with depth, but have an infinite horizontal extent. For more complicated scenarios, a Monte Carlo (MC) statistical approach to the radiative transfer (RT) solution is often utilized. The MC model can handle irregular geometries not suited toward plane-parallel solutions of the RTE, although it is acknowledged that geometric flexibility is bought at the price of time consuming computations (Thomas and Stamnes 1999). This is because the Monte Carlo approach to the RT solution simulates photon trajectories using probability distributions described by the material's inherent optical properties (IOPs), and depending upon the complexity of the scenario, a very large number of such trajectories may be necessary in order to accurately render a given phenomenon.

Although Monte Carlo methods have existed since the 1940's (Mobley 1994), ever-increasing computer efficiency and resources have recently enabled integration of MC RTE solutions into synthetic image generation programs. The technique produces visibly convincing simulations of volume and subsurface scattering, and has been used for a variety of applications in the field of Computer Graphics. Although at heart a Monte Carlo method, the technique of photon mapping (PM) has drastically increased the computational efficiency of MC models (Goodenough 2007). The increased efficiency is achieved by using a two-pass technique that first constructs the photon map using a MC solution to the RTE, and then later uses the map to reconstruct the in-scattered radiance distribution during a traditional raytracing pass. Given the successful utilization of the technique in

Computer Graphics’ applications, Monte Carlo photon mapping capabilities were recently added to the Rochester Institute of Technology’s Digital Imaging and Remote Sensing Image Generation (DIRSIG) model (Schott et al. 1999). The DIRSIG model is a ray-tracer with sophisticated Radiometry Solvers that strives to render synthetic imagery that is both physically and radiometrically accurate. The addition of PM capabilities to DIRSIG extends the model’s ability to render various participating media and associated phenomenon, including complex coastal water volumes.

Anecdotal and empirical evidence indicates that PM techniques recreate radiometrically accurate radiance distributions. This is primarily supported in two ways. First, the MC PM solution to the RTE is based on a physically and mathematically realistic solution of the RTE. Second, the technique produces visibly realistic renderings, and, it is logical to conclude that *visibly realistic* solutions are likely to be *radiometrically realistic*. Despite this (and as with any computational solution), the actual implementation of the technique requires assumptions, simplifications and integration within a cohesive rendering model. Moreover, the realistic simulation of any environment requires several other radiometric solutions that are not directly related to the photon mapped in-scattered radiance. For example, if one considers a coastal water environment, then the sensor reaching radiance will be a function of a host of boundary interface effects, medium effects (including refractive index, in-scattering and extinction properties) and submerged or floating object effects. Therefore, rendering truly interesting scenes that leverage the benefits of a photon mapped in-scattered radiance distribution goes well beyond validation of the Monte Carlo solution of the RTE.

To facilitate the validation of radiative transfer codes, a group of researchers in the 1990’s published the description and radiative transfer solution for several stylized in-water canonical problems (Mobley et al. 1993). Given the time period and preponderance of plane-parallel models in the literature, it is not surprising that the proposed validation suite lacks problems and solutions with inherent 3D complexity, and therefore is unable to tax or leverage the benefits of a photon mapped solution. Despite this, there is still great value in ensuring that DIRSIG’s photon mapped solution is intrinsically consistent with reported results for plane-parallel geometries. As such, DIRSIG’s photon mapping capabilities were compared against 1D numerical RT solution models for in-water radiative transfer (Goodenough 2007). The existing validations (see Goodenough (2007) Chapter 7) have provided very promising results. Since details concerning the DIRSIG model, the photon mapping theory, and the radiometric validation for plane-parallel radiative transfer problems can be found in Schott (2007) Chapter 14, Schott et al. (1999) and Goodenough (2007), they will not be reiterated here. Instead, the proposed research described in subsequent chapters will discuss methods that extend the existing validation of DIRSIG’s photon mapping module to include complex boundary, object and medium interactions. The goal here is not the singular validation of the photon mapping radiometry module since this can reasonably be assumed given DIRSIG’s success with the set of canonical problems

as shown by Goodenough (2007). Instead, a medium-specific *systems* based approach is sought that validates radiometric solutions that are a function of boundary, medium, object and Radiometry Solver interactions, where the photon mapped in-scattered radiance is but part of the sensor reaching radiance solution.

Conceivably, DIRSIG’s photon mapping module can be exploited to simulate any participating medium of interest. However, ocean water and coastal environments provide a convenient simulation system for validation. First, radiative transfer in coastal waters was the original motivation for incorporation of the PM module into DIRSIG and therefore contains the requisite interface necessary to utilize ocean water IOPs. Second, there is extensive historical data describing the optical and radiometric properties of water, which in turn allows physical and radiometric fidelity comparisons between simulated DIRSIG results and documented phenomenon. Third, when combined with boundary and submerged and/or floating object effects, the proposed simulation scenarios present complex and interesting 3D radiative transfer problems.

Techniques will be developed in subsequent chapters to assess not only the PM module’s scattering code, but also its integration and description of *system* related phenomenon, where the *system* is a specific medium of interest with additional boundary and object effects. The work proposed here is meant to be both complimentary and additive in content to the former validation efforts carried out by Goodenough (2007). With these goals in mind, several interesting and relevant in-water radiative transfer scenarios are considered. Each proposed scenario illustrates typical radiometric solutions and/or interactions that might be encountered when modeling coastal ocean waters. The remainder of this chapter summarizes the proposed work statement and the validation approach.

1.1 Work Statement

Since it is unreasonable to attempt to model every conceivable ray path and interaction that might occur during in-water radiative transfer, the validation effort proposed here considers a limited number of scenarios of interest. Figure 1.1 is an attempt to group the validation scenarios or solutions into simple conceptual categories related to *medium*, *boundary interface* and *submerged or floating object effects*. The *medium effects* are further subdivided between *source function* and *non-source function effects*. The source function refers to in-scattered gains in radiance that occur during a ray’s transit through the medium, while the non-source related phenomena are a function of the medium’s refractive index and extinction coefficient.

Alternatively, the individual significance of each proposed validation test can be illustrated by considering a ray’s transit through a water volume. Even for a simple scene (see Figure 1.2 for the following discussion), a multitude of fundamentally different ray paths and interactions can occur. The phenomenon are denoted with the letters (a.) through (h.), and seven individual ray paths are depicted (for simplicity and clarity, the

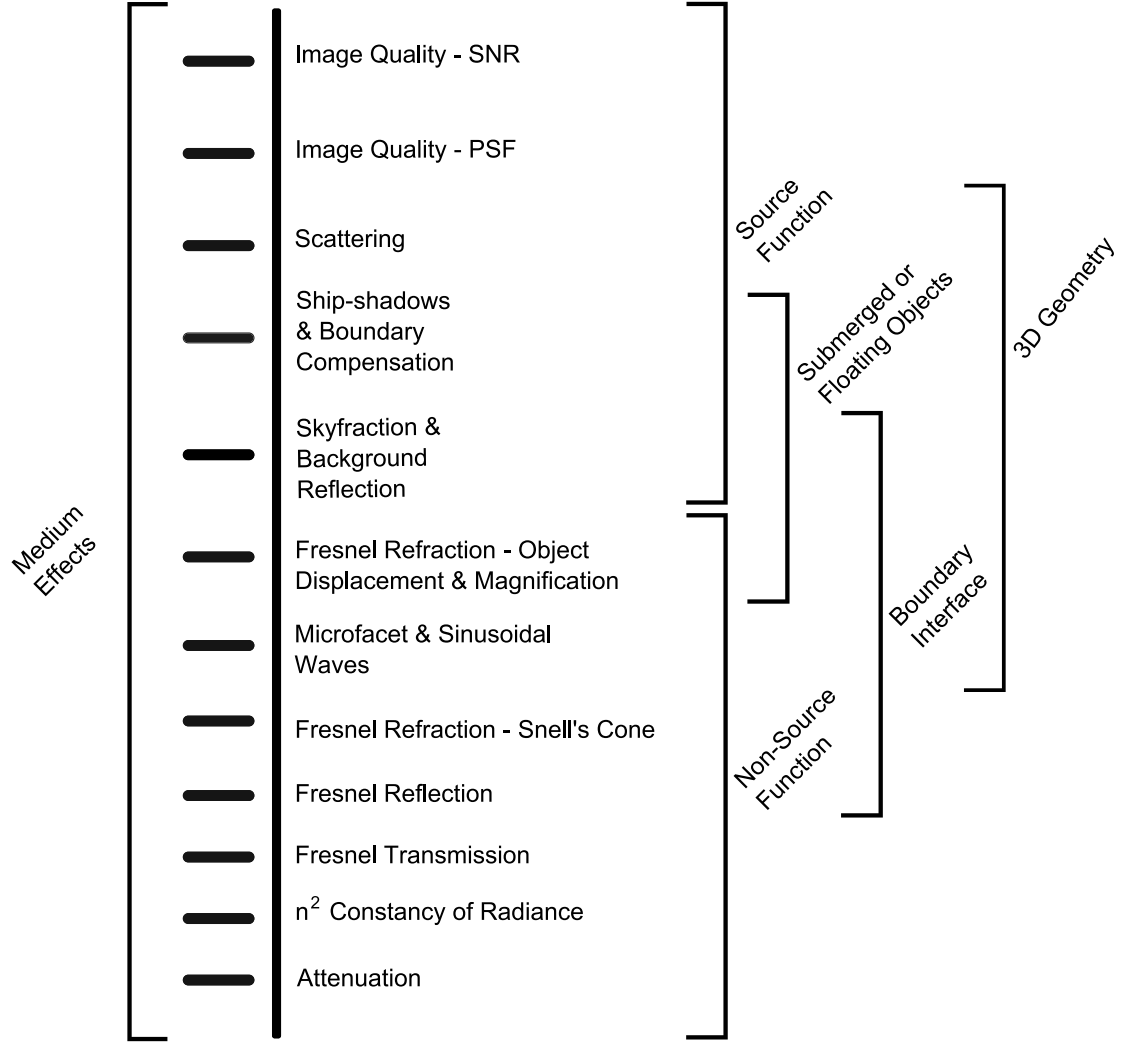


Figure 1.1: Phenomenon to be validated.

ray's propagation toward the detector is not illustrated). The phenomenon are as follows: (a.) reflection, (b.) refraction, transmission, object displacement and magnification (c.) attenuation, (d.) in-scattering, (e.) skyfraction, background reflection and boundary compensation, (f.) volumetric ship-shadows, self-shading and boundary compensation, (g.) image contrast and signal to noise ratio (SNR) as a function of target albedo and volume IOPs, and (h.) Snell's cone and the n-squared constancy of radiance. [Note that phenomenon (a.) through (d.) are pertinent to each and every ray that ends up inside

the medium, however for clarity they are listed as progressive contributions for rays (1.) through (3.).]




Figure 1.2: Ray paths and interactions for a simple scene.

Each ray path corresponds to one or more of the phenomenon conceptually grouped in Figure 1.1. For example, when a ray enters the participating medium, it must traverse the boundary interface. Within DIRSIG, the Fresnel Radiometry Solver determines the radiance distribution as a function of (a.) reflection and (b.) refraction and transmission at the air-water interface for planar or sinusoidal gravity waves. If a microfacet bidirectional reflectance distribution function (BRDF) is employed, then the Generic Radiometry Solver is queried instead. The reflected and transmitted radiance contributes to sensor reaching radiance, and passage through the interface redirects the ray as a function of refraction, which will impact the magnification and displacement of objects viewed through the water volume.

Once inside the medium, attenuation occurs as a function of absorption and out-scattering (c.). This is partially balanced by photon mapped in-scattering (d.). The solution is a function of the illumination, geometry and inherent optical properties of the medium. For three-dimensional submerged objects, skyfraction, background reflection and boundary compensation effects (e.) must also be accounted for. For submerged in-water

targets, image contrast and transfer (g.) through the scattering volume will also be a function of the location and turbidity of the water layers, the target/surround albedo, and possible caustics. For submerged detectors, the refractive index of the water volume, and the n-squared constancy of radiance (h.) account for the magnitude and distribution of sensor reaching radiance. Finally, submerged and floating objects induce high frequency radiance features, such as self-shading and ship-shadows (f.). Since image quality is a global attribute, it is not specifically illustrated by any particular ray, but must account for the multitude of possible illumination, geometry and inherent optical properties that might be encountered in typical ocean water simulations.

Collectively, the interesting ray paths (as illustrated in Figure 1.2) and conceptual phenomena (as grouped in Figure 1.1) represent a narrow-scope, end-to-end validation of in-water radiative transfer. Unfortunately, the scenarios described in Figures 1.1 and 1.2 appear deceptively simple. The deception exists because simulating interesting littoral scenes is inherently a three-dimensional radiative transfer problem, and its solution is based on *global* interactions that cannot be reasonably emulated analytically except under highly contrived situations. Since the proposed research does not include generation of yet another *numerical* algorithm to solve the *numerical* problems that DIRSIG's state-of-the-art photon mapping and Radiometry Solvers are designed to handle, the proposed scenarios must be validated using a creative combination of analytical evaluation, inference, visual demonstration, and comparison with experimental (and numerically predicted) observations. To elaborate, the following section will briefly describe the proposed methodological approach to validation.

1.1.1 Overview of Proposed Methodology

The combined validation of *boundary interface effects*, *submerged or floating object effects* and both *source* and *non-source medium effects* represents a comprehensive radiometric validation for typical and interesting in-water radiative transfer solutions. The remainder of this work is divided into chapters that provide introductory and background information on specific scenarios or solutions to be validated. Each proposed validation methodology is referenced according to the phenomenon under investigation (*e.g.* Fresnel reflection, Fresnel transmission, refraction, skyfraction, etc.), and associated with either the *boundary interface*, *submerged or floating objects*, *medium effects* and/or *3D geometry*.

The validation effort for any given scenario or solution will involve one or more of the following approaches: *DIRSIG model versus analytical prediction*, *DIRSIG model versus alternative numerical solution* and/or *DIRSIG model versus observed phenomenology* (*e.g.* *field or laboratory*). For brevity of expression, the aforementioned approaches will be referred to as *model-vs-analytical prediction*, *model-vs-model comparison* (assumed numerical), and *model-vs-observed phenomenology*, respectively. The model-vs-analytical prediction approach will be utilized for scenarios that lend themselves toward description by a governing equation, either precisely, or with approximate simplifying assumptions.

The analytical computations and approximations are deemed valuable to a validation effort since they provide *transparent*, *concise* and *rapid* predictions for phenomenon that are otherwise numerically determined in a time-consuming and opaque manner. Furthermore, if these scenarios are thoughtfully developed, then relatively simple analytical results can be extended to validate more complex phenomenon through inference. For example, refraction through a cube of water handles three possible fates for a given ray (see Figure 1.3): (a.) refraction at the horizontal air-water interface, (b.) refraction at the horizontal air-water interface followed by total internal reflection from the vertical air-water interface, and (c.) refraction at the horizontal air-water interface, followed by refraction at the vertical air-water interface. However, validation of these comparatively simple refraction events using a cube of water theoretically validates complex refraction events that might occur for a wavy air-water interface. This means that a simple refraction validation using a cube can indirectly validate refraction events for arbitrary geometries (see Figure 1.3).




Figure 1.3: Three major refraction events that can be analytically computed for a cube of water, and indirectly extended to validate refraction through a wavy air-water interface. See text for a description of paths (a.) through (c.). Note that path (c.) for the water cube eventually strikes the ground plane beneath the cubed water volume, and for the wavy interface a ray that might strike an object floating near the ocean surface.

The model-vs-model comparison will primarily involve comparison of computational mechanisms used to numerically integrate sensor reaching radiance for a scene with fixed physics-based interactions. When appropriate, DIRSIG results will be compared to published data generated by extrinsic (independent) numerical models. In all other circumstances, DIRSIG results will be generated and compared using two or more intrinsic DIRSIG Radiometry Solvers. The intrinsic DIRSIG Radiometry Solvers of interest are the Generic Radiometry Solver (GRS) and the SurfacePM Radiometry Solver (SRS). The premise is that for a fixed simulation with constant underlying physics-based interactions, numerically convergent sensor reaching radiance as a function of GRS and SRS serves to

validate both numerical models simultaneously. Since it is unlikely that all phenomenology will be equally well reproduced by both Radiometry Solvers, this comparison will also highlight the strengths and weaknesses of each computational model, and therefore equip the user with insight so that he or she may select the appropriate Radiometry Solver based on user-defined goals. If appropriate, the model-vs-model comparison will be supplemented by model-vs-analytical predictions and/or model-vs-observed phenomenology. Of particular interest here will be a comparison of the sensor reaching radiance computed by Generic and SurfacePM in the presence of (1.) skyfraction and (2.) background reflections. For clarity, and whenever reasonable, the solutions will be validated for scenes void of water before extension to in-water predictions.

Although the analytical approximations provide a “sanity-check” and insight into expected trends, they may fail to predict details when applied to complicated geometries expected from real world scenarios. Moreover, a comparison of one numerical model to another, although a powerful indicator that numerically disparate models can provide a consistent description of some form of underlying physics, they can equally (or differentially) suffer from numerical assumptions and therefore fail to mimic real world physics. The real test is how well a model reproduces empirical observations. Unfortunately, independent experimental validation, especially for ocean water scenarios, is inherently difficult. The latter requires extensive laboratory equipment to assess optical parameters and measure radiance distributions. Furthermore, experimental studies require well-conceptualized radiative transfer problems in which to simulate experimentally, and then emulate synthetically. On the other hand, oceanographic research has repeatedly used a limited number of field and tank experiments for comparative purposes [*e.g.* comparing measured and small angle scattering approximation (SAA) computed point spread functions (PSFs) and beam spread functions (BSFs)]. Therefore, independent field experiments will be cursory at best.

1.1.2 Scope

To clearly define scope, the following itemized list provides an overview of the proposed in-water radiative transfer validation, divided into three main phases of research:

- *Phase One: Validation of contributions to sensor reaching radiance as a function of boundary interface, medium and in-water objects.*
 - Evaluate the radiance distributions as a function of reflection, refraction and transmission.
 - Evaluate Snell’s cone and the n-squared constancy of radiance.
 - Verify boundary compensation, skyfraction and background reflection.
 - Compute beam attenuation as a function of absorption and out-scattering.

- Evaluate in-scattering using the two-stream approximation and stylized canonical problems with reflected benthic-plane contributions.
 - Determine irradiance distributions as a function of ship-shadows and compare to published numerical solutions.
 - Analytically evaluate DIRSIG’s microfacet bidirectional reflectance (and transmittance) distribution functions (BRDFs and BTDFs).
 - Analytically evaluate caustics generated using a sinusoidal interface.
- *Phase Two: Investigate how the user influences the rendering process.*
 - Examine the relationship between image quality and user-defined photon mapping parameters (the maximum number of photon histories to be traced and the maximum search radius to be used when querying the photon map for in-scattered radiance).
 - Compare and contrast the strengths and weaknesses of the Generic and SurfacePM Radiometry Solvers in terms of each solver’s ability to reproduce skyfraction, background reflection and shadows.
- *Phase Three: Illustrate (through realistic simulations) DIRSIG’s ability to model coastal environments.*
 - Generate real and useful shore and coastal imagery that highlights the phenomena investigated in phase one and two.

Collectively, the proposed research described above and in subsequent chapters will serve to validate in-water radiative transfer as modeled by DIRSIG Radiometry Solvers, and using a Monte Carlo photon map to describe the medium’s source function. The research is unique since it evaluates *system* related phenomenon that include complex boundary, medium and object interactions, validated using a combination of analytical solutions, inference, visual demonstration, and comparison with numerical and experimental observations.

A summary of several expected contributions of the proposed research is shown below:

- Generate DIRSIG computed radiance distributions for simplified (water “boxes”, hemispherical droplets) and complex geometries (ship-shadow).
- Illustrate the characteristics of image transfer in participating media.
- Provide proof-of-concept for scalar and time-independent radiance calculations using DIRSIG.

- Compare DIRSIG computed scalar and time-independent radiance distributions to analytical approximations.
- Compare DIRSIG computed scalar and time-independent radiance distributions to experimental (field) and numerical data for scenarios that do not admit an analytic solution.
- Illustrate how user-assigned variables impact performance, including a comparison of DIRSIG’s Generic and SurfacePM Radiometry Solvers.
- Provide first-order estimates for variance and bias reduction for volumetric photon mapped in-scattered radiance solutions.
- Generate real and useful imagery that proves that the *“whole is equal to the sum of the parts”*.

Finally, the proposed research will lay the foundation for future work concerning *nearly* any type of application-specific modeling of coastal ocean water environments, with particular emphasis on image transfer in scattering media, including the ability to predict in-water target signatures, visibility and detectability.

1.2 Outline of Subsequent Chapters

- Chapter 2 introduces the relevant optical properties of absorption, scattering and extinction.
- Chapter 3 broadly develops the radiative transfer equation.
- Chapter 4 handles reflective, transmissive and refractive contributions to sensor reaching radiance that are a function of a simple, planar air-water interface.
- Chapter 5 extends the discussion in Chapter 4 to handle reflective, transmissive and refractive issues related to sinusoidal and capillary waves.
- Chapter 6 continues the discussion of in-water objects and the resulting skyfraction and background reflected contributions to sensor reaching radiance expected for 3D underwater scenarios, using the Generic and SurfacePM Radiometry Solvers.
- Chapter 7 extends the solutions of Chapter 6 to include a two-stream approximate solution to in-scattered radiance, a comparison of absolute radiance to published numerical results from Mobley et al. (1993), and DIRSIG computed irradiance versus cursory field observations.

- Chapter 8 illustrates contributions to sensor reaching radiance that are a function of in-water objects via the specific 3D radiative transfer problem of ship-shadow and self-shading. Understanding the radiance distribution as a function of illumination conditions and adjacency to floating objects is an important three-dimensional radiative transfer problem that cannot be solved using plane-parallel assumptions.
- Chapter 9 provides the theoretical foundations for the radiance reconstruction that occurs during the second pass of the photon mapping process. Of specific interest is the expected bias, variance and convergence of the solution as a function of user-defined variables.
- Finally, Chapter 10 takes the piece-wise validation from all previous chapters and uses visual demonstrations and to illustrate DIRSIG's ability to model radiative transfer in coastal ocean-water *systems*.
- Chapter 11 briefly presents concluding remarks.

Chapter 2

Water

Since water has been chosen as the participating medium of interest, we begin with a brief qualitative and quantitative description of radiative transfer in ocean bodies. As light propagates through this volume it is attenuated by the independent mechanisms of absorption and elastic scattering. Absorption occurs as a result of thermodynamically irreversible mechanisms that convert a photon's radiant energy into thermal, kinetic or chemical energy. In the ocean, the major absorption mechanism is the conversion of photons into thermal energy. Electro-optically, pure water possesses a single transmission window with a peak near 480 *nm* which results in an overall low reflectance and blue color (Duntley 1963). At longer wavelengths pure water is highly absorptive which gives rise to its virtual zero reflectance in the near infrared (Schott 2007). Duntley (1963) indicates that clear ocean water is so selective in its absorption that only a narrow band of blue-green radiation is capable of deeply penetrating the sea “but that this radiation has been detected at depths greater than 600 m”.

The addition of phytoplankton, characterized by chlorophyll, causes additional absorption in the blue and red wavelength regions, causing the water to appear greenish, with a continued near zero reflectance in the near infrared (Schott 2007). Decomposing organic materials release humic and tannic acids, characterized by the concentration of colored dissolved organic materials (CDOM) with strong blue absorption. This shifts the transmission window toward the green and even yellow/brown in the presence of high CDOM concentrations (Schott 2007), (Duntley 1963). In addition to phytoplankton and CDOM, suspended materials (SM) may be present. SM includes inorganic and large-diameter organic materials that non-selectively contribute to absorption (Schott 2007).

Scattering refers to the random process by which a photon's direction of propagation is altered (Duntley 1963). Pure water is considered a poor scatterer. Near the 480 *nm* transmission window approximately 60% of the extinction coefficient (descriptor for beam attenuation) for clear water is a function of scattering, while the remaining 40% is a function of absorption; but in all other wavelength regions “absorption is overwhelmingly predominant” (Duntley 1963).

Pure liquids, including water, are characterized by nearly uniform or isotropic molecular scattering (Hodara 1973). The addition of phytoplankton and suspended sediments causes widespread scattering across the visible and near infrared regions of the electromagnetic spectrum (CDOM is believed to contribute very little to the scattering properties of water) (Schott 2007). The mechanism by which scattering occurs in turbid waters is largely a function of the interaction between photons and biological organisms or inorganic particles that are large compared to the wavelength of light. This results in a *highly forward-peaked scattering phase function* and *low dispersion*. The lack of dispersion means that the scattering magnitude is largely independent of wavelength (Duntley 1963). This is especially true in the forward-peak, but Hodara (1973) suggests that since the backscattered component is influenced by uniform molecular scattering, some spectral dispersion is expected and it typically follows a λ^{-4} law. As such, the scattering properties of water are most readily examined via investigations with “blue” light in clear ocean waters and “green” light in coastal and lake waters (*e.g.* where absorption is at a minimum) (Duntley 1963).

The high degree of scattering at small scattering angles is believed to be a function of both refraction and diffraction. Diffractive scattering is the result of deviations in the light beam’s wavefront as it is diffracted around particles much larger than the radiation’s wavelength (Wells 1973). Conversely, refractive scattering results from the minor refractive index deviations between the water and suspended transparent plankton with an index of refraction nearly equal to that of water (Duntley 1963).

The shape of the forward portion of the scattering function (angles less than 90 degrees from the original beam) remains relatively constant across many water types. Beyond approximately 90 degrees, the backscattered portion can exhibit greater variability (Hodara 1973). The degree of variability in the backscattered characteristics of the scattering indicatrix is a function of variations in molecular and particle-based scattering. Since clear water is dominated by molecular scattering, while turbid waters exhibit both molecular and particle-based scattering, variations in concentration determine the dominant mechanism responsible for the backscattered radiation, and therefore the degree of variability observed.

The remainder of this chapter formally develops the concept of scattered power, the scattering phase function, and the practical implication of global scattering in radiative transfer problems.

2.1 The Scattering Phase Function (Indicatrix) and Absorption Coefficient

Mechanistically, Wells (1973) suggests that it is convenient to think of a particle that scatters light as a small target with a cross sectional area s_p that is somewhat larger than the geometric cross section of the particle itself. Light that falls on this cross sectional

area is scattered, and light that falls near the particle is deflected by diffraction when the presence of the particle disturbs the radiation's propagating wavefront. The extra diffractive power of the particle partly accounts for the increase in the particle's effective cross sectional area of influence beyond the particle's actual geometric size.

For a collection of particles in a volume element, illuminated by an incident beam with cross sectional area A , the degree of obscuration of A will be a function of the total effective area of all the particles in the volume of interest. If the length of the volume element is denoted as Δ , then the total volume is $A\Delta$, chosen such that no particle obscures its neighbor, and that no incident ray is scattered more than once. For n_p particles per cubic length, the total number of particles in the volume element is $n_p A\Delta$, which encompass a total cross sectional area of $s_p n_p A\Delta$ as illustrated in Figure 2.1.




Figure 2.1: Total obscured cross sectional area due to n_p particles (per cubic length) each with cross sectional area s_p within a beam of cross sectional area A and along a length of Δ .

The fraction of power (Φ) scattered due to passage through this volume must be proportional to the fraction of total cross sectional area obscured (Wells 1973) or:

$$\frac{\Delta\Phi}{\Phi} = \frac{s_p n_p A \Delta}{A} = s_p n_p \Delta \quad (2.1)$$

In seawater, there are several species of scattering particles and often n_p is not known independent of s_p , nor is the product $s_p n_p$ known for each separate species. Instead, seawater is described by the global *phenomenological scattering constant* $[\sigma_s]$, more commonly referred to as the *scattering coefficient* (with dimensions of inverse length) over the sum of all particles ($N = p$) (Wells 1973):

$$\sigma_s = \sum_p s_p n_p \quad \left[\text{length}^2 \frac{1}{\text{length}^3} = \frac{1}{\text{length}} \right] \quad (2.2)$$

The loss in power can then be expressed as $\frac{\Delta\Phi}{\Phi} = \sigma_s \Delta$ which is the fraction of power scattered (lost) per unit propagation distance. In the limit, the loss is $\frac{d\Phi}{\Phi} = -\sigma_s(r)dr$, for

an element of path length dr , making the dependence of σ_s on location explicit, and using a minus sign to signify the loss. For propagation through a finite path starting at zero with initial flux Φ_o and through a distance r where the resultant flux is Φ_r , the following is the result of integration for a homogeneous medium where $\sigma_s(r) = \sigma_s$ (Schott 2007):

$$\begin{aligned}
\int_{\Phi_o}^{\Phi_r} \frac{d\Phi}{\Phi} &= \int_o^r -\sigma_s(r) dr \\
&= \ln \Phi \Big|_{\Phi_o}^{\Phi_r} = \int_o^r -\sigma_s(r) dr \\
&= \ln \Phi_r - \ln \Phi_o = \ln \left[\frac{\Phi_r}{\Phi_o} \right] = \int_o^r -\sigma_s(r) dr \\
&= \frac{\Phi_r}{\Phi_o} = \exp \left[\int_o^r -\sigma_s(r) dr \right] = \exp \left[-\sigma_s \int_o^r dr \right] \\
&= e^{-\sigma_s r}
\end{aligned} \tag{2.3}$$

Therefore, the power that remains in the unscattered beam as a function of distance r (along the beam) is equal to $\Phi_r = \Phi_o e^{-\sigma_s r}$.

Using the same logic, incident radiation is absorbed as a function of a *bulk absorption constant* or *absorption coefficient* $[\sigma_a]$. The combined sum of the absorption and scattering coefficients is termed the total extinction coefficient $[\sigma]$ or $\sigma = \sigma_a + \sigma_s$. The unscattered/unabsorbed power that remains in the beam as a function of the distance r is then:

$$\Phi_r = \Phi_o e^{-\sigma r} \tag{2.4}$$

The inverse of the total extinction coefficient or $\frac{1}{\sigma}$ is termed the “mean free path” (Duntley 1963). For *average* ocean waters, Honey (1979) reports that $\frac{1}{\sigma}$ is 10m, or that the beam transmission is nearly 90% over a 1m path.

As shown above, the scattering coefficient can be used to describe the fraction of power that is out-scattered along the propagation distance r . However, it does not indicate the direction into which power is scattered. To determine the probability of being scattered into a given direction, the cross sectional area of the scattering particle s_p can be divided into infinitesimal pieces ds_p , such that each subsection of s_p scatters light into an infinitesimal solid angle $d\Omega$. The differential scattering cross section is then expressed as $ds_p = s_p(\theta, \phi) d\Omega$, where θ is the zenith angle and ϕ is the azimuthal angle through which light is scattered. Of course, the total scattering cross section can be computed from integration of the differential scattering cross section or (Wells 1973):

$$s_p = \int ds_p = \int s_p(\theta, \phi) d\Omega \tag{2.5}$$

In terms of bulk properties, the *volumetric scattering function* $\sigma_s(\theta)$ (also known as the *differential scattering coefficient* or the *angular scattering coefficient*) is the product of each particle species differential scattering cross section and its number density:

$$\sigma_s(\theta) = \sum_p n_p s_p(\theta, \phi) \quad (2.6)$$

$\sigma_s(\theta)$ is interpreted as the power scattered per steradian per unit distance traveled in the beam, and where the dependence on ϕ has been omitted (valid for unpolarized light and randomly oriented particles) (Wells 1973). The quantity σ_s can be derived through integration of the volume scattering function:

$$\begin{aligned} \sigma_s &= \sum_p s_p n_p \\ &= \sum_p n_p \int ds_p \\ &= \sum_p n_p \int s_p(\theta, \phi) d\Omega \\ &= \int \sigma_s(\theta) d\Omega \end{aligned} \quad (2.7)$$

Therefore, radiation traveling within a volume element dV of a participating medium, illuminating a solid angle $d\Omega$ with radiance L , will be subject to both absorption and scattering, and the radiant intensity I_{sca} scattered in the direction β with respect to the incident beam, is given by:

$$dI_{sca} = \sigma_s(\beta) L d\Omega dV \quad (2.8)$$

where $\sigma_s(\beta)$ is the *differential scattering coefficient* or *volume scattering function* in the β direction. The flux (power) scattered in all directions can be found from Eq. (2.9):

$$\Phi_s = \int_{4\pi} dI_{sca} d\Omega' = I d\Omega dV \int_{4\pi} \sigma(\beta) d\Omega' = \sigma_s I d\Omega dV \quad (2.9)$$

where $\sigma_s = \int_{4\pi} \sigma(\beta) d\Omega'$ is the *scattering coefficient*. This suggests that the relative angular distribution of the scattered intensity can be described as:

$$p(\beta) = \frac{4\pi\sigma_s(\beta)}{\sigma_s} \quad (2.10)$$

where $p(\beta)$ is the *scattering phase function* or *scattering indicatrix*, normalized to 4π when integrated over all solid angles:

$$\frac{1}{4\pi} \int_{4\pi} p(\beta) d\Omega' = 1 \quad (2.11)$$

To illustrate, consider an isotropic scattering indicatrix. This implies a uniform scattering indicatrix or $p(\beta) = 1$, since there is equal probability that a photon will be scattered in any direction. Realistic phase functions are often very complex and essentially all numerical methods (and of course, analytical methods) to solve the radiative transfer equation strive to find a suitable approximation or discretization for the scattering phase function.

The terminology used to describe the absorption and scattering from a small volume element can be extended and manipulated in various analogous expressions in order to convey information to the reader, or re-define a scattering problem. For instance, by analogy with the extinction coefficient, the broad beam attenuation coefficient can be defined as:

$$\gamma = \sigma_a + s(> 10^\circ) \text{ where } s(> 10^\circ) = \int_{\frac{1}{6}}^{\pi} \sigma_s(\theta) d\Omega \quad (2.12)$$

which gives the rate of power lost from a cone that is 10 degrees in angular radius [or $\frac{1}{6}$ radians] for small angular scattering problems, and where $e^{-\gamma r}$ is the fraction of light that remains in a 20 degree cone after propagation a distance r along the incident beam. In other words, the broad beam attenuation coefficient attributes small angular scattering events to direct radiance. A similar approach is taken within DIRSIG to handle forward scattered photons.

The solution for the power remaining in the unscattered beam (shown in Eq. (2.3)) is valid for single or multiply scattered radiation, but there is no simple solution that describes the *distribution of the scattered radiation*. Although the scattering phase function describes the probability that an incident photon is scattered into a specific direction, without extrapolation and manipulation, it does not describe the radiance distribution after each photon has been multiply scattered. If Eq. (2.4) describes the power that remains in a beam after losses due to absorption and out-scattering, a full solution for the power in the beam at point r must include gains due to in-scattered (and, if relevant, emitted) radiation. In a large volume with a high probability of scattering, the sensor reaching radiance as a function of in-scattered energy is clearly a global rather than local phenomenon. Mathematically obtaining a solution for such a global quantity is indeed very difficult. Chapter 3 extends Eq. (2.4) to include radiance gains, thereby forming a complete description of the radiance distribution, also known as the radiative transfer equation (RTE).

Chapter 3

The Radiative Transfer Equation

There are two distinct approaches to handling multiple scattering problems. *Analytical theory* (not to be confused with analytical approximations) uses Maxwell or wave equations combined with particle scattering and absorption characteristics in order to arrive at solutions. It is considered mathematically rigorous and includes a theoretical basis for diffraction and interference effects, although in practice only approximate solutions are possible (Ishimaru 1978). Conversely, *radiative transport theory* solutions deal directly with the transport of energy, and therefore do not include diffraction effects. To use *radiative transport theory* the transfer problem must consist of uncorrelated wave fields that allow for the addition of power, rather than requiring the addition of wave fields (Ishimaru 1978). It is this second approach that will be considered in all subsequent discussions.

3.1 Variable & Clarification

A valuable description of the radiative transfer equation (RTE) can be found in Mobley (1994) [and summarized by Goodenough (2007)], Thomas and Stamnes (1999), Zege et al. (1991) and Ishimaru (1978). Of these, the work of Mobley (1994) is probably the most commonly referenced within the oceanic remote sensing community. However, as a matter of preference, the RTE development presented by Thomas and Stamnes (1999) and Shu (1991) was considered most helpful to the author when formulating the following description.

As a general cautionary note, the reader should be warned that the term *specific intensity* with the variable notation I is used in various published descriptions of the RTE. The remote sensing reader will find that *specific intensity* is mathematically consistent with the more commonly used remote sensing variable of *radiance*. In addition, radiometric descriptions of the RTE often use the variable E to describe *energy*, which is in contrast to the use of E to describe *irradiance* in remote sensing applications. In applications

that use E to describe *energy*, the variable F refers to *irradiance*. Keeping this variable clarification in mind, this chapter will broadly develop the RTE.

3.2 Overview of the Radiative Transfer Equation

In a lossless isotropic medium void of interactions that result in absorption, scattering and emission events, the radiance falling on two area elements dA' and dA separated by an arbitrary distance r is constant so that $dL' = dL$ (Schott 2007). This concept can be extended using the transport equation to describe the change in radiance across two area elements as a result of sources and sinks, as shown in Eq. (3.1) where c is a finite speed of propagation, dt is a unit of time, and \vec{k} is a directional vector (Shu 1991).

$$\frac{dL_v}{dt} + c \vec{k} \cdot \nabla L_v = \text{sources} - \text{sinks} \quad (3.1)$$

If the direction vector \vec{k} is expressed as $\mathbf{k}(x, y, z) = k_x \mathbf{i} + k_y \mathbf{j} + k_z \mathbf{k}$, then the directional derivative of the scalar field $L_v(x, y, z)$ in the direction $\mathbf{k}(x, y, z)$ is:

$$\vec{k} \cdot \text{grad } L_v = k_x \frac{dL_v}{dx} + k_y \frac{dL_v}{dy} + k_z \frac{dL_v}{dz} = \vec{k} \cdot \nabla L_v \quad (3.2)$$

Therefore, Eq. (3.1) can be re-written as follows:

$$\frac{dL_v}{dt} + c \left[k_x \frac{dL_v}{dx} + k_y \frac{dL_v}{dy} + k_z \frac{dL_v}{dz} \right] = \text{sources} - \text{sinks} \quad (3.3)$$

Since it is more intuitive to describe the RT equation per unit length rather than per unit time, the equation of transfer can be alternatively described according to Eq. (3.4):

$$\frac{1}{c} \frac{dL_v}{dt} + \vec{k} \cdot \nabla L_v = \text{sources} - \text{sinks} \quad (3.4)$$

Explicit description of any source or sink along a ray's transit will vary depending upon the attributes of the material and the presence or absence of external forms of illumination. For a *medium void of external sources*, the RTE sinks and sources can be expressed as shown in Eq. (3.5) where ρ is mass density per unit volume for the participating medium, j_v is the emissivity per unit mass, κ_v^{abs} is the absorption cross section per unit mass, κ_v^{sca} is the scattering cross section per unit mass and $p_v(\vec{k}, \vec{k}')$ is the scattering phase function or the probability density function that describes scattering from \vec{k}' to \vec{k} , normalized to one when integrated over all solid angles (Shu 1991).

$$\begin{aligned}
\frac{1}{c} \frac{dL_v}{dt} + \vec{k} \cdot \nabla L_v &= \frac{1}{4\pi} \rho j_v + \rho \kappa_v^{sca} \int p_v(\vec{k}, \vec{k}') L_v(\vec{k}') d\Omega' \\
&- \rho \kappa_v^{abs} L_v - \rho \kappa_v^{sca} L_v
\end{aligned} \tag{3.5}$$

In Eq. (3.5) the first two terms are sources, and the latter two are sinks. More specifically, $-\rho \kappa_v^{abs} L_v$ is the radiation removed per unit length of travel due to absorption, $-\rho \kappa_v^{sca} L_v$ is the radiation removed per unit length of travel due to out-scattering, ρj_v is the source emission (per steradian) where $j_v = 4\pi \kappa_v^{abs} B_v(T)$ and $B_v(T)$ is the Planck function, and $-\rho \kappa_v^{sca} \int p_v(\vec{k}, \vec{k}') L_v(\vec{k}') d\Omega'$ is the amount of radiation scattered into the beam. If the radiation travels in a short interval or the average radiance is of interest, then the time dependence can be ignored and the transport equation becomes (Shu 1991):

$$\vec{k} \cdot \nabla L_v + \rho \kappa_v L_v = \rho \left(\frac{j_v}{4\pi} + \kappa_v^{sca} \int p_v(\vec{k}, \vec{k}') L_v(\vec{k}') d\Omega' \right) \tag{3.6}$$

where $\kappa_v = \kappa_v^{abs} + \kappa_v^{sca}$. If the source function is expressed according to Eq. (3.7), then Eq. (3.6) reduces to Eq. (3.8) when the ray-path derivative $\vec{k} \cdot \nabla$ is defined as d/dr .

$$S_v(\vec{k}', x) = \frac{1}{\kappa_v} \left(\frac{j_v}{4\pi} + \kappa_v^{sca} \int p_v(\vec{k}, \vec{k}') L_v(\vec{k}') d\Omega' \right) \tag{3.7}$$

$$\frac{dL_v}{dr} + \rho \kappa_v L_v = \rho \kappa_v S_v \tag{3.8}$$

Since particle density and size may vary from one location to the next, it is convenient to express distance using a nondimensional term called the optical depth. The optical depth is defined as:

$$\tau_v = \int_{r_o}^r \rho \kappa_v dr \tag{3.9}$$

For example, if $\tau_v = 1$ then over a given optical distance the intensity decreases by scattering and absorption to a value equal to $\exp(-1)$ of the original intensity (Ishimaru 1978). Furthermore, the RT equation can be rewritten using an integrating factor (Shu 1991).

$$\frac{d}{dr} (e^{\tau_v} L_v) = e^{\tau_v} \rho \kappa_v S_v \tag{3.10}$$

Since $\rho \kappa_v dr = d\tau_v$, integration from r_o to r (optical depth from 0 to τ_v) results in Eq. (3.11).

$$\begin{aligned}
L_v(\vec{k}, \tau_v)e^{\tau_v} - L_v(\vec{k}, 0) &= \int_0^{\tau_v} S_v e^{t_v} dt_v \\
L_v(\vec{k}, \tau_v) &= L_v(\vec{k}, 0)e^{-\tau_v} + e^{-\tau_v} \int_0^{\tau_v} S_v e^{t_v} dt_v
\end{aligned} \tag{3.11}$$

If an integration variable is introduced such that $\tau'_v = \tau_v - t_v$, then Eq. (3.11) becomes:

$$L_v(\vec{k}, \tau_v) = L_v(\vec{k}, 0)e^{-\tau_v} + \int_0^{\tau_v} S_v e^{-\tau'_v} d\tau'_v \tag{3.12}$$

Eq. (3.12) is the deceptively simple integro-differential equation of radiation transfer. Obtaining a solution is exceptionally difficult because the source term is a global variable. The mathematics become intractable since the scattered radiance distribution is a function of nonlocal processes (Marshak and Davis 2005). Eq. (3.12) cannot be solved analytically without additional assumptions and approximations, or numerically without some mechanism to populate or sample the scattered radiance distribution.

3.3 Overview of Analytical Approximations to Radiative Transfer

According to Zege et al. (1991), there are three major analytical approaches to the solution of the radiative transfer equation. The first approach is to make use of average radiance values over the upwelling and downwelling hemispheres, such as that done in the two-stream approximation. The second is based on the assumption that the properties of the multiply scattered light can be determined using integral parameters of the scattering indicatrix (*e.g.* asymmetry parameters, the ratio of back to forward scattering, etc.) and that the radiance distribution is only weakly determined by the details of the angular structure described by the phase function. This allows use of a simplified approximation to the scattering indicatrix, but which contains the same integral parameters. The third approach handles highly forward-peaked phase functions. For media with phase functions that satisfy this requirement, scattering is estimated using the small angle scattering theory, where it is assumed that the majority of the intensity remains parallel to the incident beam, and then rapidly decreases as a function of the scattering angle.

Ishimaru (1978) divides approximate solutions in a much more intuitive fashion, although still based on the attributes of the scattering media, and fully consistent with the divisions put forth by Zege et al. (1991). For example, scattering can be described depending upon the size of the particles, or the distribution of the particles. For size characterizations, there are two solutions. One solution can be employed when particle sizes are much smaller than the wavelength of illumination, and the other when particle sizes

are much larger than the wavelength of illumination. Particles that are small with respect to the wavelength of illumination have scattering characteristics that are largely independent of the scattering angle, and can be assumed to be described by a nearly isotropic scattering phase function (or at least a very simplified phase function). Conversely, when particles sizes are much larger than the wavelength of illumination, the scattering occurs in small angles around the incident ray, allowing the small angle scattering approximation solution to the RTE.

For density characterizations, the medium can be divided into *tenuous* or *dense* distributions. A tenuous distribution means that single scattering dominates and therefore single scattering approximations to the RTE can be employed since the incident wave encounters few particles during transit between the source and receiver. These approximations are applicable when the volume density, or the ratio of the volume occupied by particles to the total volume of the medium, is smaller than 0.1% (Ishimaru 1978). At the other extreme is a dense medium where diffusion approximations are valid (*e.g.* the radiance distribution is nearly isotropic or at least only weakly dependent on the details of the scattering phase function after undergoing numerous scattering events). This includes volume densities that are much greater than 1%. For media with volume densities between 0.1 and 1%, Ishimaru (1978) reports that the full RTE must be solved. Since water bodies can be extremely variable, it is possible to find volumes that span this entire range, from very clear and shallow coastal waters, to very turbid lakes and ponds.

The details of individual approximate solutions will be discussed as necessary throughout the remaining chapters. Of specific interest will be the small angle scattering approximation (SAA) described by Zege et al. (1991), and the two-stream approximation described in Chapter 7 (and Appendix A.2). The reader is referred to Mobley (1994), Thomas and Stamnes (1999), Zege et al. (1991), Ishimaru (1978) and Shu (1991) for a more in depth discussion of radiative transfer solutions including: (1.) practical considerations, (2.) simplifications of the scattering indicatrix, (3.) the Eddington approximation, (4.) approximations of the specific intensity, (5.) discrete ordinates, (6.) the diffusion approximation (DA) [see Maffione (1998) and Tam and Zardecki (1980)], and finally, (7.) the multicomponent method that combines the SAA and DA [see Zege et al. (1993), Gerstl et al. (1987) and Tam and Zardecki (1980)].

3.4 Overview of Photon Mapping

For a detailed description of photon mapping and its integration within DIRSIG, the reader is referred to related works by Jensen (2001), Mobley (1994) and Goodenough (2007). Briefly, solution of the RTE is broken up into two steps or passes. During the first pass, photons are generated at sources and propagated into the scene using Monte Carlo techniques. The sampling process considers sources within the scene, the medium IOPs and surface interactions (including the transmissive and refractive air-water interface, as

well as interactions with non-transmissive objects within the scene). The scattered photons in the map are associated with an $[x,y,z]$ location and direction of travel as a result of the sampling process. This map represents the scattered light field for the volume of interest.

During the second pass, the “numerical” RTE is formulated by first imagining that the RT equation represents an actual (physical) propagating ray. In a method termed “ray marching”, this ray is broken into segments where the radiance leaving one segment is equal to the radiance entering the next segment. Within an individual segment, the in-scattered radiance is readily determined by querying the photon map. The stored photon flux, in combination with the probability that the presampled flux will be scattered into the propagating ray’s line of sight, allows for an elegant solution of the elusive global source term shown in Eq. (3.12).

3.5 The “*Big Equation*”

Our discussion thus far has focused largely on the manner in which water IOPs influence the radiance distribution within a volume of interest. Although obvious, it is worth mentioning that in addition to being influenced by water IOPs, the direct and scattered radiance from Eq. (3.12) is also a product of all radiometric interactions described by the “*big equation*” (see Schott (2007) for a complete description of sensor reaching radiance). In scenarios with significant contributions from in-water objects, the direct term generally dominates the sensor reaching radiance, while the scattered term redirects this radiance, and supplies a modest backscattered return. The remainder of this work attempts to validate a selective collection of these interactions.

Chapter 4

Simple Boundary Interface Effects

This chapter describes reflective, refractive and transmissive results that are a product of the air-water interface. The majority of the chapter is devoted to methodological considerations, since the background theory is fairly basic. The results are extended in Chapter 5 to include complicated sinusoidal and capillary waves.

4.1 Proposed Methodology - Fresnel Reflectance

Goal: Evaluate the radiance magnitude of ray's reflected from an air-water interface.

Source Function		
Non-Source Function	✓	Model-vs-Model Comparison
Submerged or Floating Objects		Model-vs-Analytical Prediction ✓
Boundary Interface	✓	Model-vs-Observation Comparison
3D Geometry		

4.1.1 Uniform Illumination

The magnitude of the Fresnel reflectance will be evaluated using a hemispherical droplet of water ($n_w \approx 1.34$), above a totally absorbing ground plane, illuminated by a completely diffuse atmosphere (see Figure 4.1). Since Fresnel reflectance is specular, *verification of the reflected sensor reaching radiance for a uniform atmosphere is equivalent to verification of reflection using an infinite number of direct (e.g. “solar” or “laser”) beams*. The totally absorbing ground plane ensures zero magnitude benthic-reflected radiance for the sub-ray transmitted inside the water volume, thereby effectively isolating and measuring the magnitude of the interface-reflected ray. The solution can just as easily be repeated using polygon geometry (e.g. a water “box”), although performing the analysis for a curved surface effectively evaluates *most angular intersections that can conceivably occur*

for a ray and air-water interface (flat or wavy). In the interest of simplicity, the spectral properties of the bulk media will be held constant and examined using a panchromatic band between 0.4 and 0.6 microns.




Figure 4.1: Illustration of scenario to verify Fresnel reflection.

Predicted Results

The reflected ray will obey the law of reflection where the angle of incidence is equal to the angle of reflection. The transmitted ray will be refracted as a function of Snell's law (see Eq. (4.1) where n represents the refractive index of the specific medium, i represents the incident material/ray, and r represents the refractive material/ray).

$$n_i \sin \theta_i = n_r \sin \theta_r \quad (4.1)$$

The actual fate of the transmitted ray is unimportant in this particular scenario, although its direction must be determined in order to solve the Fresnel equation and obtain the magnitude of the reflected and transmitted radiance with respect to the magnitude of the incident radiance (Martin 2004). The Fresnel reflectance for unpolarized illumination and any arbitrary angle of incidence is shown in Eq. (4.2), where θ_i represents the angle of incidence and θ_r represents the angle of refraction. The special case of normal incidence is handled by Eq. (4.3) (Mobley 1994).

$$r_f = \frac{1}{2} \left[\left(\frac{\sin(\theta_i - \theta_r)}{\sin(\theta_i + \theta_r)} \right)^2 + \left(\frac{\tan(\theta_i - \theta_r)}{\tan(\theta_i + \theta_r)} \right)^2 \right] \quad (4.2)$$

$$r_f = \left[\frac{n_r - n_i}{n_r + n_i} \right]^2 = \left[\frac{n_{water} - n_{air}}{n_{water} + n_{air}} \right]^2 \quad \text{normal incidence} \quad (4.3)$$

For a ray propagating from air-to-water, the reflectance can be predicted from the plot shown in Figure 4.2. Mobley (1994) reports that the reflectance is 0.02 to 0.03 for rays with incidence less than 30 degrees, and only climbs to 0.1 for rays that possess incident angles greater than 65 degrees. For a sphere with a 10m-radius [and centered at (0,0,0)], the expected results are shown in Figure 4.3 for the scene attributes listed in Table 4.1, and using Eq. (4.4) where E_h is the total hemispherical downwelling irradiance (uniform or Lambertian) and r_f is the Fresnel reflectance.

$$L_{reflected} = \frac{E_h}{\pi} r_f \quad (4.4)$$




Figure 4.2: Illustration of Fresnel reflection as a function of angle of incidence (in degrees) for an air-incident ray, generated using Eq. (4.2).

Table 4.1: Scene characteristics to test Fresnel reflectance.

Object Geometry	Material Properties	Illumination	Atmosphere	Image Capture
Ground Plane	$r_g = 0$	Uniform Downwelling	Simple Atmosphere	Vis-Panchromatic
10m-Radius	$n_w = 1.34$	$E_h = 1.0$	Transmission = 1.0	0.4-0.6 μm
Water Sphere		SkyFraction = 1.0		

r_g is the ground reflectance, n_w is the refractive index of the water volume, and E_h is the total hemisphere irradiance. Note that “skyfraction” as used in the “simple atmosphere” DIRSIG class, divides the user-defined hemisphere irradiance between skylight and solar contributions. A value less than 1.0 includes solar contributions.

DIRSIG Radiometry Solvers Utilized: Generic, Fresnel, MediumPM.

Expected Fresnel Reflection From 10 m Radius Sphere

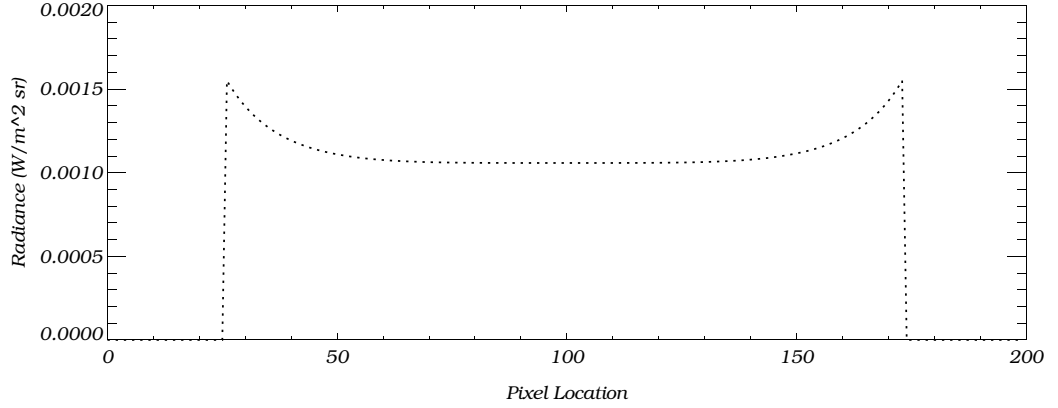


Figure 4.3: Analytically predicted sensor reaching radiance [$W m^{-2} sr^{-1}$] for a 10m-radius hemispherical droplet with refractive index of $n_w = 1.34$.

To determine r_f analytically, the hemispherical interface was handled using the triangle sine law. Referring to Figure 4.4, the generalized angle of incidence is defined as $180^\circ - X$ where X can be determined from the triangle law of sines and based on the geometry of the simulation and the detector.




Figure 4.4: Using the triangle law of sines, the angular variable ‘X’ can be calculated. From this, the incident angle ($180-X$) can be computed.

If the computed angle of incidence is equal to or greater than approximately 45° , the

sensor reaching radiance will be zero. This is because rays reflected at angles greater than 45° eventually impact the totally absorbing ground plane (or a perceptual void within DIRSIG). This is also true for rays with an incidence angle of 90° (and therefore tangent to the hemisphere's surface as illustrated in Figure 4.5).




Figure 4.5: Rays with incident angles of 45° degrees or greater (left), [including those that are tangent to the droplet (right)] generate zero reflected radiance.

4.1.2 Results

The DIRSIG simulated (left) and analytically predicted imagery (right) are shown in Figure 4.6. A transect of the sensor reaching radiance near the midpoint of the imagery is shown in Figure 4.7. It was found that a cut-off angle of 47° (instead of 45°) provided a better fit to the DIRSIG data for this particular geometry (*e.g.* assuming a ray impacts the ground plane or perceptual void beyond an incident angle of 47°). This difference (2°) is considered trivial and perhaps indicative of the construction of the skydome within DIRSIG.

4.1.3 Solar and Skylight

Since the Fresnel reflectance is purely specular, validation using a uniform atmosphere is equivalent to validation of Fresnel reflectance as a function of an infinite number of incident beams (*e.g.* “solar” or “laser”) radiating from individual skyquads. However, the solar disc in DIRSIG does not necessarily occupy a full skyquad. Therefore, computation of specularly reflected solar illumination was examined (note that the solar disc has a diameter of 0.53 degrees in DIRSIG). The proposed scenario for simulation combines diffuse downwelling and solar illumination, setting the skyfraction (within the simple atmosphere DIRSIG class this variable divides the hemisphere irradiance into solar and skylight) to



Figure 4.6: Analytically generated (right) versus DIRSIG generated (left) radiance imagery reflected from a hemispherical droplet of water.

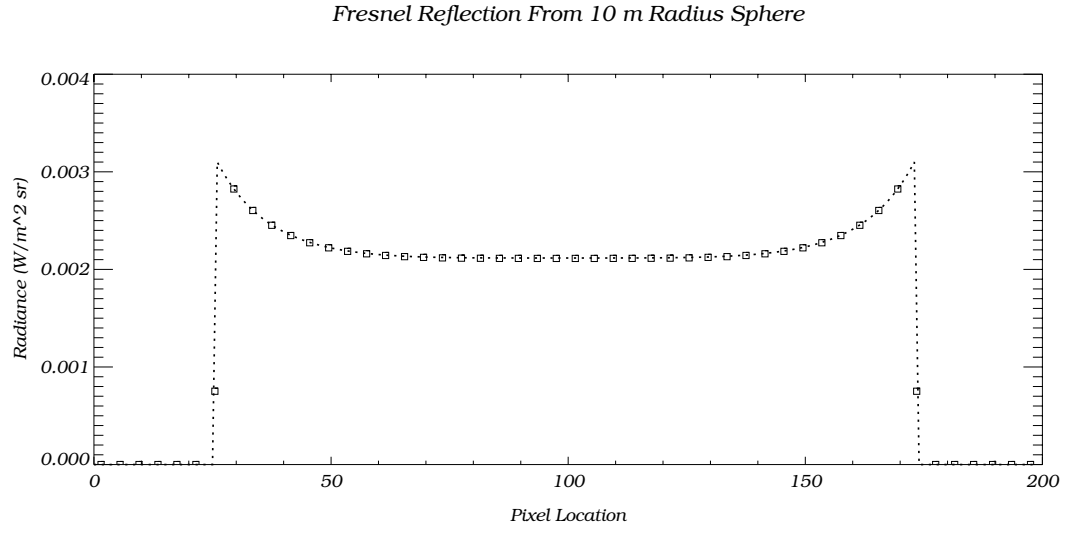


Figure 4.7: Analytically computed (dotted line) versus DIRSIG reflected radiance (squares).

0.5 to collect both skylight and solar contributions for a sun zenith (θ) angle of 0 and 5° (azimuth (ϕ) = 0°). The scenario under investigation is shown in Table 4.2.

Table 4.2: Scene characteristics to test Fresnel Reflectance for combined solar and skylight contributions.

Object Geometry	Material Properties	Illumination	Atmosphere	Image Capture
Ground Plane, Variable Height Water Box/Sphere	$r_g = 0.0$ $n_w = 1.34$	Uniform Downwelling, $E_h = 1.0$, SkyFraction = 0.5 Sun Zenith = 0, Sun Azimuth = 0	Simple Atmosphere, Transmission = 1.0	Spectral 0.4-0.6 μm
		Sun Zenith = 5, Sun Azimuth = 0		

r_g is the ground reflectance, n_w is the refractive index of the water volume, and E_h is the hemisphere irradiance. Note that “skyfraction” as used in the “simple atmosphere” DIRSIG class, divides the user-defined hemisphere irradiance between skylight and solar contributions. A value less than 1.0 includes solar contributions.
Radiometry Solvers Utilized: Generic, Fresnel, MediumPM.

Predicted Results


For a uniform atmosphere, the solar irradiance is computed from the user-defined hemisphere irradiance value. When a ray strikes the Fresnel interface and is specularly reflected toward the solar disc, the entire solid angle of the ray is assumed to be occupied by the sun, and then augmented by a ratio of the solar solid angle versus the “normal” solid angle. The “normal” solid angle is the detector element’s solid angle when using the Fresnel Radiometry Solver (FRS) [or the division of the skydome into skyquads (*e.g.* using μ [number of cosine zenith angles] and ϕ [number of azimuthal angles]) if using the Generic Radiometry Solver (GRS)]. The solid angle of the sun is computed inside the FRS (CDFresnelRadSolver.cpp) according to the following pseudo-code:

$$\begin{aligned}
 \text{solarIrradiance} &= (1.0 - \text{skyfraction}) \text{hemisphereIrradianceVector} \\
 \text{solarRadius} &= \tan(\text{solarHalfAngle}) \\
 \text{solarSolidAngle} &= \pi \text{solarRadius}^2
 \end{aligned} \tag{4.5}$$


The ratio of the solid angle of the sun versus the solid angle of the detector element augments the weight of the reflected radiance by accounting for the fact that the sun does not (likely) occupy the entire solid angle for the ray in question. However, since the Fresnel equation quantifies the magnitude of reflected and transmitted “radiance”, the actual solid angle of the sun cancels out and the calculation proceeds as shown in Eq. (4.6), where E_s is the solar irradiance, Ω_s is the solid angle of the sun, Ω_d is the solid angle of the detector, r_f is the Fresnel reflectance, F_u is the user-defined skyfraction for the uniform atmosphere, and E_h is the user-defined total hemisphere irradiance:

$$\begin{aligned}
L_r &= \frac{E_s}{\Omega_s} r_f \frac{\Omega_s}{\Omega_d} \\
&= \frac{E_s r_f}{\Omega_d} \\
&= \frac{(1.0 - F_u) E_h r_f}{\Omega_d}
\end{aligned}
\tag{4.6}$$

For the scenario listed in Table 4.2 (row one) Figure 4.8 (a.) illustrates the predicted sensor reaching radiance. As the zenith angle of the sun is varied, the location of the specularly reflected solar glint will move appropriately, although its radiance magnitude will stay relatively constant until the incident angle (defined as the angle a “view ray” makes with the volume and the normal to the interface) becomes fairly large and Fresnel reflection increases (see Figure 4.2). The basic geometry of the scene is shown in Figure 4.8 (b.).



(a.) Predicted sensor reaching radiance.



(b.) Geometry configuration.

Figure 4.8: (a.) Predicted sensor reaching radiance [$W m^{-2} sr^{-1}$] for a skyfraction of 0.5, total hemisphere irradiance of 1.0 and the sun at θ (zenith) = ϕ (azimuth) = 0° (the sensor reaching radiance corresponds to a central horizontal slice of a 300 x 300 image). (b.) Geometry of scenario corresponding to Table 4.2 (row one), and the sensor reaching radiance shown above.

4.1.4 Results

The analytical prediction is compared to DIRSIG generated results in Figure 4.9 for sun zenith (θ) = azimuth (ϕ) = 0° . For a planar volume with a depth of 99m, and a platform at 100m from the ground plane, the sun glint occupies 10 x 10 pixels [horizontal x vertical]. For an equivalent sensor-ground geometry, but with a sphere centered at (0,0,0) and having a radius of 99m, the sun glint likewise occupies 10 x 10 pixels. Moreover, the image size of the solar glint matches the analytical prediction and corresponds to the physical size of the sun in DIRSIG (*e.g.* 0.265° on either side of the optic axis of the sensor for a 0.53° diameter DIRSIG sun).




Figure 4.9: Analytically predicted (dotted line) and DIRSIG generated (squares) sensor reaching radiance for a skyfraction of 0.5, total hemisphere irradiance of 1.0 and the sun at $\theta = \phi = 0^\circ$.

However, of greater interest is the change in the physical size of the sun glint when the sun has a non-zero zenith angle, and when reflected from a highly curved surface. If the depth of the planar water volume is reduced to 50m, or even 25m, the sun glint for a solar disc at nadir still occupies 10 x 10 image pixels. Restricting our interest to the 25m water depth, the sun glint for a solar disc that is off-axis is 10 x 9 pixels at solar $\theta = 5^\circ$ and 8° (where the non-symmetrical size is due to sampling since the “lost” pixel changes from the vertical to the horizontal direction when the solar disc is at $\theta = 5^\circ$ and $\phi = 90^\circ$). This is illustrated in Figure 4.10 by the dotted line and asterisks.

However, the same cannot be said if the radius (and/or sun zenith angle) is varied for a spherical droplet. Rather, the spherical interface reduces the size of the sun glint, and

Fresnel Reflected Sensor Reaching Radiance

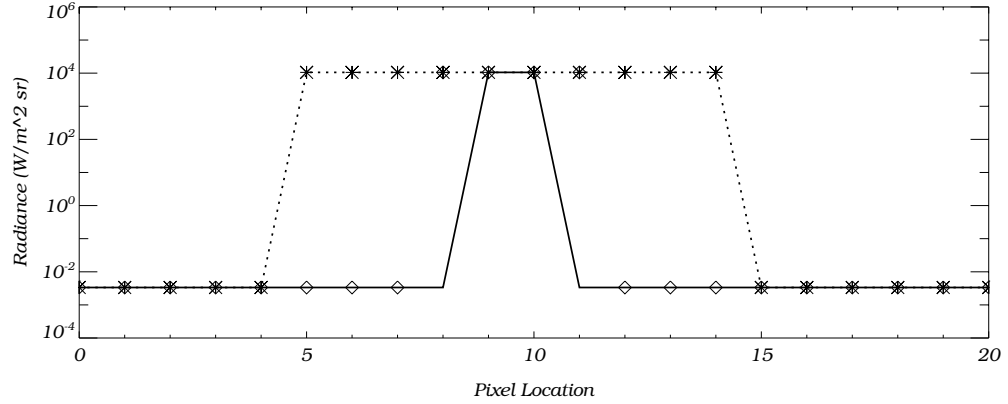


Figure 4.10: Change in size of sun glint when reflected from a planar surface (dotted line = $\theta = 0^\circ$, asterisks = $\theta = 5^\circ$) and spherical interface (solid line = $\theta = 0^\circ$, diamonds = $\theta = 5^\circ$) for a 25m deep water volume and a 25m radius sphere.

the amount of compression is a function of the degree of curvature of the interface. To illustrate, consider $\theta = 5^\circ$ or 8° with a 25m water depth, which produces a sun glint that occupies 10×9 pixels for a planar interface, but 4×3 pixels ($\theta = 0^\circ$) and 2×1 pixels ($\theta = 5^\circ$) for the curved interface (see Figure 4.10).

This leads to two conclusions. The first is that the curved surface reduces the size of the solar glint, and that the reduction is a function of the curvature. This is illustrated in 4.11. Defining the second derivative of the surface as the rate at which the surface normal changes with respect to the viewer, then it can be shown that the higher the second derivative, the larger the compression (Fleming et al. 2004). This means that a feature (the solar disc) that occupies a particular size in a planar reflector, will actually be compressed in a curved reflector, and the amount of compression will be a function of the curvature of the interface (Fleming et al. 2004). Therefore, decreasing the sphere radius will reduce the size of the sun glint in the image as illustrated above.

The second conclusion is that the size of the solar glint is most reduced for a nadir sun. The geometry is shown in Figure 4.12, and the results were previously shown in Figure 4.10. The solid line and diamonds illustrate reflections from a spherical interface with a radius of 25m, while the dotted line and asterisks show reflections from a planar interface with a depth of 25m. Comparison of the dotted line with the solid line shows the size of




Figure 4.11: As the curvature of the reflective surface increases, so does the amount of compression. That is, the “object” is of constant size, but its image varies with surface curvature. Image adapted from Fleming et al. (2004).

the sun glint for a nadir solar disc, while the triangles and asterisks show the size of the sun glint for a solar disc at $\theta = 5^\circ$. In all cases, the spherical interface has reduced the size of the glint when compared to the planar interface for a sun at the same zenith angle, but the discrepancy is greatest for the nadir sun.




Figure 4.12: For a sun at $\theta > 0^\circ$, the image of the solar disc is larger than that that obtained for θ equal to 0° . This is illustrated by the white vertical line which depicts the size of the image of the solar disc when the sun is at nadir (left), and off-axis (right). For comparative purposes, the image size of the solar disc when on-axis (smaller white vertical line in image at right) is shown alongside the expected size when off-axis (larger white vertical in line).




Figure 4.13: Analytical (a.) and DIRSIG [(b.) - (e.)] generated imagery depicting the Fresnel reflected solar glint. The top row [(a.) - (c.)] depicts the Fresnel reflected radiance from solar and skylight contributions from a water box [(a.) and (b.)] and sphere (c.) when sun $\theta = \phi = 0^\circ$ (top row). The second row depicts DIRSIG generated imagery for Fresnel reflected radiance from solar and skylight contributions from a box (d.) and sphere (e.) when sun $\theta = 5^\circ$ and $\phi = 0^\circ$. Since the sun glint has a much higher radiance return than the rest of the interface, the scaled intensity images appear black and white or to have zero and non-zero sensor reaching radiance, but Figure 4.9 shows that this is not the case. Note that the size of the sun glint is smaller when reflected from the curved droplet of water as compared to the planar water interface.

4.2 Proposed Methodology - Fresnel Transmittance & Attenuation

Goal: Evaluate the radiance magnitude of ray's transmitted through an air-water interface. The solution also validates the n-squared constancy of radiance, attenuation due to absorption and out-scattering during transit through the scattering medium, Snell's Cone, and the interaction between the Fresnel Radiometry Solver (FRS), the MediumPM Radiometry Solver (MRS) (non-source function effects) and the Generic Radiometry Solver (GRS).

Source Function			
Non-Source Function	✓	Model-vs-Model Comparison	
Submerged or Floating Objects		Model-vs-Analytical Prediction	✓
Boundary Interface	✓	Model-vs-Observation Comparison	
3D Geometry			

The magnitude of the transmitted ray as a function of Fresnel transmittance is more difficult to isolate than the reflected component. For a *submerged* camera scenario observing the water-air interface, and oriented so that the focal plane is parallel to the water-air boundary, the sensor reaching radiance for a totally absorbing ground plane is a function of Fresnel transmittance, absorption and out-scattering in the medium (in the absence of photon mapping), the n-squared constancy of radiance, and total internal reflection beyond Snell's window.

From a *nadir* viewing geometry, the sensor reaching radiance is a function of the Fresnel transmittance at the air-water interface, absorption and out-scattering in the medium (in the absence of photon mapping), reflection from an in-water object, and re-transmittance through the water-air interface. Furthermore, this nadir viewing geometry will include a contribution due to Fresnel reflectance from the air-water boundary.

The validation will take place in two steps, beginning with the submerged camera scenario, and progressing to the more complicated nadir-viewing camera scenario where reflection from the water-ground plane will be modeled as a completely Lambertian bidirectional reflectance distribution function (Ward BRDF) and computed using GRS.

4.2.1 Fresnel Transmittance & Snell's Cone

Figure 4.14 illustrates the submerged camera scenario. The illustration demonstrates (1.) the n-squared constancy of radiance or the change in radiance for a ray traversing media with different refractive indices (as a function of the change in solid angle), and (2.) total internal reflection beyond Snell's window. The Fresnel Radiometry Solver handles the transmission and refraction at the water-air interface, the MediumPM Radiometry Solver handles attenuation of the ray while inside the water volume, and the illumination is the

product of a completely diffuse downwelling atmosphere (chosen for ease in analytical calculation).




Figure 4.14: Illustration of scenario to verify Fresnel transmittance, Snell's Cone, and the n -squared constancy of radiance.

Predicted Results

From Snell's law (Eq. (4.1)), the angle at which a ray is internally reflected can be determined as a function of the refractive indices of the materials that straddle the interface. For the submerged-camera scenario described in Table 4.3, this will occur beyond $\approx 48.3^\circ$ for water-incident rays. For a 200×200 pixel image generated using a $5000 \mu m$ focal length (and pixel width = $100 \mu m$) total internal reflection will occur beyond ≈ 56 pixels on either side of the optical axis [*e.g.* pixel locations 44 and 156 for a horizontal (or vertical) transect of the image]. This is apparent in Figure 4.15. For non-zero medium IOPs the function shown in Figure 4.15 will be modified by attenuation (and in-scattering if volumetric photon mapping is being used). For a fixed physical depth of the sensor, the optical depth will vary along perpendicular and slanted paths between the sensor and the water-air interface. This modifies the sensor reaching radiance as illustrated in Figure 4.16 for fixed coefficients of $\sigma_s = 0.09$ and $\sigma_a = 0.05$.

The analytical prediction shown in Figures 4.15 and 4.16 is based on the computation of sensor reaching radiance shown in Eq. (4.7), where $L_{transmitted}$ is the sensor reaching radiance as a function of transmission through the water volume, E_h is the hemisphere irradiance, T_f is the Fresnel transmittance equal to $1.0 - r_f$ (where r_f is the Fresnel reflectance shown in Eq. (4.2) and (4.3)), n_w is the refractive index of the water volume, z is the water depth or distance between the platform and water-air interface, σ is the extinction coefficient, and θ_i is the angle of incidence at which rays, generated at detector elements, impact the water-air interface.

Table 4.3: Fresnel transmittance for a submerged camera.

Object Geometry	Material Properties	Illumination	Atmosphere	Image Capture
Ground Plane, 10m-Deep Volume	$r_g = 0$ $n_w = 1.34$	Uniform Downwelling, $E_h = 1.0$, SkyFraction = 1.0	Simple Atmosphere, Transmission = 1.0	Vis-Panchromatic, 0.4-0.6 μm Submerged

r_g is the ground reflectance, n_w is the refractive index of the water volume, and E_h is the hemisphere irradiance. Note that “skyfraction” as used in the “simple atmosphere” DIRSIG class, divides the user-defined hemisphere irradiance between skylight and solar contributions.

Radiometry Solvers Utilized: Generic, Fresnel, MediumPM.

$$L_{transmitted} = \frac{E_h}{\pi} T_f n_w^2 \exp \left[-\sigma \left(\frac{z}{\cos(\theta_i)} \right) \right] \quad (4.7)$$

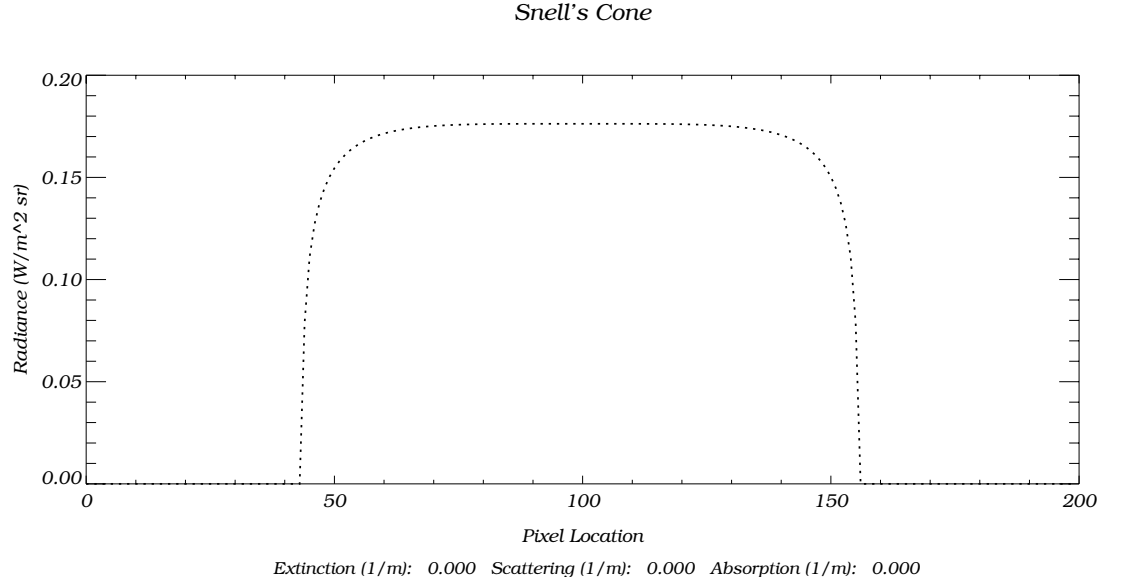


Figure 4.15: Geometric prediction of sensor reaching radiance with total internal reflection as a consequence of the critical angle for a submerged camera and $n_w = 1.34$.




Figure 4.16: The general function shown in Figure 4.15 is modified by the optical path along perpendicular and slanted paths for $\sigma_s = 0.09$ and $\sigma_a = 0.05$.

4.2.2 Results

The Snell window is the valid consequence of total internal reflection when an incident ray strikes the air-water interface at an angle equal to or greater than the critical angle. The DIRSIG radiance results as a function of constant, uniform downwelling illumination and $n_{air} = 1.0$ and $n_{water} = 1.34$, as compared to the predicted direct radiance ($\sigma_s = \sigma_a = 0$) is shown in Figure 4.17. The results, which are truly a measure of refraction, are in excellent agreement.




Figure 4.17: Radiance comparison for DIRSIG (squares) versus the predicted direct radiance solution (solid line) with zero extinction.

The successful comparison of the refractive effects as shown in Figure 4.17 were extended to include an evaluation of attenuation as a function of out-scattering and absorption. A simple analytical code was developed using Eq. (4.7), capable of calculating extinction along a given ray's path within the water volume, including compensation for Fresnel interface effects and the n-squared constancy of radiance. The successful comparisons are shown in Figure 4.18 for two coefficient combinations. This trio of comparisons (shown in Figures 4.17 and 4.18) nicely validates DIRSIG's ability to model Fresnel events, including Snell's cone, and the n-squared constancy of radiance, which are extremely pertinent to submerged sensor configurations. The following section extends this validation to include a nadir sensing platform with reflected ground plane contributions.





Figure 4.18: A comparison of the analytical prediction (dotted line) of direct reduced radiance and DIRSIG computed results (squares). See plots for coefficient information.

4.2.3 Proposed Methodology - Fresnel Transmittance, MediumPM & the Generic Radiometry Solver


The second scenario is much more complicated since it involves at minimum, three user-assigned Radiometry Solvers (excluding automatic DIRSIG architecture, atmosphere solvers, solar solvers, etc.), and the modeling of no less than eight possible radiance reducing interactions, enumerated below. The following discussion accompanies Figure 4.19. For the nadir-view camera, the ray will traverse the air-water interface and be both reflected and transmitted [(1) and (2)]. The reflected component will strike the solar disc or a sky quad, and contribute to sensor reaching radiance. The transmitted ray will be attenuated by absorption [(3)] and out-scattering [(4)] which is handled by the MediumPM Radiometry Solver, before striking the Lambertian ground plane which will be assigned the GenericRadSolver. For a material assigned GRS, the radiance is computed by evaluating the material's BRDF, and appropriately (in terms of importance) sampling possible scene sources. These importance samples will pass back through the water volume, undergoing absorption and out-scattering, followed by Fresnel reflectance and transmittance at the air-water interface [(5), (6), (7) and (8)], before finally sampling the downwelling atmosphere or solar disc.

Analytically, this process is significantly simplified by using a hemispherical droplet of water, assigning the ground plane a completely diffuse Ward BRDF, and using a uniform atmosphere with constant radiance per solid angle (or skyquad), thereby negating the need to numerically compute variable radiance from each solid angle. Furthermore, the analytical sensor reaching radiance will be computed only for the central position of the droplet (or the sphere's origin), since each secondary ray will travel exactly the same optical path (which geometrically corresponds to the radius of the water droplet) and intersect the water-air interface with normal incidence [see Figure 4.19 (b.)]. Conversely, prediction of the sensor reaching radiance at all other positions within the droplet requires a numerical computation, or at least additional analytical assumptions, which will not be pursued here.


Predicted Results

Using the parameters described in Table 4.4, the predicted sensor reaching radiance is shown in Figure 4.20, as calculated using Eq. (4.8), where E_h is the downwelling irradiance, r_f is the Fresnel reflection, T_f is the Fresnel transmittance, σ is the extinction coefficient, r_g is the ground reflectance and R is the radius of the sphere.

$$\begin{aligned} L &= L_{transmitted} + L_{reflected} \\ L_{reflected} &= \frac{E_h}{\pi} r_f \end{aligned}$$



(a.) The sensor will collect both reflected and transmitted contributions to sensor reaching radiance.



(b.) Illustration of importance samples utilized by GRS.

Figure 4.19: Illustration of scenario to verify Fresnel transmittance. (a.) Illustrates the interactions that are involved for a nadir-viewing sensor. For this scenario, the transmitted portion cannot be separated from the reflected radiance. (b.) Illustrates importance samples utilized by the Generic Radiometry Solver (GRS). The analytical model seeks to solve the radiance at the center of the hemisphere, where each importance sample travels the radius of the droplet, and impacts the water-air interface with normal incidence. A solution for all other points within the droplet would require a numerical model, or at least additional analytical approximations, which has not been attempted here.

$$L_{transmitted} = E_h T_f^2 \frac{r_g}{\pi} \exp[-\sigma R^2] \quad (4.8)$$

The plot shown in Figure 4.20 indicates that DIRSIG's results should match the analytical prediction in the very center of the droplet, and approach (but be slightly less than), the analytical approximation at the periphery of the image. The sensor reaching radiance for edge pixels corresponds to ground pixels beyond the actual boundary of the water volume, which have a reduced skyfraction and a consequent object fraction that consists of the transmissive water volume. Therefore, when Generic samples the downwelling hemisphere,

some rays will traverse the droplet before reaching the skydome, which will decrease the overall magnitude of the sensor reaching radiance. Since the analytical prediction does not account for attenuation due to the object fraction, the analytical results will be slightly higher for edge pixels (and instead represent the sensor reaching radiance for pixels far removed from the water droplet).

Table 4.4: Fresnel transmittance for a nadir camera.

Object Geometry	Material Properties	Illumination	Atmosphere	Image Capture
Ground Plane	$r_g = 0.2$	Uniform Downwelling	Simple Atmosphere	Vis-Panchromatic
10m-Radius	$n_w = 1.34$	$E_h = 1.0$	Transmission = 1.0	0.4-0.6 μm
Sphere		SkyFraction = 1.0		Nadir

r_g is the ground reflectance, n_w is the refractive index of the water volume, and E_h is the hemisphere irradiance. Note that “skyfraction” as used in the “simple atmosphere” DIRSIG class, divides the user-defined hemisphere irradiance between skylight and solar contributions. Radiometry Solvers Utilized: Generic, Fresnel, MediumPM.

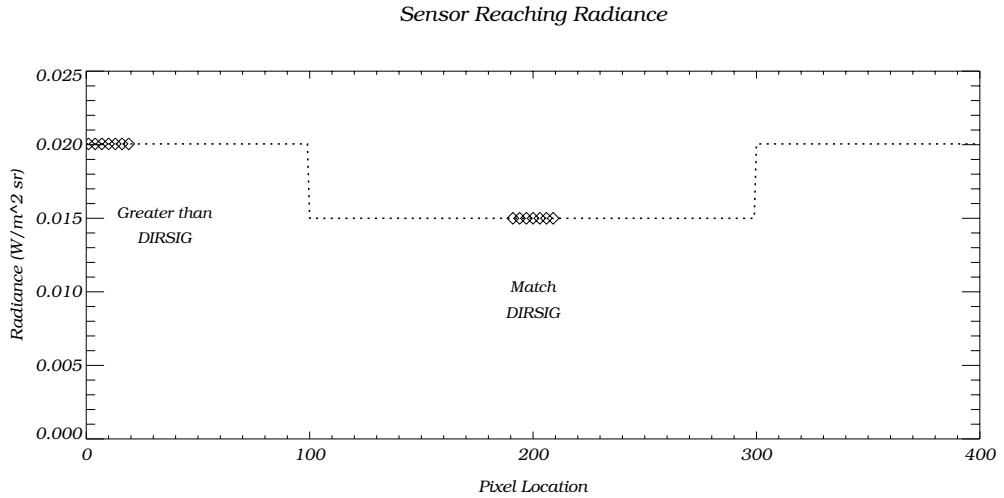


Figure 4.20: The analytically predicted sensor reaching radiance should match DIRSIG generated results at the very central portion of the image, but be greater at the image edges. The numerical results will not be “flat” as might be inferred from the analytical prediction. To clarify, the analytical result is only being computed at the central location (pixel 200) and the extreme edges (pixel values of 0 and 400) as denoted by diamonds. The dotted line connects these predicted results to illustrate the edge of the hemispherical water volume as shown by the “steps” in the plot.

4.2.4 Results

Using the parameters described in Table 4.4, Figure 4.21 shows the analytically predicted and DIRSIG generated radiance for the combined Fresnel, MediumPM and Generic Radiometry Solvers. As formerly noted, the DIRSIG generated radiance should match at the center of the water droplet, and be less than the analytical approximation for pixels at the edge of the image (and outside of the water droplet). Again, the results are in excellent agreement. Note that the decrease in sensor reaching radiance at the edge of the sphere is a function of refraction off of the curved interface. The slight increase in sensor reaching radiance near the edge (and interior) of the sphere is a function of attenuation. The contributions due to refraction and attenuation are best interpreted by inspection of Figure 4.22 which shows a hemispherical droplet of water as a function of refraction only (left), and refraction plus attenuation (right).




Figure 4.21: Analytically predicted (dotted line) versus DIRSIG generated (squares) sensor reaching radiance for $\sigma_s = \sigma_a = 0.01$.

Overall, the results thus far indicate that DIRSIG successfully models the magnitude of direct radiance that is both reflected and transmitted after passage through the air-water interface. The next section evaluates the apparent shift in an object's magnification and displacement due to refraction at this boundary, thereby completing the validation of events that can be expected when modeling a simple boundary interface.




Figure 4.22: DIRSIG imagery for refractive only (left) and refraction plus attenuation $\sigma_s = \sigma_a = 0.01$ (right)).

4.3 Proposed Methodology - Refraction

Goal: Evaluate the refractive induced change in displacement and magnification of objects submerged in a water volume.

Source Function			
Non-Source Function	✓	Model-vs-Model Comparison	
Submerged or Floating Objects	✓	Model-vs-Analytical Prediction	✓
Boundary Interface	✓	Model-vs-Observation Comparison	
3D Geometry	✓		

Refraction at the air-water interface is an important phenomenon since it controls the directional path taken by a ray during transit within the water volume. It is important to validate that the ray not only carries the correct radiance magnitude, but also that it travels in a direction that satisfies Snell's law.

An analytically viable evaluation of refraction can be conducted using simple geometry consisting of a polygon volume (*e.g.* rectangle) of water, between a nadir-viewing sensor and a checkerboard ground plane (Figure 4.23 illustrates a side-view of the simulation). This case study scenario was chosen for validation purposes since the simulated scene, although simplistic, allows the user to easily evaluate the two main consequences of refraction for submerged objects: object displacement and magnification. When applied to the checkerboard ground plane, the refractive displacement and magnification of the *checks* can be easily evaluated for comparative purposes. In addition, the fate of a ray that is (a.) refracted at the horizontal interface, (b.) refracted at the horizontal interface before impacting the vertical interface and suffering total internal reflection, and (c.) refracted at the horizontal and vertical interface can collectively be assessed [see Figure 4.23 (left)]. Furthermore, these interactions are similar to those expected in more complicated geometries, such as refraction through a wavy interface, as illustrated in Figure 4.23 (right). As such, validation of refraction through more complicated interfaces can be made via deduction and inference.

4.3.1 Predicted Results

The logic behind the analytical computation of refraction for this scenario is described in Figure 4.24. For the scene in question, the platform is treated as a pinhole camera, and a ray is trigonometrically cast out from each pixel's midpoint. The propagating ray will have an angle of incidence with all horizontal (and possibly vertical) interfaces within the scene that can be readily computed. Based upon the scene geometry, the initial hit for a ray will consist of either (1.) the horizontal air/water interface, or (2.) the air/ground interface.




Figure 4.23: Illustration of scenario to verify refraction. The Figure at left shows a side-view of the simulation used to verify refraction. Three rays are illustrated: (a.) refraction at the horizontal interface, (b.) refraction at the horizontal interface before impacting the vertical interface and suffering total internal reflection, and (c.) refraction at the horizontal and vertical interface. The Figure at right shows the same ray fates, illustrating that validation of a simpler geometry (water polygon with checkerboard ground plane) can be extended to validate more complex interactions on wavy sea surfaces.

If the ray strikes the air/ground interface and remains outside the waterbox at all times along its descent into the scene, then the ray's (x,y) hit location can be trigonometrically calculated, and the pixel can be assigned an intensity to produce a checkerboard appearance. The mechanism by which the checkerboard pattern was created was borrowed from the DIRSIG procedural CDGroundPlane class. This class allows the user to specify a non-facetized piece of planar geometry, with or without checks. If checks are assigned and a propagating ray intersects the procedural ground plane, then the CDGroundPlane::intersect() function determines which of the two checkerboard materials (*e.g.* white check or black check) the propagating ray will intersect using a Boolean function, described by the pseudo code shown by Eq. (4.9), where *floor* returns the closest integer less than or equal to its argument and *fmod* returns the floating point remainder of a division. From the material ID, a hitList is created so that appropriate material properties and radiometry solvers can be called.

$$\text{checkTest} = (\text{bool}) \left(\text{fmod} \left[\text{floor} \left(\frac{x}{\text{checkWidth}} \right) + \text{floor} \left(\frac{y}{\text{checkWidth}} \right) \right], 2 \right) \quad (4.9)$$

The equivalent IDL interpretation of the checkTest function is shown in pseudo code by

Eq. (4.10). The alternative to borrowing this checkTest function would be to create a facetized surface (in DIRSIG) or a binary matrix (in IDL) to describe the alternating checks, which would be extremely time consuming. Instead of updating a material list as is done in DIRSIG, a *material* variable will be used in IDL to assign a pixel value (*e.g.* if material = 0 then pixel = 0 else pixel = 255).

$$\text{material} = ([\text{floor}(\frac{x}{\text{width}}) + \text{floor}(\frac{y}{\text{width}})] \bmod 2) \quad (4.10)$$

If the propagating ray strikes the top of the water box, the angle of refraction is calculated using Snell's Law (Eq. (4.1)), and the ray is propagated to the ground plane. If the ray strikes the air/water interface near the limits of the waterbox, the ray may intersect the vertical boundary of the waterbox on its way toward the ground plane. Should this occur, the ray's incidence on the vertical water/air interface is evaluated. If the ray is less than the critical angle, the ray is refracted again as it exits the water box and before it eventually impacts the ground plane. If the incident angle is at the critical angle, the ray is propagated along the boundary of the waterbox toward the ground plane. Finally, if the angle of incidence is greater than the critical angle, then the ray is totally internally reflected back inside the water box, before hitting the ground plane. For each impact location, the checkerboard is reformed using Eq. (4.10).

The frequency of the checks on the ground plane will be both displaced and altered in size as a function of the water depth for $n_w = 1.34$. For square checks (width = 0.3m), this is evident from the results predicted in Figure 4.25 and Figure 4.26 for a 4.9m and 9.9m deep water volume, respectively. The predicted results will be compared to a binary DIRSIG image (*e.g.* absolute radiance will not be evaluated).

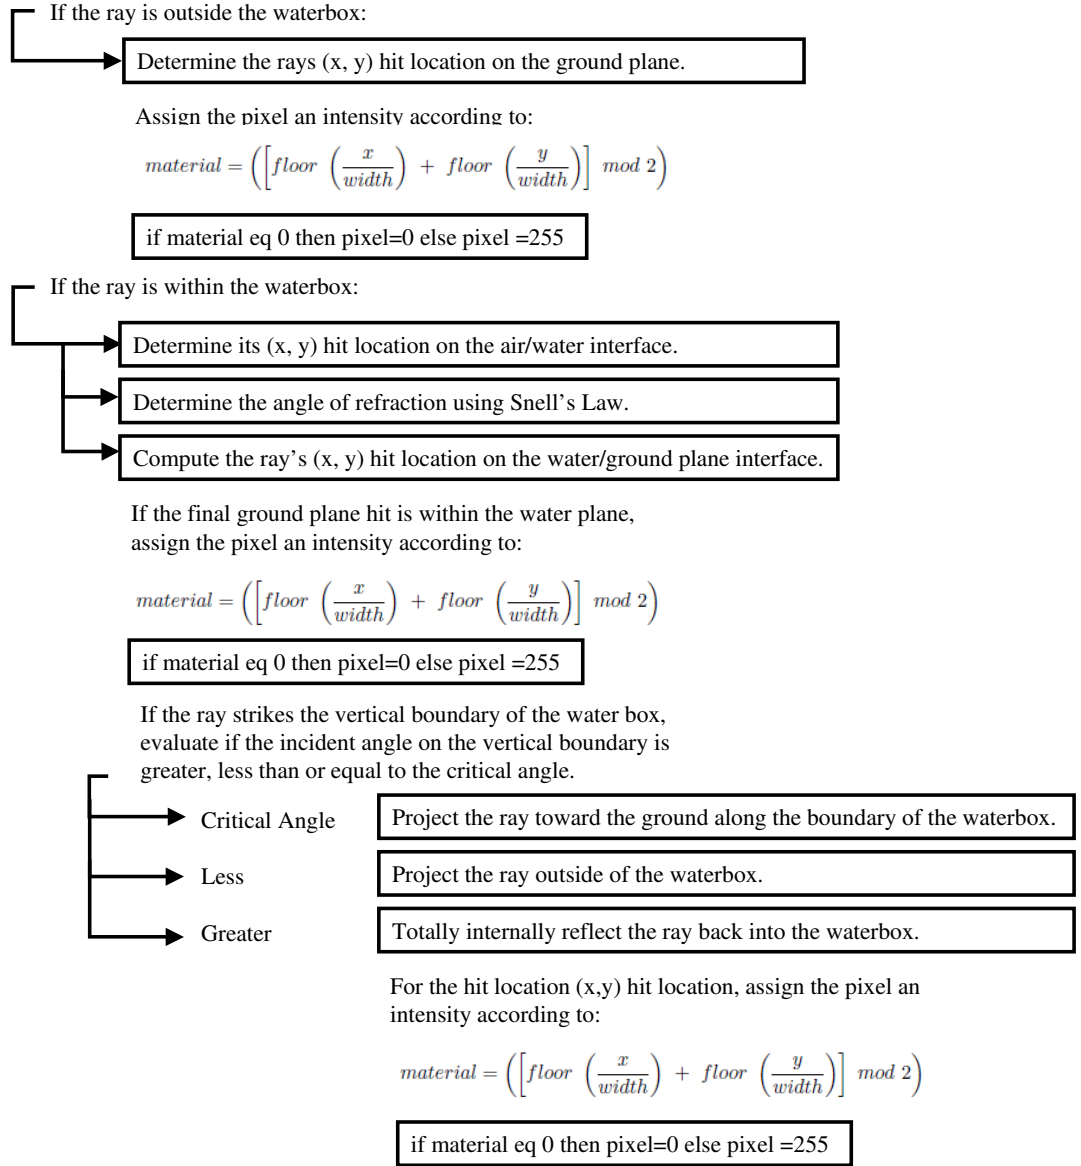

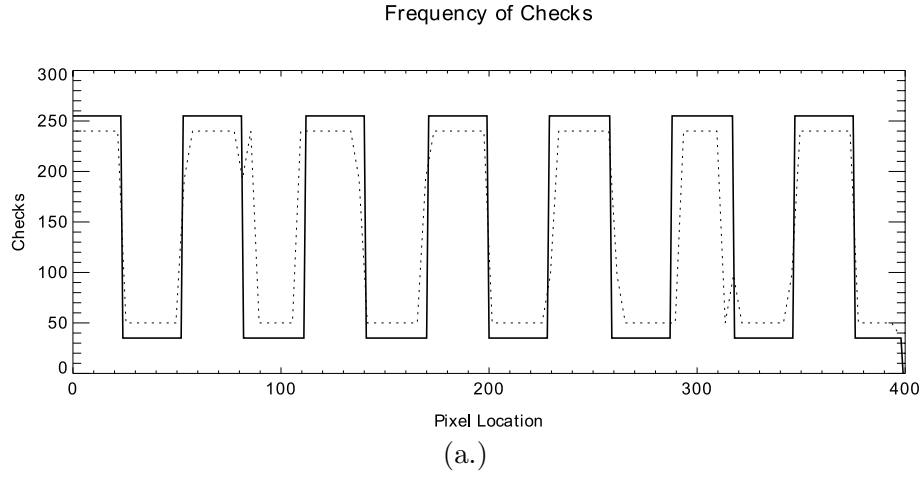



Figure 4.24: Flowchart summarizing computation of checkerboard pattern following refraction.



(b.)



(c.)

Figure 4.25: (a.) Displacement and change in size for 0.3m wide square white and black checks viewed through a water volume. The solid line shows the frequency of white and black checks on the checkerboard ground plane, while the dotted line shows the frequency of white and black checks on the checkerboard ground plane as viewed through a water volume. Note that the magnitude of reflected radiance is not being compared, only the frequency of oscillation in the magnitude. (b.) and (c.) Show the image of the checkerboard ground plane as seen with and without an intervening water volume (4.9m in depth), respectively.






Figure 4.26: Repeat of Figure 4.25 but for a 9.9m water volume.

4.3.2 Results

The results depicted in Figure 4.25 and 4.26 for a 4.9m and 9.9m deep water volume, respectively, were compared to DIRSIG generated results. Since the frequency, size and displacement of the checks is under examination, and not the absolute radiance, the DIRSIG simulations were converted to binary images. Figure 4.27 depicts the oscillation of the checks in the absence of water (analytical (solid line), DIRSIG (diamonds), with pixel values between 300 and 350) and the alteration of the checks due to inclusion of a 4.9m deep water volume (analytical (solid line), DIRSIG (asterisk), with pixel values between 50 and 240). The DIRSIG and analytical imagery are shown in Figure 4.28. The comparable plot for a 9.9m deep water volume is depicted in Figure 4.29, and the DIRSIG and analytically generated imagery are shown in Figure 4.30. In both cases, the DIRSIG modeled results are in excellent agreement with those predicted analytically. A comparison of the check size in the presence and absence of the water volume clearly shows the refractive induced magnification of in-water objects, most notably for the central checks near pixel location 200. In addition, check size and frequency closer to the edge of the water cube (near pixels 100 and 300) nicely illustrate the combined effects of double refraction at both the horizontal and vertical air-water interfaces. Overall, the DIRSIG numerical computation matches both intuition, and the analytical prediction of check magnification and displacement as a function of passage through the air-water interface.




Figure 4.27: Approximate central transect of imagery to illustrate check frequency in the absence of water (analytical (solid line), DIRSIG (diamonds), with pixel values between 300 and 350) and the alteration of the checks due to the inclusion of a 4.9m deep water volume (analytical (solid line), DIRSIG (asterisk), with pixel values between 50 and 240).




Figure 4.28: DIRSIG (right) and analytically generated (left) imagery for a 4.9m deep water volume.




Figure 4.29: Approximate central transect of imagery to illustrate check frequency in the absence of water (analytical (solid line), DIRSIG (diamonds), with pixel values between 300 and 350) and the alteration of the checks due to the inclusion of a 9.9m deep water volume (analytical (solid line), DIRSIG (asterisk), with pixel values between 50 and 240).




Figure 4.30: DIRSIG (right) and analytically generated (left) imagery for a 9.9m deep water volume.

4.4 Transition From Simple Interface Problems to More Complicated Boundaries

This chapter focused on validation of the direct reduced radiance as a function of passage through a simple air-water interface. Aspects of reflection, transmission, refraction, and the n-squared constancy of radiance were evaluated, including total internal reflection and Snell's cone. The following chapter considers the air-water interface when roughened by the addition of capillary and sinusoidal-like gravity waves.

Chapter 5

Waves & Caustics

Goal: Evaluate reflection and transmission through a a wind-roughened air-water interface.

Source Function			
Non-Source Function	✓	Model-vs-Model Comparison	✓
Submerged or Floating Objects		Model-vs-Analytical Prediction	✓
Boundary Interface	✓	Model-vs-Observation Comparison	✓
3D Geometry	✓		

This chapter continues the discussion of contributions to sensor reaching radiance that are a function of the air-water interface. However, instead of a simple planar or spherical interface, gravity and capillary waves will be considered. DIRSIG has the *intrinsic* ability to model sinusoidal gravity and probabilistic capillary waves (although any arbitrary air-water interface can be realized by a user willing to generate facetized geometry). The sinusoidal gravity waves are the product of procedural geometry described in the simulation ODB file for a given wave period, phase shift and amplitude. Conversely, the capillary waves are the product of a microfacet BRDF/BTDF model that predicts Fresnel reflection and transmission based on sub-pixel and purely specular microfacets (Ashikhmin et al. 2000). The following discussion will examine Fresnel reflection and transmission as a product of both sinusoidal and capillary sub-pixel wave models.

5.1 Overview

A planar interface gives rise to a mirror-like reflection of the solar disc located at the *horizontal specular point* (Cox and Munk 1956). Since the sea surface is rarely (never) planar, sun glitter is evident. Glints are the result of solar rays being reflected toward the observer from a multitude of facets with varying inclination. The counterpart of the reflected ray is a transmitted and refracted ray whose magnitude and direction are a function of the

inclination of the interface. It is the combination of many transmitted and refracted rays which result in the commonly observed subsurface caustics. Since these interface events are responsible for a multitude of radiative transfer phenomena, their validation is vital to nearly all water-simulation goals. Although multiple interface and time-dependent events are possible (and inevitable) in real coastal ocean water environments, the remainder of this discussion will neglect time-dependent and secondary reflection/refraction phenomena by rendering simulations with a fairly high sun and/or an interface with low inclination (*e.g.* Cox and Munk (1954) report that secondary reflection events can occur when the slope of the surface facet becomes greater than approximately $\frac{1}{2}$ of the sun’s elevation angle).

5.2 Sub-Pixel Capillary Waves

5.2.1 Bidirectional Reflectance Distribution Function (BRDF)

The bidirectional reflectance distribution function (BRDF) relates reflectance for all combinations of input and output angles (Schott 2007). More specifically, the bidirectional reflectance is the ratio of the radiance reflected in an outgoing direction (θ_r, ϕ_r) to the irradiance from a specified incoming direction (θ_i, ϕ_i) as shown in Eq. (5.1) with units of sr^{-1} where θ represents the zenith angle, ϕ represents the azimuthal angle, and β is the inclination angle of the surface.

$$r_{BRDF} = \frac{L(\theta_r, \phi_r)}{E(\theta_i, \phi_i)} = \frac{dL(\theta_r, \phi_r)}{L(\theta_i, \phi_i) \cos(\beta) d\omega_i} = \frac{dL(\theta_r, \phi_r)}{L(\theta_i, \phi_i) \cos(\beta) \sin(\theta_i) d\theta_i d\phi_i} \quad (5.1)$$

If the surface is treated as a collection of subpixel microfacets, then each facet can be described by an upwind and crosswind slope, denoted as z_x and z_y , respectively. Assuming that the sun is nearly a point source, then a fixed solar and sensor position will require a facet with a predefined slope in order to generate a Fresnel reflection event. The magnitude of the reflection event $[\rho(\alpha)]$ will be a function of the angle of incidence $[\alpha]$ of the solar illumination, and the probability that the wind-roughened surface has the required slope $[p_s(z_x, z_y, W) dz_x dz_y]$, where W represents wind speed. The magnitude of the reflected radiance can therefore be described according to Eq. (5.2).

$$dL(\theta_r, \phi_r) = L(\theta_i, \phi_i) \rho(\alpha) p_s(z_x, z_y, W) dz_x dz_y \quad (5.2)$$

The slope probability density function $p_s(z_x, z_y, W) dz_x dz_y$ describes the fraction of “a small (horizontal) unit area of sea surface for which the slope is within the limits $z_x \pm \frac{1}{2} dz_x$ and $z_y \pm \frac{1}{2} dz_y$ ” (Cox and Munk 1954). Therefore, the “actual” area is $\sec(\beta) p_s(z_x, z_y, W) dz_x dz_y$ (Cox and Munk 1954), which is accounted for by the $\cos(\beta)$ in the denominator of Eq. (5.1). If Eq. (5.2) is inserted into Eq. (5.1), the BRDF expression is:

$$\begin{aligned}
r_{BRDF} &= \frac{d L(\theta_r, \phi_r)}{L(\theta_i, \phi_i) \cos(\beta) \sin \theta_i d \theta_i d \phi_i} \\
&= \frac{L(\theta_i, \phi_i) \rho(\alpha) p_s(z_x, z_y, W) d z_x d z_y}{L(\theta_i, \phi_i) \cos(\beta) \sin(\theta_i) d \theta_i d \phi_i} \\
&= \frac{\rho(\alpha) p_s(z_x, z_y, W) d z_x d z_y}{\cos(\beta) \sin(\theta_i) d \theta_i d \phi_i}
\end{aligned} \tag{5.3}$$

If the slope increments in Eq. (5.3) are converted to angular increments, the resulting analytical expression is shown in Eq. (5.4) where β was previously defined as the *tilt* or *inclination* of a facet and θ_r as the zenith angle of the reflected ray.


$$\begin{aligned}
r_{BRDF} &= \frac{p_s(z_x, z_y, W) \rho(\alpha)}{4 \cos^4(\beta) \cos(\theta_r)} \\
&= \frac{p_s(z_x, z_y, W) \rho(\alpha) \sec^4(\beta)}{4 \cos(\theta_r)}
\end{aligned} \tag{5.4}$$

Research indicates that the *probability density of the occurrence of a given wave slope* (or p_s) can be modeled (to first-order) as an anisotropic two-dimensional Gaussian function that varies with wind speed (Cox and Munk 1956). Examples of this probability density function (PDF) at wind speeds of 5, 10, 30 and 40 $\frac{m}{s}$ can be seen in Figure 5.1.

The expected BRDF, in spherical and Cartesian coordinates, as a result of the generalized slope PDF illustrated in Figure 5.1, is shown in Figure 5.2 and 5.3 for a solar zenith angle of 24° and azimuthal angle of 180° with low to moderate wind speeds between 2 and 8 $\frac{m}{s}$. Validation of DIRSIG's implementation of the microfacet BRDF can easily be achieved by comparison with the analytical results expected from Eq. (5.4). However, a more intuitive investigation will be considered that directly evaluates the expected point spread function associated with the reflection of the solar disc from the simulated facetized surface. This approach is equally applicable to evaluation of the expected bidirectional transmittance distribution function (BTDF) by considering the point spread function associated with solar rays that are transmitted through the air-water interface.


The Cox and Munk Probability Density Function

Wind Speed (m/s) @ 12.5 m above MSL: 5.0
Std Dev X-slope: 0.126 Std Dev Y-slope: 0.098




The Cox and Munk Probability Density Function

Wind Speed (m/s) @ 12.5 m above MSL: 10.0
Std Dev X-slope: 0.178 Std Dev Y-slope: 0.139



The Cox and Munk Probability Density Function

Wind Speed (m/s) @ 12.5 m above MSL: 30.0
Std Dev X-slope: 0.308 Std Dev Y-slope: 0.240



The Cox and Munk Probability Density Function

Wind Speed (m/s) @ 12.5 m above MSL: 40.0
Std Dev X-slope: 0.356 Std Dev Y-slope: 0.277






Figure 5.1: Cox and Munk 2D anisotropic Gaussian PDF for capillary waves at various winds speeds [defined at 12.5m above mean seal level (MSL)]. Note that the x-slope corresponds to upwind, and the y-slope to cross wind directions.


Wind Speed (m/s) @ 12.5 m above MSL: 2.0
Std Dev X-slope: 0.079 Std Dev Y-slope: 0.062 SolZen: 24.0 SolAz: 180.0



Wind Speed (m/s) @ 12.5 m above MSL: 2.0
Std Dev X-slope: 0.079 Std Dev Y-slope: 0.062 SolZen: 24.0 SolAz: 180.0



Wind Speed (m/s) @ 12.5 m above MSL: 4.0
Std Dev X-slope: 0.112 Std Dev Y-slope: 0.088 SolZen: 24.0 SolAz: 180.0



Wind Speed (m/s) @ 12.5 m above MSL: 4.0
Std Dev X-slope: 0.112 Std Dev Y-slope: 0.088 SolZen: 24.0 SolAz: 180.0






Figure 5.2: BRDF for various wind speeds shown in spherical and Cartesian coordinates. Zero azimuthal and zenith angle corresponds to the specular direction. [Note that the maximum reflectance value was compared to that shown by Salinas and Liew (2007) (although they used an isotropic Gaussian PDF) and found to be comparable (*e.g.* $W = 2.0 \frac{m}{s}$, maximum ≈ 0.35 , $W = 4.0 \frac{m}{s}$, maximum ≈ 0.22) (Salinas and Liew 2007)].


Wind Speed (m/s) @ 12.5 m above MSL: 6.0
Std Dev X-slope: 0.138 Std Dev Y-slope: 0.107 SolZen: 24.0 SolAz: 180.0



Wind Speed (m/s) @ 12.5 m above MSL: 6.0
Std Dev X-slope: 0.138 Std Dev Y-slope: 0.107 SolZen: 24.0 SolAz: 180.0



Wind Speed (m/s) @ 12.5 m above MSL: 8.0
Std Dev X-slope: 0.159 Std Dev Y-slope: 0.124 SolZen: 24.0 SolAz: 180.0



Wind Speed (m/s) @ 12.5 m above MSL: 8.0
Std Dev X-slope: 0.159 Std Dev Y-slope: 0.124 SolZen: 24.0 SolAz: 180.0




Figure 5.3: BRDF for various wind speeds. Zero azimuthal and zenith angle corresponds to the specular direction. [Note that the maximum reflectance value was compared to that shown by Salinas and Liew (2007) (although they used an isotropic Gaussian PDF) and found to be comparable (*e.g.* $W = 6.0 \frac{m}{s}$, maximum ≈ 0.16 , $W = 8.0 \frac{m}{s}$, maximum ≈ 0.13) (Salinas and Liew 2007)].

Predicted Visual Results

Figure 5.4 provides a simple visual prediction of the solar glint from the sun with a solar zenith and azimuth angles of zero degrees, a nadir sensing platform, and when the wind speed is equal to $W = 0.5 \frac{m}{s}$ (left) and $W = 1.0 \frac{m}{s}$ (right) (images are oriented such that horizontal is coincident with the upwind direction).




Figure 5.4: Left: Sun glint with $W = 0.5 \frac{m}{s}$, solar zenith and azimuth = 0, and nadir sensing platform. Right: Same as image at left except $W = 1.0 \frac{m}{s}$.

The DIRSIG results will be visually compared to the predicted sun glint as shown in Figure 5.4, and numerically validated by comparison to the predicted magnitude of the reflected sensor reaching radiance using Eq. (5.4), evaluated along transects in the upwind and crosswind directions.

5.2.2 Results

Figure 5.5 plots the predicted and DIRSIG modeled reflected sensor reaching radiance along upwind and crosswind transects for a microfacet air-water interface with variable wind speeds between 2.0, 3.0, 5.0 and 8.0 $\frac{m}{s}$. The DIRSIG simulated imagery that accompanies Figure 5.5 is shown in Figure 5.6, including one extra image that shows the sun glint at the very low wind speed of 0.5 $\frac{m}{s}$.




Figure 5.5: Analytical prediction (solid line) of solar glints versus DIRSIG modeled (dotted line) results, upwind and crosswind at wind speeds of 2.0, 3.0, 5.0 and 8.0 $\frac{m}{s}$.




Figure 5.6: Solar glint DIRSIG generated images at 0.5, 2.0, 3.0, 5.0 and 8.0 $\frac{m}{s}$ (horizontal is crosswind, vertical is upwind).

Although the DIRSIG modeled results are in very good agreement with the analytical prediction for all wind speeds, the fit is slightly better at lower wind speeds. The slight

deviation at the tails of the glint (and with increasing wind speed) is believed to be a product of assumptions associated with the analytical prediction, rather than a shortcoming of either model. For example, the Cox-Munk analytical prediction is believed to be a better approximation of actual physics when used to model low wind speeds. This is because the model assumes that the facet is tilted by an angle that is less than the elevation of the incident ray, and in the absence of any self-shading, which is increasingly valid as the wind speed decreases. According to Zeisse (1995), if the sensor has an elevation of 10° or greater, then the error in the Cox-Munk prediction is on the order of 10% or less for wind speeds less than $11 \frac{m}{s}$, and dropping to less than 1% for wind speeds below $3.5 \frac{m}{s}$. In contrast, DIRSIG's numerical computation is based on raytraced interactions, which need not make any inherent assumptions about the inclination of facets. Given this caveat regarding the analytical prediction, the validation shown in Figure 5.5 is considered extremely successful.

5.2.3 Bidirectional Transmittance Distribution Function (BTDF)

Based on the analytical derivation of the BRDF described in the previous section, an analytical expression for the BTDF is proposed. The goal is to describe the transmittance as a function of the solar and sensor positions directly, when the sensor is submerged and viewing transmitted radiance with an angular distribution as a function of both refraction and the slope distribution of microfacets on a wind-roughened air-water interface. To do this, the facet slopes (z_x and z_y) can be used to relate a given angle of incidence to a given angle of *refraction*. When this is combined with the n-squared constancy of radiance and the magnitude of the transmitted radiance (or $T = 1 - p(\alpha)$), then a suitable BTDF can be formulated.

To begin, consider the relationship between a sloped facet, and its corresponding normal vector (Ross et al. 2005):

$$\begin{pmatrix} x_n \\ y_n \\ z_n \end{pmatrix} = \frac{1}{\sqrt{1 + z_x^2 + z_y^2}} \begin{pmatrix} -z_x \\ -z_y \\ 1 \end{pmatrix} \quad (5.5)$$

If Eq. (5.5) is rearranged, then the slopes can be expressed as shown below:

$$\begin{aligned} z_x &= -\frac{x_n}{z_n} \\ z_y &= -\frac{y_n}{z_n} \end{aligned} \quad (5.6)$$

If the law of reflection is employed, as in the preceding section, then the incident and reflected vectors can be described using the facet normal vector as shown in Eq. (5.7)

where \vec{I} , \vec{N} and \vec{R} are the source, normal and receiver (reflected) vectors, respectively, and α is the local angle of incidence.

$$\vec{I} = 2 \cos(\alpha) \vec{N} - \vec{R} \quad (5.7)$$

Using the conversion between Cartesian and polar coordinates shown in Eq. (5.8), and solving for x_n , y_n and z_n , Eq. (5.6) can be used to relate the facet normal to the source and receiver directions as shown in Eq. (5.9).

$$\begin{pmatrix} x_i \\ y_i \\ z_i \end{pmatrix} = \begin{pmatrix} \sin(\theta_i) \cos(\phi_i) \\ \sin(\theta_i) \sin(\phi_i) \\ \cos(\theta_i) \end{pmatrix} \quad (5.8)$$

$$\begin{aligned} z_x &= -\frac{x_n}{z_n} = -\frac{(x_i + x_t)}{(z_i + z_t)} \\ &= -\frac{\sin(\theta_i) \cos(\phi_i) + \sin(\theta_t) \cos(\phi_t)}{\cos(\theta_i) + \cos(\theta_t)} \\ z_y &= -\frac{y_n}{z_n} = -\frac{(y_i + y_t)}{(z_i + z_t)} \\ &= -\frac{\sin(\theta_i) \sin(\phi_i) + \sin(\theta_t) \sin(\phi_t)}{\cos(\theta_i) + \cos(\theta_t)} \end{aligned} \quad (5.9)$$

The angle of incidence (α) can then be found using the relationship shown in Eq. (5.10), substituting the expressions for upwind and crosswind slope from Eq. (5.5).

$$\begin{aligned} \cos(\alpha) &= \vec{R} \cdot \vec{N} = \vec{I} \cdot \vec{N} \\ \vec{R}^T \vec{N} &= x_n \sin(\theta_t) \cos(\phi_t) + y_n \sin(\theta_t) \sin(\phi_t) + z_n \cos(\theta_t) = \cos(\alpha) \\ \cos(\alpha) &= \frac{-\sin(\theta_t) [z_x \cos(\phi_t) + z_y \sin(\phi_t)] + \cos(\theta_t)}{\sqrt{1 + z_x^2 + z_y^2}} \end{aligned} \quad (5.10)$$

In an analogous manner, the vector relationship between the facet normal, and the source (incident) and receiver (*refracted or transmitted*) directions can be found using Snell's law. In vector form, this is shown in Eq. (5.11) where γ is the angle of refraction, \vec{I} , \vec{N} and \vec{T} are the source, normal and receiver (transmitted) vectors, and n is the index of refraction (assuming that the incident, transmitted and facet normal all point toward the transmitting material (Hecht 2002)).

$$n_t \vec{T} - n_i \vec{T} = (n_t \cos(\gamma) - n_i \cos(\alpha)) \vec{N} \quad (5.11)$$

For convenience, Eq. (5.11) can be re-expressed, assuming that the incident and facet normal point in the same direction, and opposite to the refracted direction, or:

$$n_t \vec{T} - n_i (-\vec{T}) = (n_t \cos(\gamma) - n_i \cos(\alpha)) (-\vec{N}) \quad (5.12)$$

The slopes that “marry” the incident and refracted directions are shown in Eq. (5.13).

$$\begin{aligned} z_x &= - \left[\frac{-x_n}{-z_n} \right] \\ &= - \left[\frac{n_t x_t + n_i x_i}{n_t z_t + n_i z_i} \right] \\ &= - \left[\frac{n_t \sin(\theta_t) \cos(\phi_t) + n_i \sin(\theta_i) \cos(\phi_i)}{n_t \cos(\theta_t) + n_i \cos(\theta_i)} \right] \\ z_y &= - \left[\frac{-y_n}{-z_n} \right] \\ &= - \left[\frac{n_t y_t + n_i y_i}{n_t z_t + n_i z_i} \right] \\ &= - \frac{n_t \sin(\theta_t) \sin(\phi_t) + n_i \sin(\theta_i) \sin(\phi_i)}{n_t \cos(\theta_t) + n_i \cos(\theta_i)} \end{aligned} \quad (5.13)$$

The angle of incidence (α) can be expressed as shown in Eq. (5.14), and then analogous to Eq. (5.2), the transmitted radiance can be expressed as shown in Eq. (5.15), making use of the n-squared constancy of radiance.

$$\begin{aligned} \cos(\alpha) &= I \cdot N \\ \cos(\alpha) &= \frac{-\sin(\theta_i) [z_x \cos(\phi_i) + z_y \sin(\phi_i)] + \cos(\theta_i)}{\sqrt{1 + z_x^2 + z_y^2}} \end{aligned} \quad (5.14)$$

$$dL(\theta_t, \phi_t) = \frac{n_t^2}{n_i^2} L(\theta_i, \phi_i) (1 - \rho(\alpha)) p_s(z_x, z_y, W) dz_x dz_y \quad (5.15)$$

All that remains is to convert the slope increments to angular increments to generate the desired analytical BTDF expression. For the BRDF, the final result is shown in Eq. (5.4),

which was obtained using the determinant of the Jacobian or $dz_x dz_y = |J| d\theta_i d\phi_i$, as shown in Eq. (5.16).

$$J = \begin{vmatrix} \frac{\delta z_x}{\delta \theta_i} & \frac{\delta z_x}{\delta \phi_i} \\ \frac{\delta z_y}{\delta \theta_i} & \frac{\delta z_y}{\delta \phi_i} \end{vmatrix} \quad (5.16)$$

For refraction, the Jacobian of the transformation between facet slope and source position is found to be:

$$|J| = n_i^2 \sin(\theta_i) \frac{\cos(\alpha)}{z_n[n_t |\cos(\theta_t)| + n_i |\cos(\theta_i)|]} \quad (5.17)$$

When Eq. (5.17) is combined with Eq. (5.15), the result is:

$$dL(\theta_t, \phi_t) = \frac{n_t^2 L(\theta_i, \phi_i) (1 - \rho(\alpha)) p_s(z_x, z_y, W) \cos(\alpha)}{z_n[n_t |\cos(\theta_t)| + n_i |\cos(\theta_i)|]} \sin(\theta_i) d\theta_i d\phi_i \quad (5.18)$$

If Eq. (5.18) is integrated over the solid angle of the sun, then the following result is obtained (assuming constant source radiance):

$$\begin{aligned} \frac{L(\theta_t, \phi_t)}{L(\theta_i, \phi_i)} &= \int_{\Omega_{sun}} n_t^2 \frac{(1 - \rho(\alpha)) p_s(z_x, z_y, W) \cos(\alpha)}{z_n[n_t |\cos(\theta_t)| + n_i |\cos(\theta_i)|]} \sin(\theta_i) d\theta_i d\phi_i \\ \frac{L(\theta_t, \phi_t)}{L(\theta_i, \phi_i) \Omega_{sun}} &= n_t^2 \frac{(1 - \rho(\alpha)) p_s(z_x, z_y, W) \cos(\alpha)}{z_n[n_t |\cos(\theta_t)| + n_i |\cos(\theta_i)|]} \end{aligned} \quad (5.19)$$

Eq. (5.19) is the expression for the BTDF, as shown below:

$$\begin{aligned} t_{BTDF} &= \frac{L(\theta_t, \phi_t)}{L(\theta_i, \phi_i) \Omega_{sun}} = \frac{L(\theta_t, \phi_t)}{E(\theta_i, \phi_i)} \\ &= n_t^2 \frac{(1 - \rho(\alpha)) p_s(z_x, z_y, W) \cos(\alpha)}{z_n[n_t |\cos(\theta_t)| + n_i |\cos(\theta_i)|]} \end{aligned} \quad (5.20)$$

Although the derivation of the BTDF described above was attempted from first principles, and using the concepts developed by Cox and Munk (1956), the most recent revision of the DIRSIG BTDF code makes use of a BTDF expression taken from the field of Computer Graphics, and fully described by Walter et al. (2007). For low wind speeds and a high sun the BTDF shown in Eq. (5.20) and that provided by Walter et al. (2007) are consistent (*e.g.* the condition of low wind speed and high sun represents the conditions in which the

original Cox and Munk (1956) BRDF, and therefore the derived BTDF expression shown above, are considered reasonable approximations of the observed physical phenomena). The BTDF expressions only begin to diverge when the wind speed increases, and/or the sun zenith angle increases, which represents conditions where the facet may not actually present its front face to the source, but which is assumed true in the BRDF and BTDF approximations shown above. Bearing this in mind, the DIRSIG generated results were compared to the analytical approximation described by Eq. (5.20).

5.2.4 Results

Figure 5.7 is similar in content to Figure 5.5, showing transects of sensor reaching radiance along the upwind and crosswind directions. However, Figure 5.7 compares the analytical approximation and the DIRSIG modeled results for transmitted sensor reaching radiance when the receiver is submerged within the water volume, parallel with the water-air interface, and looking up or toward the hemisphere above. The DIRSIG simulated imagery that accompanies Figure 5.7 is shown in Figure 5.8, including one extra image that shows the solar transmission at the very low wind speed of $0.5 \frac{m}{s}$.




Figure 5.7: Analytical prediction (solid line) of solar transmission versus DIRSIG modeled (dotted line) results, upwind and crosswind at wind speeds of 2.0, 3.0, 5.0 and $8.0 \frac{m}{s}$.




Figure 5.8: Solar transmission DIRSIG generated images at 0.5, 2.0, 3.0, 5.0 and $8.0 \frac{m}{s}$ (horizontal is crosswind, vertical is upwind).

The DIRSIG transect of the solar subsurface caustic shown in Figure 5.7 is more noisy than its reflected counterpart (shown in Figure 5.5). The reason for the increased noise cannot be explained, despite consultation with DIRSIG development staff. This noise persists even at the very low wind speed of $0.01 \frac{m}{s}$, as shown in Figure 5.9 which compares solar transmission using the Fresnel Radiometry Solver (FRS) (which can handle any planar interface), and the Generic Radiometry Solver (GRS) (which handles the microfacet interface). This plot indicates that the GRS results converge toward the FRS modeled results as the wind speed decreases, thereby providing a simple internal consistency check, despite the persistence of the unexplained noise. Although the noise is somewhat perplexing, its absence in both the FRS caustic and GRS glint solution does in deed support the assertion that the noise is restricted to some sampling or computational factor that is specific to the BTDF, rather than the generalized submerged detector simulation scenario. However, the glint and caustic share the same numerical sampling mechanisms, and other than total internal reflection, projected area effects, and the n-squared constancy of radiance, there are no obvious computational considerations that might explain the noise associated with the subsurface caustic.

Keeping in mind that Figure 5.7 utilizes an infrequent remote sensing scenario that involves an *upward looking* submerged detector, and that a horizontal or downward looking submerged detector solution will include a degree of noise as a result of the photon mapped in-scattered radiance solution, the interface-dependent fluctuations are likely to make only a minor contribution to a simulation's final signal-to-noise ratio (SNR). With this in mind, the details regarding the noise can be further investigated if a practical application for this *upward looking* scenario becomes relevant to future simulations.




Figure 5.9: Fresnel transmittance (black faceted surface outline) versus microfacet BTDF (solid gray shaded surface) with wind speed = $0.01 \frac{m}{s}$.

5.3 Sinusoidal Gravity Waves - Caustics

Procedural sinusoidal gravity waves are implemented in DIRSIG using the following ODB directive, where the bounding box describes the overall dimensions of the volume to be simulated, the magnitude defines the amplitude of the oscillation, the period relates to the wavelength, and the phase shift to the actual *origin* of the first trough with respect to the global geometry origin (0,0,0) as defined in the ODB file (*e.g.* the first trough is at 0,0,0 unless shifted by the phase shift variable). The orientation variable defines the “direction” in which the wave is radiating from (*e.g.* along the X, along the Y, or any arbitrary direction vector). The default is to propagate along the X-direction or (1,0,0). Note that the Z-variable should always be set to zero (Goodenough 2009).

```
SINUSOID {
    BOUNDING_BOX = -500, -500, -0.1, 500, 500, 7
    MAGNITUDE = 0.5
    PERIOD = 2.0
    PHASE_SHIFT = 0.0
    #ORIENTATION = <x> <y> <z>
    MATERIAL_ID = 5
}
```

Image simulation programs such as DIRSIG are capable of following the fate of each ray through raytracing numerical techniques. Since the goal here is validation, rather than the re-invention and/or re-implementation of a numerical technique, we will restrict ourselves to a simplified scenario with a valid analytical approximation as described by Schenck (1957) and illustrated in Figure 5.10.

To begin, we need to ascertain the slope of the interface at all possible points along the air-water boundary, restricted to a one-dimensional sinusoidal wave propagating along the X-axis. The air-water interface can be described by Eq. (5.21) where A defines the wave’s amplitude, and λ its wavelength.

$$Y = A \sin\left(\frac{2\pi X}{\lambda}\right) \quad (5.21)$$

If Eq. (5.21) is differentiated, the result is Eq. (5.22).

$$\frac{dY}{dX} = \frac{2\pi A}{\lambda} \cos\left(\frac{2\pi X}{\lambda}\right) \quad (5.22)$$

The angle of incidence for a *vertical* ray will be equal to the slope as shown in Eq. (5.23).

$$\theta_i = \tan^{-1}\left(\frac{dY}{dX}\right) \quad (5.23)$$




Figure 5.10: Sinusoidal gravity wave and variables. Illustration adapted from Schenck (1957).

For simplicity in analytical expression, we will restrict our interest solely to vertical rays so that Eq. (5.23) and the law of reflection can be used to determine the angle of reflection or θ_{re} (see Figure 5.10) for a solar glint viewed above a sinusoidal air-water interface. The angle of refraction θ_r can likewise be determined using Snell's law. It is then trivial to determine the angle θ_{rV} between the refracted ray and the original vertical ray as shown in Eq. (5.24) and illustrated in Figure 5.10.

$$\theta_{rV} = \theta_i - \theta_r \quad (5.24)$$

If D is defined as the mean depth of the water volume, then the actual depth H of the water volume at point P considering the oscillating waveform is given by Eq. (5.25), where Y can be computed from Eq. (5.21).

$$H = D + Y_{at\ P} \quad (5.25)$$

Finally, Schenck (1957) shows that the displacement of the refracted ray X_d can be determined using Eq. (5.26), and as illustrated in Figure 5.10.

$$X_d = X_{at\ P} + [H \tan(\theta_{r_V})] \quad (5.26)$$

If the system of Eqs. comprised of (5.21) through (5.26) is solved for several vertically incident rays, and the bottom surface is divided into equally spaced bins that correspond to the projection of the sensor's pixel size onto the ground plane, then Eq. (5.26) will determine the number of rays that terminate within any given pixel and therefore the relative intensity of a given pixel (*e.g.* caustics).

5.3.1 Results

Using the “fish-tank” vantage point as shown in Figure 5.11, the caustics generated from variable wavelength and amplitude sinusoidal-gravity waves are shown in Figures 5.12, 5.13 and 5.14. In Figures 5.12 through 5.14 (top), the asterisks depict the actual sinusoidal surface, while the dashed line illustrates the shift in continuous vertically incident rays. The squares depict a subset of these rays, sampled in 2.0 unit increments, which can be matched to the relative intensity plot (solid line) using the predicted refractive-induced shift (*e.g.* in Figure 5.12 the first input sample has zero shift, while the second is shifted backward by approximately 0.5 units, generating the first and second peaks on the relative intensity plot shown as the solid line).

The middle plot in Figures 5.12 through 5.14 replots the solid line (or the analytical prediction of the refractive-induced shift in the vertically incident sample rays) versus the normalized caustic intensity (just above the ground plane) as simulated by DIRSIG (asterisks). Finally, the DIRSIG generated imagery is shown at bottom, but with contrast reversal so that the input “laser” samples are shown as dark gray, and the volume itself appears as light gray. It is readily apparent that the DIRSIG results are in excellent agreement with the analytical prediction. In addition, inspection of the accompanying DIRSIG imagery shows the water depth at which certain incident rays will converge and therefore generate caustics with greater intensity than their neighbors. Of course, the caustic intensity will be diffused by the addition of scattering with the volume of interest, which is investigated in the following section.




Figure 5.11: Preview image. The sinusoidal water body (with exaggerated wave height) is shown in white.






Figure 5.12: Top: Surface wave, refractive-induced shift in caustic, and caustic intensity for several normally incident sample rays. Middle: Transect of normalized caustic intensity (analytical versus DIRSIG) for the wave and water depth depicted above. Bottom: DIRSIG imagery (with contrast reversal such that the actual lasers are shown in dark gray).

Surface Wave (Y), Degree of Shift & Caustic Intensity

GSD (m): 0.050 Amplitude (m): 2.0 Wavelength (m): 5.0
 No. Incident Rays: 14.0 Sampling Interval (m): 2.0 Water Depth (m): 6.0
 Surface (asterisks), Shift (dashed), Relative Intensity (solid), Samples (squares)



Refraction of Caustics from a Sinusoidal Interface

Predicted (solid), DIRSIG (asterisks)






Figure 5.13: Same as Figure 5.12.

Surface Wave (Y), Degree of Shift & Caustic Intensity

GSD (m): 0.050 Amplitude (m): 2.0 Wavelength (m): 5.0
 No. Incident Rays: 14.0 Sampling Interval (m): 2.0 Water Depth (m): 10.0
 Surface (asterisks), Shift (dashed), Relative Intensity (solid), Samples (squares)



Refraction of Caustics from a Sinusoidal Interface

Predicted (solid), DIRSIG (asterisks)





Figure 5.14: Same as Figure 5.12.

5.4 Sinusoidal Gravity Waves - Combination of the Interface and Scattering

If the in-water radiance beneath a sinusoidal interface is compared at more than one optical depth, the evolution of the radiance distribution can be used to evaluate radiance as a function of the interface (optical depth of zero where zero is the air-water interface), and as a function of the interface and scattering (optical depth greater than zero). To examine this phenomena, the independent 3D Monte Carlo simulation results from Zhai et al. (2008) will be used as a reference. A scenario *similar* to that used by Zhai et al. (2008) will be considered, including a slowly varying (20m wavelength) and low amplitude (0.3m) sinusoidal air-water interface. Similar to the work by Zhai et al. (2008), a 10m water volume will be simulated, with a 0.5 reflective ground plane, a scattering coefficient of 0.99, a scattering albedo of 0.5, and a Henyey-Greenstein phase function with an average cosine of 0.95.

Using several incident rays with variable angular orientation, the influence of the interface below a peak, a trough, and between a peak and trough can be examined when observations are taken just below the air-water interface. If observations are taken at progressively deeper depths, the results will be a combination of the interface above, and scattering within the medium.

Figure 5.15 shows the results of Zhai et al. (2008) for a *similar* scenario to that described above. The figure plots radiance as a function of gray-scale values where the central portion of the plot represents radiance with a zenith angle of 180° (coincident with the Z-axis), while the periphery represents radiance at a zenith angle of 90° . Figure 5.15 shows the radiance distribution sampled beneath a peak, a trough, and between peaks and troughs, at geometric depths of 10 (top), 9 (middle) and 5 (bottom) meters above the ground plane, where the ground is located at $Z = 0\text{m}$ and the air-water interface at $Z = 10\text{m}$. The sampling methodology is shown in Figure 5.16 where the squares depict the four locations of interest (below a peak, a trough and between peaks and troughs).

At 10m, or just below the air-water interface, the radiance is expected to be a function of refraction through the sinusoidal wave. At a depth of 9m, focusing (at peaks) and defocusing (at troughs) can be observed. Sampling closer to the ground plane shows radiance that is a function of the direct incident beams, and multiple scattering within the volume. By the time the radiance reaches a geometric depth of 5m from the ocean floor, the distributions are symmetric, dominated by multiple scattering, and no longer show perturbations due to the interface above. This indicates that the asymptotic region is reached somewhere above or near a geometric depth of 5m for the IOPs specified in this scenario (*e.g.* the radiance no longer has an angular dependence, and only varies with depth due to absorption and scattering (Piening and McCormick 2003)).




Figure 5.15: Radiance as a function of the sinusoidal interface and scattering, just below the interface at 10m above (top), at 9m above (middle) and at 5m above the ground plane. The arrow is a vector illustrating the surface normal in the central portion of the image. Figure from (Zhai et al. 2008).




Figure 5.16: Locations and depths at which radiance observations are considered within a 10m water volume. The work of Zhai et al. (2008) considers all three depths, while the DIRSIG results are limited to two depths; just beneath the air-water interface, and at 5m from the ground plane.

Figure 5.17 depicts the radiance distribution obtained from a related DIRSIG simulation. Results are shown directly below the interface, and at approximately 5m above the ground plane. Near the interface, refraction is evident between peaks and troughs, although focusing (beneath a peak) and defocusing (beneath a trough) is not evident (despite re-simulation using variable incident beam widths, sampling angles, sampling locations, and observer sampling area). Although focusing/defocusing was not apparent despite the aforementioned re-simulation efforts, the deficit is not believed to be inherent. Instead, the lack of observable focusing/defocusing is believed to be a product of the simulation configuration, making continued re-simulation efforts of little benefit in light of the successfully predicted refractive effects illustrated in Figures 5.12, 5.13 and 5.14. Of greater interest, however, is the observation that the asymptotic limit is reached somewhere between the interface and a depth of 5m from the ground plane, as predicted by Zhai et al. (2008). This is depicted in the bottom row of Figure 5.17. When these distributions are compared to their counterparts at the more shallow depth (top row of Figure 5.17), it is apparent that perturbations associated with the interface have been lost due to multiple scattering, as predicted by Zhai et al. (2008).

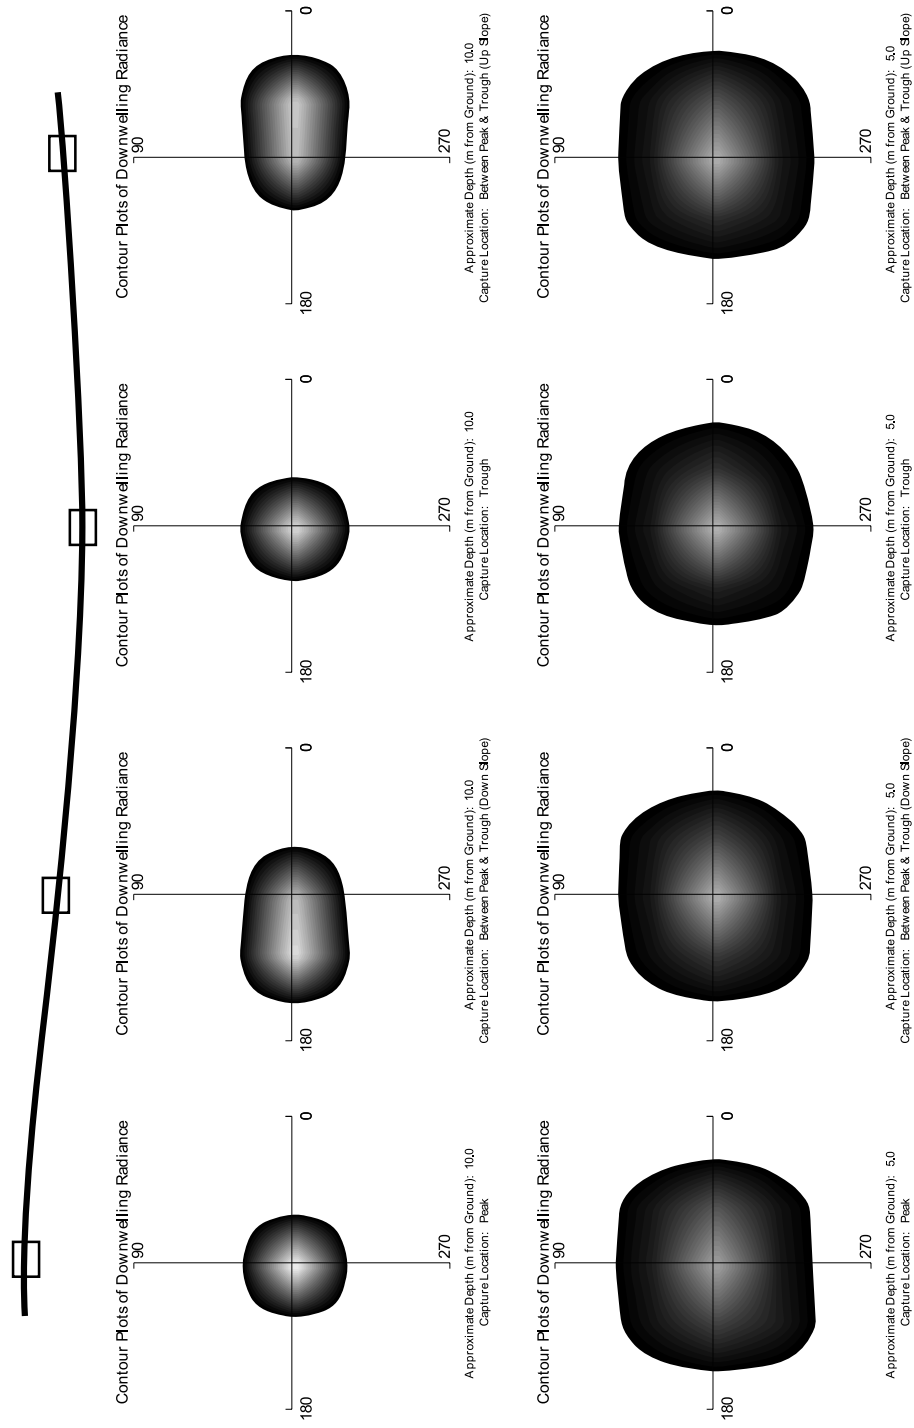


Figure 5.17: Downwelling radiance contours at two heights and four locations beneath a sinusoidal interface (wavelength = 20.0m and magnitude = 0.3m). The Figure captions denote the vertical depth at which the radiance was collected, and the horizontal location along the oscillating sinusoidal wave (*e.g.* collected beneath a peak, a trough, and between peaks and troughs, where down slope denotes the midpoint between a peak and trough, and up slope denotes the midpoint between a trough and a peak).

5.5 Transition to Bottom Effects

This chapter and the one that immediately preceded it validate contributions to sensor reaching radiance that are a function of the air-water interface, including simple interactions and those complicated by a wind-roughened boundary. In geometric and optically shallow volumes indicative of coastal environments, a significant contribution to sensor reaching radiance is expected to originate from bottom effects, skyfraction and background reflections. The following chapter extends the existing validation to include interactions associated with in-water objects.

Chapter 6

Bottom Effects, Skyfraction & Background Reflection

Goal: Evaluate bottom effects, skyfraction and background reflection using the Generic and SurfacePM Radiometry Solvers.

Source Function			
Non-Source Function	✓	Model-vs-Model Comparison	
Submerged or Floating Objects	✓	Model-vs-Analytical Prediction	✓
Boundary Interface	✓	Model-vs-Observation Comparison	
3D Geometry	✓		

This chapter fully describes bottom effects, skyfraction and background reflected contributions to sensor reaching radiance as a function of the Generic and SurfacePM Radiometry Solvers. The solution will begin by considering an analytical prediction of sensor reaching radiance for a simple 3D scenario, before submersion in a water volume (accounting for the additional effects of a refractive interface and attenuation within the participating volume). Chapter 7 extends the solution by adding *two-stream approximate in-scattered radiance* to the *direct reduced radiance* modeled in this chapter, and by considering the overall convergence of the total solution to sensor reaching radiance as modeled by the Generic and SurfacePM Radiometry Solvers when combined with in-scattered radiance modeled by the MediumPM Radiometry Solver.

6.1 Overview

For three-dimensional in-water objects, DIRSIG's Radiometry Solvers should accurately evaluate contributions to sensor reaching radiance that are a function of (1.) simple reflective bottom effects, (2.) an object's skyfraction, and (3.) potential background

reflected contributions due to neighboring objects. Such realistic scenarios would include modeling the light field in shallow waters with significant contributions from the sea floor, including three dimensional natural and man-made objects or structures that can be found on the sea floor (*e.g.* rocks, dock pilings, etc.) In addition, a fortunate consequence of this validation is the evaluation of refractive interface effects such as apparent in-water object displacement and magnification.

Two DIRSIG Radiometry Solvers will be employed and tested; the Generic Radiometry Solver (GRS), and the SurfacePM Radiometry Solver (SRS). These solvers employ raytracing and photon mapping models, respectively, during computation of sensor reaching radiance. Agreement between GRS and SRS generated results will serve to validate both DIRSIG models and show that mechanistically different computational procedures can indeed produce convergent results. To begin, each Radiometry Solver will be briefly described, including qualitative and quantitative predictions for sensor reaching radiance (as a function of user defined variables when appropriate).

6.2 Comparison of GRS, SRS and Analytical Predictions for Skyfraction & Background Reflection in the Absence of a Water Volume

6.2.1 The Generic Radiometry Solver

In an ideal situation, the user would like to sample downwelling irradiance on all targets of interest using an infinite number of samples and an infinitely small sampling step size. However, when limited by computational efficiency considerations, it is more reasonable to utilize prior knowledge of the target's BRDF in order to determine how best to sample all significant downwelling irradiance contributions. This is the job of the Generic Radiometry Solver. For a detailed discussion concerning the computational aspects of GRS, the reader is referred to the related Technical Brief by Goodenough (2007). Briefly, the skydome above a target is divided into equal solid angle sky quads based on user-defined μ (number of zenith angle) and ϕ (number of azimuthal angle) parameters. The assigned BRDF for a given material is then importance mapped, which in turn dictates how best the skydome above the material should be sampled in order to evaluate all significant contributions to downwelling irradiance. Depending upon the nature of the BRDF, more important sky quads (as determined by important regions of the BRDF) are preferentially sampled. The limited number of importance samples (and the fixed sampling step size), combined with the characteristics of the BRDF, may result in some sky quads that are completely missed. For a sky quad that has not been utilized during importance sampling, the user has the option to specify `min_quad` or quick samples that are cast toward the missed quads and which grab a representation of the *emitted radiance* (but are truncated in the presence of geometry that is strictly reflective.)

Based upon the importance mapping of the BRDF, an important sample may intersect another piece of reflective geometry rather than a sky quad. If this happens, the bounce count is incremented. The maximum number of bounce counts represents the maximum number of reflected contributions that will be used to determine the downwelling radiance falling on a given hit point. If the bounce is less than the maximum, then the sample decay rate is used to decrease the number of importance samples and the process begins again. If the maximum number of bounce counts has been reached, and the reflected ray impacts a non-emissive object, then the result is truncated and no additional radiometric contributions are calculated.

For the simple geometry under investigation here, the bounce count will be varied between 1 and 3. To illustrate the impact of this variable consider Figure 6.1. A ray is traced from a detector element and impacts the ground plane (bounce 0). The BRDF is importance sampled (in this case, a completely Lambertian Ward BRDF). If an importance sample strikes the reflective box (bounce 1), a new generation of importance samples (decremented by the sample decay rate) are created. If the BRDF for the box is such that an importance sample impacts the ground plane again (bounce 2), the process will repeat, ad infinitum, unless truncated by the user-assigned max bounce count.

If the bounce count is set equal to 0, then the process will return zero radiance. This is because rays are no longer generated if the bounce is *equal to or greater than 0*. If the bounce count is set equal to 1 [Figure 6.1 (a.)], then the skyfraction is obtained (even in the presence of reflective neighboring geometry). If the bounce count is set equal to 2 [Figure 6.1 (b.)], then hits 0 and 1 are collected. This provides the skyfraction and the *background reflection of the skyfraction that belongs to neighboring objects*. If the bounce count is increased to 3 [Figure 6.1 (c.)], then in addition to the skyfraction and the background reflection of the skyfraction that belongs to the neighboring object, bounce 2 will also collect the *background reflection of the object fraction that belongs to the neighboring geometry*.

Therefore, increasing the bounce count from 1 to 3 for shadowed ground pixels should result in an increase in the sensor reaching radiance. The magnitude of the increase is a function of the skyfraction for both the pixel in question (ground plane), and the skyfraction for the object intersected during bounce 1 (the box). Since the box represents the object fraction for any given ground pixel, and since the box has a skyfraction of 0.5, half of the radiance falling on the box is a product of former reflections. If these former reflections are excluded, then the magnitude of reflected radiance from the object fraction and onto the ground plane will be diminished. For ground pixels with small skyfractions (near the box), this exclusion will be significant and the impact of increasing bounce counts will be fairly important for ground pixels in this shadowed region. Conversely, ground pixels with large skyfractions (far removed from the box) have very little downwelling contributions from reflections off of neighboring geometry so the impact of increasing bounce count will be small to negligible.




Figure 6.1: Illustration of how bounce count and importance samples impact sensor reaching radiance. The rays with two cross hatches represent “quick” rays which only collect source or emitted radiance. For the non-emissive box and ground plane, these rays would be truncated for a bounce count of 1 (a.), a bounce of 2 (b.) and a bounce count of 3 (c.).

For the side of the box, a bounce count of 1 will collect the box’s skyfraction. For a bounce count of 2, the skyfraction of the ground plane will be collected and reflected onto the vertical facets of the box. Since the side of the box has a skyfraction of 0.5, increasing the bounce count to 2 should increase the sensor reaching radiance substantially. However, increasing the bounce count to 3 should not have a major impact on the sensor reaching radiance. This is because the majority of the ground pixels have a very large skyfraction (excluding those in the shadow region of the box). Therefore, inclusion (or exclusions) of the ground’s object fraction will generate only a minor increase (or decrease) in the sensor reaching radiance over that generated with a bounce count of 2. This is illustrated in Figure 6.2.

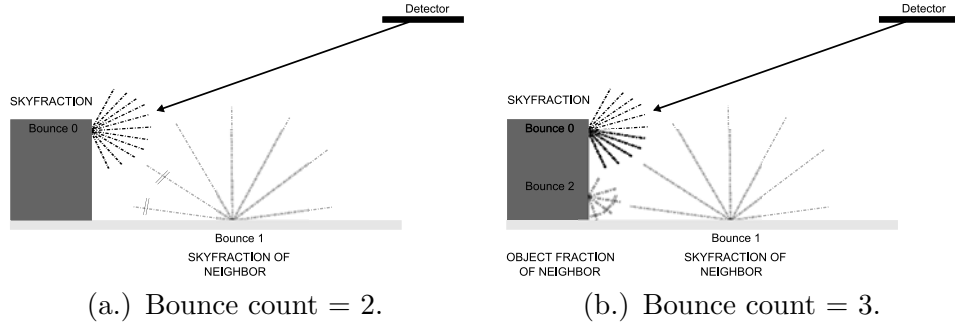


Figure 6.2: The same illustration as shown in Figure 6.1, but for rays with a primary hit (from detector element to first piece of geometry) on the side of the box, rather than the ground plane. Setting the bounce count to 1 collects bounce 0 or the skyfraction (not illustrated). Setting the bounce count to 2 collects the skyfraction for both the box, and the ground plane (a.). Increasing the bounce count to 3 collects hit 0, 1 and 2, but hit 2 represents the ground’s object fraction, which is small for most ground pixels (b.).

Intuitively, it would seem reasonable to assume that as the number of importance samples are increased (for a fixed μ and ϕ), the greater the chance that all quads will be sampled (at least for a uniform Lambertian BRDF), rather than requiring `min_quad` quick samples. If this is true, then there is a greater chance that a sample ray will intersect a neighboring piece of geometry (depending upon the geometry present), and for a fixed bounce count, the sensor reaching radiance may increase. Conversely, if the μ and ϕ samples are increased (for a fixed number of importance samples), a greater number of `min_quad` samples may be needed. If an object has a large skyfraction, the impact of the increased number of `min_quad` samples will be less dramatic than for an object with a smaller skyfraction. This is illustrated in Figure 6.3 for a hit point with a small skyfraction (left) and a large skyfraction (right). The `min_quad` samples are illustrated as solid arrows radiating from hit points, and those with cross hatches will be truncated since neither piece of geometry (the ground plane or box) is emissive. If the importance samples are then increased for the smaller integration step size, the radiance from pixels with small skyfractions will increase more dramatically than those with larger skyfractions since less ‘quick’ samples will (conceivably) be necessary. Of course, this all depends on whether or not the pixel in question has an object fraction with a non-zero reflectance.




Figure 6.3: Illustration of how increasing the μ and ϕ sampling for a fixed number of importance samples will cause a decrease in radiance for objects with a small skyfraction (a.), but will be less significant for objects with a large skyfraction (b.). The quick samples are illustrated as solid arrows radiating from hit points, and those with cross hatches will be truncated since neither piece of geometry is emissive.

6.2.2 GRS Configuration

To test these hypotheses, the geometry shown in Figure 6.4 (top) was simulated using various Radiometry Solver configurations. Several subsequent plots show a transect through the central portion of the image, depicting sensor reaching radiance from pixels on the top of the box, the vertical facet of the box, and along the ground plane. Figure 6.4 (bottom) is meant to orient the reader so that all subsequent plots can be easily interpreted.




Figure 6.4: Illustration of scenario to verify skyfraction and background reflected effects. The simulation includes a 3D box resting on a ground plane. The box is given zero reflectance to directly model skyfraction effects. To evaluate the combined skyfraction and background reflected effects, the box will be given a non-zero reflectance and therefore both the box and ground plane contribute a radiance load onto each other via background reflected contributions (*e.g.* for a uniform skydome, the radiance on the side of the box is half of that on the top of the box, and the radiance on the ground plane adjacent to the box is shadowed by the presence of the box, with skyfraction increasing to nearly 0.99 for ground pixels far removed from the box).

Importance Samples

Referring to Figure 6.5, as the number of importance samples increases from 100 (triangles) to 1500 (solid line), the radiance on the side of the box and on the ground plane in the shadow region increases. This is because an increasing number of importance samples translates into a greater number of ‘real’ rays that can be reflected and therefore invoke additional sampling, as opposed to simple ‘quick’ rays which are invoked to evaluate missed quads (*e.g.* the `min_quad` samples). The increase in sensor reaching radiance is clearly evident for 100 versus 1500 importance samples (solid line versus triangles), but the results begin to converge for 900 versus 1500 importance samples (solid line versus dotted line which nearly coincide.) This confirms the hypothesis that increasing the number of importance samples will increase the sensor reaching radiance, and that the increase is a function of the skyfraction of the target (*e.g.* the smaller the skyfraction, the greater the increase in sensor reaching radiance when coupled with a larger number of importance samples).




Figure 6.5: Illustration of importance sample impact on DIRSIG generated sensor reaching radiance as a function of skyfraction and background reflection. Generic Variables: $\mu = 10$, $\phi = 10$, quick samples = 1, sample decay = 5, max bounce count = 2. Legend: Importance samples = 1500 (solid line), 900 (dotted), 500 (asteriks), 300 (squares), 100 (triangles).

μ and ϕ Sampling with a Fixed Number of Importance Samples

For a constant number of importance samples, the variation in sensor reaching radiance as a function of μ and ϕ is illustrated in Figure 6.6. As μ and ϕ are increased, the integration step size is decreased. In the case of intervening geometry, there is a greater chance that it will be intercepted by a ray, rather than averaged across. However, for a fixed number of importance samples, and increasing μ and ϕ , it is quite possible that a greater number of `min_quad` or ‘quick’ samples will be necessary. Since these samples only collect radiance from real sources or emitting objects, rays that intersect non-emissive objects will cause ray termination and a decrease in the observed radiance. The impact will be more noticeable for pixels with smaller skyfractions, as was illustrated in Figure 6.3. The DIRSIG generated results for `min_quad` = 1, sample decay = 5, max bounce count = 2, importance samples = 900 and a variable integration step size or $\mu = \phi = 40$ (squares), $\mu = \phi = 30$ (asterisks), $\mu = \phi = 20$ (dotted), $\mu = 10$ and $\phi = 20$ (solid line) is shown in Figure 6.6. The plot clearly indicates that the sensor reaching radiance decreases as μ and ϕ are increased for a fixed number of importance samples for shadowed pixels on the ground plane (small skyfraction), and especially for pixels on the side of the box (small skyfraction). This confirms the hypothesis that decreasing the integration step size (increasing μ and ϕ) with a fixed number of importance samples decreases the sensor reaching radiance, and that the impact is most noticeable for objects with small skyfractions.

Sensor Reaching Radiance as a Function of SkyFraction & Background Reflection




Figure 6.6: Illustration of μ and ϕ sampling impact on DIRSIG generated sensor reaching radiance as a function of skyfraction and background reflection. Generic Variables: quick samples = 1, sample decay = 5, max bounce count = 2, importance samples = 900. Legend: $\mu = \phi = 40$ (squares), $\mu = \phi = 30$ (asterisks), $\mu = \phi = 20$ (dotted), $\mu = 10$ and $\phi = 20$ (solid line).

Importance Sampling with Fixed μ and ϕ Sampling

If the number of importance samples are increased alongside an increase in μ and ϕ , then it is reasonable to assume that less ‘quick’ `min_quad` samples will be needed, and for small skyfraction pixels, the sensor reaching radiance will increase. This is illustrated in Figure 6.7 where $\mu = \phi = 40$ for 900, 1500, 2500 and 3500 importance samples, as compared to 1500 importance samples with $\mu = 10$ and $\phi = 20$. The plot illustrates that for the side of box and in the sky-shadowed region of the ground plane, as the number of importance samples increases for the smaller integration step size, the sensor reaching radiance also increases, and approaches that found for 1500 importance samples, but with less sky quads. Again, the increase in sensor reaching radiance is most evident for pixels on the side of the box. Although the results begin to converge for μ and $\phi = 40$ and 1500 versus 3500 importance samples, 1500 importance samples with $\mu = 10$ and $\phi = 20$ still shows a higher return.




Figure 6.7: Illustration of how increasing μ and ϕ as well as importance sampling impacts DIRSIG generated sensor reaching radiance as a function of skyfraction and background reflection. Generic Variables: quick samples = 1, sample decay = 5, max bounce count = 2, $\mu = \phi = 40$, importance samples = 900 (solid line), 1500 (dotted line), 2500 (astericks), 3500 (squares), and $\mu = 10$, $\phi = 20$ with importance = 1500 (triangles).

6.2.3 SurfacePM versus the Generic Radiometry Solver

Based on the preceding plots and discussion it is clear that the user can exert considerable influence on sensor reaching radiance. Understanding the computational mechanisms employed by GRS allows the user the ability to configure the Radiometry Solver while making choices between fidelity and efficiency. The same is true of the SurfacePM Radiometry Solver, although this solver determines sensor reaching radiance in a manner that is fundamentally different from the backward raytracing model employed by GRS. Instead of casting rays from detector elements into the scene, and then generating secondary rays to sample the downwelling irradiance as employed by GRS, the SurfacePMRadSolver (SRS) is a surface photon mapper. It functions in a two-pass manner that is consistent with how the volumetric photon mapper computes in-scattered radiance. During the first pass, rays are generated at sources and undergo reflections and absorptions (and scattering if a volumetric photon mapped medium exists between the surface photon mapped material and the sources). During the second traditional raytracing pass, rays are generated at detector elements and if they intersect a material assigned SRS, the user-defined spectral search radius is used to collect photons on the surface (about the ray) and compute the downwelling irradiance. This is illustrated in Figure 6.8 (b.).




Figure 6.8: Comparison of the computational mechanisms employed by the Generic (a.) and SurfacePM (b.) Radiometry Solvers.

The user influences the SRS solution by making two choices. The first is the number of photon histories that should be traced to accurately sample the irradiance distribution. The second concerns the area of influence around a propagating ray that should be queried to accurately estimate the irradiance at any given location in the scene. Both of these parameters impact the signal to noise ratio (SNR) and bias in the solution. The bias can be drastically reduced by employing boundary compensation measures. The newest version of DIRSIG computes radiance by first searching for photons within the nominal search area specified by the user, but adaptively reducing this area should rapid changes in the computed radiance be evident. The result is that the radiance is collected *along edges*, instead of averaged *across them*. This is shown in Figure 6.9. The top row depicts

the biased imagery of a step function simulated without boundary compensation and with increasing SRS search radii. The middle row depicts the same imagery, but when the SRS search radius is fixed and equal to the GIFOV of the detector. Although this provides for a reduction in bias, the SNR is likewise reduced. The bottom row illustrates the effect of the aforementioned adaptive boundary compensation. The majority of the bias is absent, while still maintaining a relatively high SNR. This affords the user the freedom to use a relatively large search radius (thereby increasing SNR), while allowing DIRSIG to adaptively adjust the nominal radius for surfaces that exhibit rapid changes in irradiance.

The remainder of this section compares the solution of sensor reaching radiance as determined by GRS and SRS for skyfraction and background reflection. Sections that follow include the addition of a refractive air-water interface, attenuation through the water volume, and in-scattered radiance contributions.

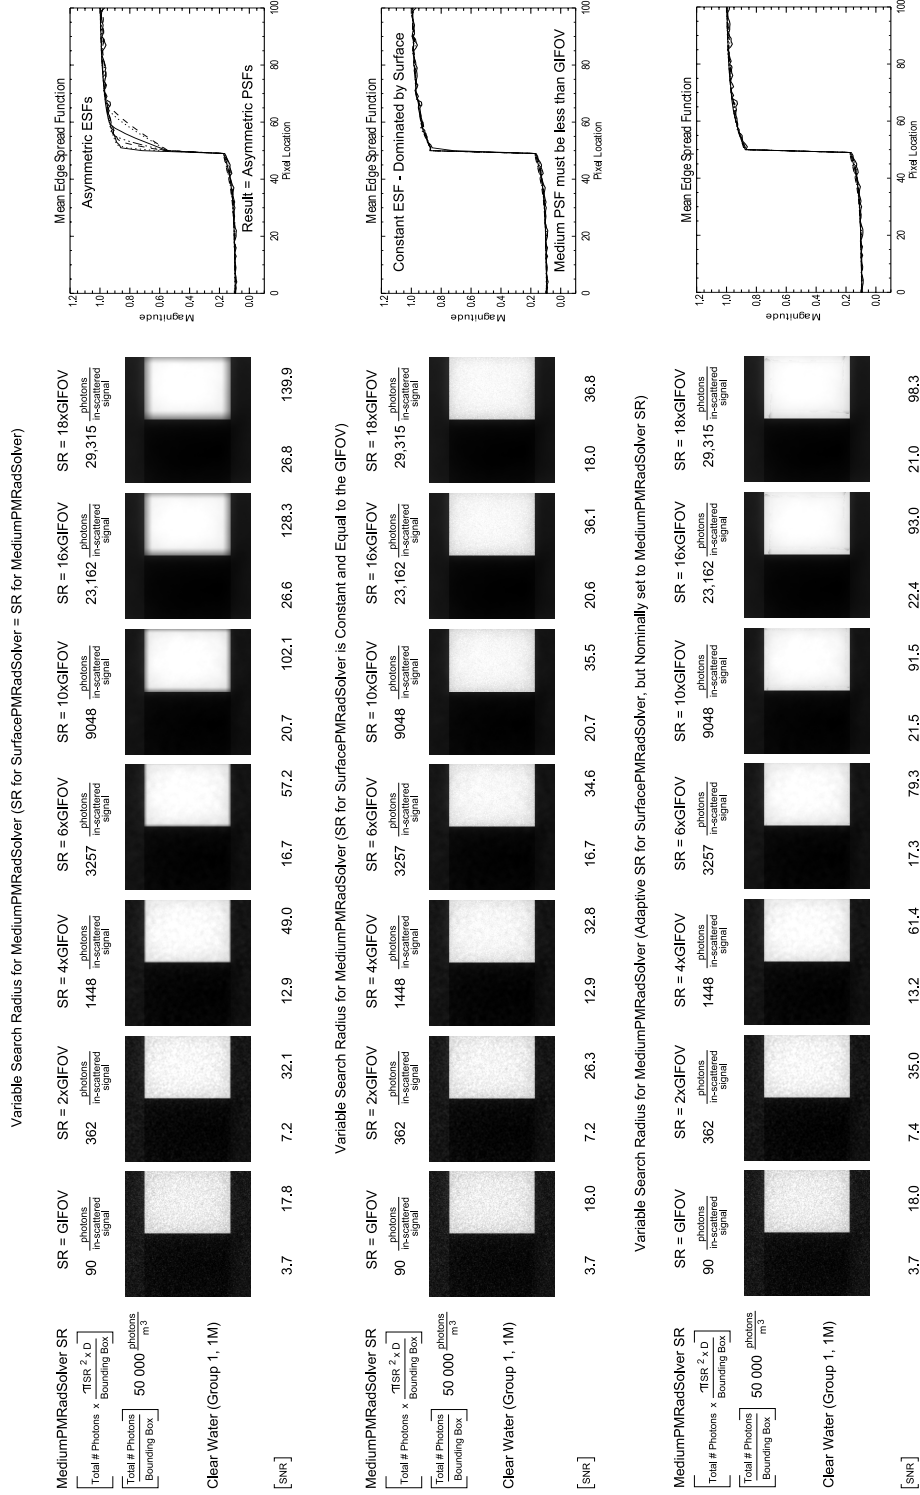


Figure 6.9: Top row: SNR and bias as a function of increasing search radius (SR). Middle row: SNR and bias reduction associated with a fixed surface search radius. Bottom row: Bias reduction and SNR enhancement as a function of the adaptive adjustment of the nominal surface search radius. Group #1 specifies the IOPs as shown in table 9.3, for all targets at 1.0m depth (from the air-water interface) and using the Henyey-Greenstein phase function with $g = 0.7$. The GIFOV is 2.4cm.

6.2.4 Analytical Predictions

The skyfraction represents the portion of the skydome that is unobstructed by intervening geometry. The total downwelling irradiance, obtained by integrating over the solid angle of the unobstructed skydome, and expressed as a fractional value of the total solid angle of the skydome, will vary between 0.5 for a pixel on the ground plane near the box or vertical obstruction, and asymptotically approach 1.0 for pixels far removed from any other piece of geometry. This is illustrated in Figure 6.10 for ground samples (moving progressively further away from a piece of 3D geometry (or a box meant to represent any type of polygonal geometry), where the skyfraction has been set to zero for ground pixels beneath the box). If a similar plot were prepared for the box, then the skyfraction would be 1.0 on the top (horizontal) facet, and 0.5 on each side (vertical) facet.




Figure 6.10: Predicted (exact) skyfraction for pixels on the ground plane. The skyfraction for ground pixels has been set to zero in the location of the box. The skyfraction then jumps to 0.5 for pixels adjacent to the box, before asymptotically approaching 1.0.

For the simple scene under investigation here (see Table 6.1), the generalized expression for sensor reaching radiance will be $L_{Sensor} = L_{SkyFraction} + L_{Background}$, which reduces to $L_{Sensor} = L_{SkyFraction}$ for a totally absorbing box.

The sensor reaching radiance (per pixel) as a function of skyfraction can be calculated as shown in Eq. (6.1), making use of a Lambertian reflector and uniform skydome, for a box located at the extreme left edge of the image, and using the following radiometric variables: L as radiance, L_d as downwelling radiance, $L_{skyfraction}$ as upwelling radiance as a function of skyfraction, L_{sensor} as sensor reaching radiance, E_d as downwelling irradiance, r as reflectance, F as skyfraction, θ as zenith angle, ϕ as azimuthal angle, θ_h as angular height

Table 6.1: Skyfraction Effects.

Object Geometry	Material Properties	Illumination	Atmosphere	Image Capture
Ground Plane, Box	$r_g = 1.0$ $r_b = 0.0$	Uniform Downwelling, Hemisphere $E_h = 1.0$ SkyFraction=1.0	Simple Atmosphere, Transmission = 1.0	Vis-Panchromatic 0.4-0.6 μm
Ground Plane, Box	$r_g = \text{variable}$ $r_b = \text{variable}$	Uniform Downwelling, Hemisphere $E_h = 1.0$ SkyFraction=1.0	Simple Atmosphere, Transmission = 1.0	Vis-Panchromatic 0.4-0.6 μm

r_g is the ground reflectance and r_b is the box reflectance. Note that skyfraction as used in the simple atmosphere DIRSIG class, divides the user-defined hemisphere irradiance between skylight and solar contributions. Note that sandy bottoms have reflectance values between 0.2 and 0.6, which is decreased by the presence of detritus and algae (Carder et al. 2003).

of obstruction, and ϕ_w as angular width of obstruction (see Figure 6.11 for clarification of the angular variables).




Figure 6.11: Illustration of angular variables shown in Eq. (6.1).

$$\begin{aligned}
 L_u &= E_d \frac{r}{\pi} \\
 L_{Skyfraction} &= F E_d \frac{r}{\pi} \\
 E_d &= \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\frac{\pi}{2}} [L_d \cos\theta \sin\theta \, d\theta \, d\phi]
 \end{aligned}$$

$$\begin{aligned}
E_d &= L_d \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\frac{\pi}{2}} \cos\theta \sin\theta \, d\theta \, d\phi \quad (\text{Lambertian assumption for } L_d) \\
L_d &= \frac{E_d}{\int_{\phi=0}^{2\pi} \int_{\theta=0}^{\frac{\pi}{2}} \cos\theta \sin\theta \, d\theta \, d\phi} = \frac{E_d}{\pi} \\
L_{Skyfraction} &= F L_d \left[\int_{\phi=0}^{2\pi} \int_{\theta=0}^{\frac{\pi}{2}} \cos\theta \sin\theta \, d\theta \, d\phi \right] \frac{r}{\pi} \\
L_{Skyfraction} &= \frac{E_d}{\pi} \underbrace{\left[\int_{\phi_w}^{2\pi} \int_0^{\frac{\pi}{2}} \cos\theta \sin\theta \, d\theta \, d\phi + \int_0^{\phi_w} \int_{\theta_h}^{\frac{\pi}{2}} \cos\theta \sin\theta \, d\theta \, d\phi \right]}_F \frac{r}{\pi}
\end{aligned} \tag{6.1}$$

The analytical expression shown in Eq. (6.1) can be extended to include background reflected contributions, provided the user is willing to accept simplifying assumptions about the skyfraction associated with both the target and the background object. As the bounce count is increased from 1 to 3, the analytically predicted results must account for the pixel's sky fraction, as well as energy that is downwelling onto the pixel of interest that is a function of reflection from adjacent objects (with their own skyfraction). A bounce count of 2 will collect energy that is a function of the target's skyfraction, and the skyfraction of the background object. If the bounce count is increased to 3, then the sensor reaching radiance will contain the additive contribution of the target's skyfraction, the skyfraction of the background, and a restricted component of the object fraction of the background (restricted since an infinite number of bounces is not being considered). If either the target or the background has a variable skyfraction (*e.g.* the ground plane's skyfraction varies from 0.5 to approximately 1.0, depending upon the pixel location), then an estimated value of \hat{F} is required in the analytical expression, which replaces the numerical sampling that is carried out by DIRSIG.

The analytical expression that approximates the numerical computation of skyfraction and background reflection is shown in Eq. (6.2), where F_T is the skyfraction of the target, r_T is the reflectance of the target, r_{Back} is the reflectance of the background, E_d is the downwelling irradiance, and \hat{F}_{Back} and \hat{F}_T are estimated values of the skyfraction of the background and the target, respectively. If the box is the target (or the background), then \hat{F}_T (or \hat{F}_{Back}) is exact and equal to 0.5. If the ground is the target (or the background), then \hat{F}_T (or \hat{F}_{Back}) should be set fairly high since the skyfraction of the ground is very high except for pixels immediately shadowed by the box itself. Since the average skyfraction of all the ground pixels is 0.925, it is reasonable to set the estimate of \hat{F}_T (or \hat{F}_{Back}) to this high average value.

$$L_{Target} = L_{Skyfraction} + L_{Background}$$

$$\begin{aligned}
L_{Skyfraction} &= F_T E_d \frac{r_T}{\pi} \\
L_{Background} &= \left(\underbrace{E_d \hat{F}_{Back} \frac{r_{Back}}{\pi}}_{\text{Bounce 2}} + \underbrace{\left[(1 - \hat{F}_{Back}) \frac{E_d r_{Back}}{\pi} \right] r_T \hat{F}_T}_{\text{Bounce 3}} \right) (1 - F_T) r_T
\end{aligned}$$

F_T = variable and exact for ground, 0.5 and exact for box
 \hat{F}_T = estimated at 0.925 for ground, exact at 0.5 for box
 \hat{F}_{Back} = estimated at 0.925 for ground, exact at 0.5 for box

(6.2)

The overall analytical methodology is described in Figure 6.12, and the predicted results for bounce count = 1, 2 and 3 are overlaid in Figure 6.13. Inspection of Figure 6.13 allows the reader to draw three major conclusions. The first is that increasing the bounce count generates an incremental increase in the sensor reaching radiance reflected from pixels in the sky-shadowed region of the ground plane. This is because these target pixels, and the background object itself (the box), both have small skyfractions. The second observation is that an increase in bounce count from 1 to 2 generates a substantial increase in sensor reaching radiance for the side of the box. This is because the side of the box has a small skyfraction. Therefore, the addition of the reflection of the skyfraction from the ground, and onto the box, has a substantial impact on the magnitude of the sensor reaching radiance. Since the analytical prediction of sensor reaching radiance for pixels on the side of the box is based on a high and constant value of $\hat{F}_{Back} = \hat{F}_g$ (where g = ground), the analytical prediction will also be constant, regardless of the target location on the side of the box. Therefore, the analytical prediction should match the DIRSIG modeled results (for at least a few pixels - *e.g.* intersect the DIRSIG generated curve), and depending upon the analytical value of \hat{F}_g , be too low for pixels near the top of the box (pixels near location 50), and too high for pixels near the bottom of the box (pixels near location 100). The third major conclusion is that increasing the bounce count from 2 to 3, does not generate a great deal of additional sensor reaching radiance reflected from the side of the box. This is because the majority of the ground pixels have a very high skyfraction, and therefore inclusion of the ground's object fraction (bounce 3) adds very little additional radiance to the total solution.

In summary, if both the target and background have small skyfractions, then increasing the bounce count will incrementally add additional radiance to the total solution (although with diminishing returns as the number of reflections increases). If the target has a small skyfraction and the majority of the background has a large skyfraction, then increasing the bounce count from 1 to 2 will generate a significant increase in sensor reaching radiance, but little additional return will be gained by increasing the bounce count from 2 to 3.

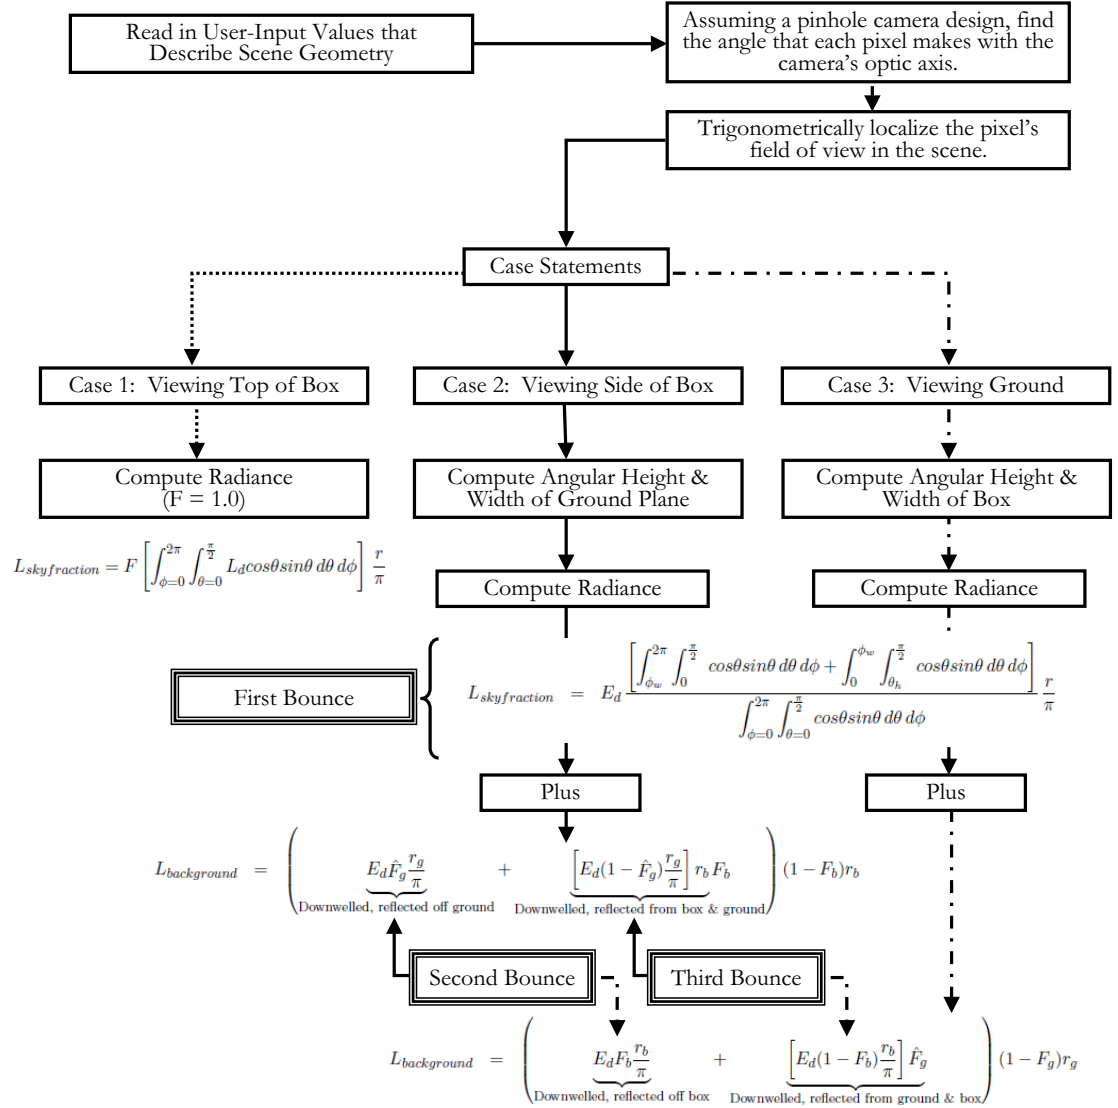


Figure 6.12: Analytical methodology used to predict sensor reaching radiance $[W m^{-2} sr^{-1}]$ as a function of skyfraction and background reflections.




Figure 6.13: Analytical prediction of sensor reaching radiance [$W m^{-2} sr^{-1}$] for bounce count = 1 (solid line), 2 (dotted-dashed line) and 3 (asterisks) for $r_g = r_b = 1.0$, and $\hat{F}_{Back} = \hat{F}_g = 0.925$.

6.2.5 Results

The agreement between an analytical approximation of sensor reaching radiance and that computed using DIRSIG's Generic and SurfacePM Radiometry Solvers will be restricted to a very small case study with constant scene and environment variables as defined in an invariant Object Database (ODB) and Configuration (CFG) file. The GRS configuration, as defined in the Material (MAT) file, is a product of the small trade study conducted in Section 6.2.2. After considering the convergence of the results, weighed by efficiency, the GRS parameters will be as follows: importance samples = 900, decay rate = 5, $\mu = 10$, $\phi = 20$ and `min_quad` = 1. For SRS, an initial trade study will not be conducted since the expected variability in the results will be restricted to variance and bias (as a function of the number of photon histories traced and the search radius adopted); these parameters will be set high (8 million) and low (1.0 units), respectively, although both can be augmented as necessary to drive down either the variance or bias.

For a variable bounce count between 1 and 3, Figure 6.14 compares the DIRSIG GRS modeled results (diamonds), to those predicted analytically (dashed line) for $r_g = r_b = 1.0$. The results were repeated for variable combinations of r_g and r_b (not shown here), and the same trend is apparent, with very good agreement between the analytical prediction and DIRSIG computed sensor reaching radiance.




Figure 6.14: Analytical prediction of sensor reaching radiance (dashed lines) versus GRS (diamonds) for bounce count = 1, 2, and 3, $r_g = r_b = 1.0$, and $\hat{F}_{Back} = \hat{F}_g = 0.925$.

When the analytical prediction is compared with the SRS modeled results (see Figure 6.15), the computation of sensor reaching radiance as a function of *skyfraction* also shows

excellent agreement. As the bounce count is increased to account for background reflection, the reflected radiance from the side of the box increases as expected (although slightly lower in magnitude than that predicted either analytically or numerically using GRS). However, as the bounce count is increased, even up to 100, the reflected radiance from the shadowed ground plane does not increase. Moreover, the lack of an apparent background reflected contribution from the side of the box and onto the shadowed ground plane does not appear to be a function of the reflectance of the box ($r_b = 0.7$ or $r_b = 1.0$), the number of photons in the map (8 millions versus 12 million), or the spectral search radius (1.0 or 0.5 units). In fact, despite the troubleshooting effort, all SRS configurations result in nearly equivalent sensor reaching radiance, which is evident from the fact that all SRS plots overlap (barring statistical variability) between pixel location 100 and 600 in Figure 6.15.




Figure 6.15: Analytical prediction of sensor reaching radiance (dashed lines) versus SRS (diamonds) for bounce count = 0 (GRS 1), 2 (GRS 3), 10, and 100 using SSR = 1.0 and 8 million photons, and for a bounce count of 10 with SSR = 0.5 and 12 million photons, with $r_g = 1.0$ and $r_b = 0.7$, and $\hat{F}_{Back} = \hat{F}_g = 0.925$.

The exact explanation for the divergence is unclear, although the issue has been discussed with the DIRSIG development team. The background reflected contribution from the ground and onto the side of the box suggests that the basic physics has been modeled correctly. This makes troubleshooting even more difficult since the explanation must include something that is specific to the ground surface map, rather than a pervasive problem.

If asked to speculate, the author would first consider a trace of individual photon events

during the first pass to determine the percentage of events that include photons reflected onto the ground from the side of the box. If completely absent, which is what the results seem to indicate, then one possible explanation is the differential filling of maps. For example, if the ground surface map fills as a function of direct sky radiance before having a chance to collect photons that have been reflected by other pieces of geometry, then it will exhibit the correct skyfraction, but little to no background reflected contribution. This was recently tested and disproved by using a single map across all surfaces within the volume. However, the more recent simulations raise new concerns since the results are drastically different from former simulations.

Reverting to two separate maps (one for each material type in the scene), Figure 6.16 repeats the basic simulation shown in Figure 6.15 (circa 2009), but with very different sensor reaching radiance near pixel 100 (circa 2010), where $SSR = 1.0$, bounce count = 2 and max photons = 8M per surface map. No explanation can be provided by the author for this change. Moreover, the change is believed to be unrelated to purposeful and directed work at the request of the author (or the team) concerning background reflections. Instead, the change appears to be a by-product of code changes that are the result of other modeling efforts. Although changes such as this cannot be foreseen or avoided given the constant evolution of the DIRSIG model, it is clear that this is an area that can benefit from continued investigation. *Note: with the exception of Figure 6.16, the results of this chapter are circa 2009 or earlier, and therefore do not reflect the more recent SRS solution.*




Figure 6.16: Compare with Figure 6.15 (generated over a year ago) and note the unexplained changes in the SRS modeled response of radiance between pixels 50 and 200. The dashed lines show the analytical prediction of sensor reaching radiance (for skyfraction [lower of two curves] and skyfraction + background reflection [higher of two curves]). The diamonds depict the SRS modeled results with bounce count = 100 using SSR = 1.0 and 8 million photons, collected January, 2009. Compare this to results collected June, 2010 (diamonds + solid line) with bounce count = 2, SSR = 1.0, 8M photons per surface map, $r_g = 1.0$ and $r_b = 0.7$. The divergence is a function of some unknown change that occurred in the last year, and believed to be unrelated to any directed or purposeful effort to account for background reflected contributions with this Radiometry Solver.

6.3 Comparison of GRS, SRS and Analytical Predictions for Skyfraction in the Presence of an Air-Water Interface

6.3.1 Refraction & Skyfraction


Immersing the scenario illustrated in Figure 6.8 further complicates the prediction of sensor reaching radiance. The air-water interface is marked by an abrupt change in refractive index which generates Fresnel reflected and transmitted radiance contributions, alters the apparent location of the box, and reduces the skyfraction to, at most, Snell's cone. Each complication is considered in more detail below.

Shift in Size & Location


The refractive induced shift in location and size is illustrated in Figure 6.17. The ground plane and box will appear closer to the detector than that predicted based on the simple design of the geometry as described in the ODB file. DIRSIG uses a ray tracer which inherently accounts for each refraction event and therefore the apparent depth of an object is foreshortened automatically. For an analytical prediction, the air-water refractive index ratio can be used to statically determine the new location of the box. This is illustrated in pseudo code below:

$$\begin{aligned}
 \text{apparentWaterDepth} &= \text{waterDepth} \frac{n_{air}}{n_{water}} \\
 \text{apparentDepthofTopofBox} &= \text{boxHeight} \frac{n_{air}}{n_{water}} \\
 \text{apparentDistanceGroundToPlatform} &= \text{DistanceGroundToPlatform} - \text{waterDepth} + \text{apparentWaterDepth} \\
 \text{apparentDistancePlatformToTopofBox} &= \text{DistanceGroundToPlatform} - \text{waterDepth} \\
 &\quad + \text{apparentWaterDepth} - \text{apparentDepthofTopofBox} \\
 \text{edgeofBoxAtBoxHeightonFocalPlane} &= \frac{x}{2} - \left(\frac{f}{\text{pixelSize}} \frac{\text{edgeofBoxFromOAAAtGround}}{\text{apparentDistancePlatformToTopofBox}} \right) \\
 \text{edgeofBoxAtGroundonFocalPlane} &= \frac{x}{2} - \left(\frac{f}{\text{pixelSize}} \frac{\text{edgeofBoxFromOAAAtGround}}{\text{apparentDistanceGroundToPlatform}} \right) \quad (6.3)
 \end{aligned}$$

The majority of variables written in pseudo code above, are self explanatory. The exceptions are: x = number of columns in the image, and f = focal length of the sensor. The calculated values serve to redefine the location of the box, as viewed from the platform, due to submersion. For example, a ground-to-platform distance of 100 units will have an apparent distance of ≈ 92 units due to the apparent location of the ground plane inside the water volume. Likewise, a water depth of 30 units will have an apparent depth of ≈ 23 units. Similarly, if the top of the box is 20 units above the bottom plane, it will have an apparent position only ≈ 15 units above the bottom plane.



(a.) Scenario to be validated.




(b.) Shift in location and size due to submersion.


Figure 6.17: Same as Figure 6.4 but submersion causes an apparent shift in location and size for in-water objects.

Change in Sensor Reaching Radiance

The angle of incidence and refractive indices on either side of the interface also impact the expected sensor reaching radiance, which are modeled as Fresnel surface effects. The planar air-water boundary will generate a 0.02 reflectance contribution (for the majority of air-incident rays), and attenuate the transmitted ray. The transmission is $T = 1 - r$, illustrated in Figure 6.18 (top) for air-incident rays. As can be seen from Figure 6.18 (top), the transmittance through the boundary (air-to-water) will be very large for the majority of air-incident rays.



(a.) Air-incident rays.



(b.) Water-incident rays.

Figure 6.18: Predicted Fresnel reflectance (dotted line) and transmittance (dashed line) as a function of angle of incidence for (a.) air-incident rays and (b.) water-incident rays when $n_{air} = 1.003$ and $n_{water} = 1.34$.

Change in Skyfraction

Submerged horizontal facets have skyfractions with a maximum zenith angle of twice the critical angle for all azimuth angles (or a solid angle equal to Snell’s cone). Figure 6.18 (bottom) shows that the skyfraction is reduced, to at most, Snell’s cone, due to total internal reflection. For GRS, importance samples are generated from the primary hit location (defined as the first non-transmissive material that is hit by a ray cast out from a detector element – *e.g.* box or ground plane). Ignoring possible attenuation and in-scattering within the volume, and instead restricting our attenuation to interface-related events, the GRS importance samples will impact the water-air boundary while attempting to sample the downwelling irradiance falling on the primary hit location. If the rays intersect the water-air boundary with an angle of incidence greater than the critical angle ($\approx 48^\circ$) then the ray will undergo total internal reflection as illustrated in Figure 6.18 (bottom) for water-incident rays.

Therefore, the reduced skyfraction, as modeled by GRS, is a product of traditional second pass events, wherein importance samples beyond Snell’s cone are internally reflected. In contrast, SRS introduces skyfraction effects during first pass events. The SRS first pass is carried out by the `CDMonteCarloTrace::walkRay` function. Rays are generated at sources, and propagated toward the material to be surface photon mapped. When an interface is present, the rays (photons) will be refracted according to Snell’s law. This means that for any point on the photon mapped surface, refraction will shrink the skyfraction to (at maximum) Snell’s cone. Therefore, the photon population, augmented by Snell’s cone, is a product of the first pass. During the second pass, rays are projected from detector elements, and those photons which accumulate within a user defined radius around the ray of interest will contribute to sensor reaching radiance. Although the reduced skyfraction is a product of second pass events with GRS, and first pass events with SRS, the overall net result is the same, as illustrated in Figure 6.19.

For rays that intersect the water-air interface with incidence angles less than the critical angle, refraction away from the boundary normal increases the *sampled* skyfraction for non-occluded in-water objects as compared to the skyfraction computed for the same object in the absence of the water volume. For a non-uniform skydome, the sampled downwelling irradiance spans a larger skyfraction than expected, although the samples will still be attributed to a solid angle that occupies Snell’s window. Furthermore, rays that are totally internally reflected may impact nearby submerged objects, and therefore contribute to a larger object fraction. These possible ray fates are illustrated in Figure 6.20, assuming the GRS model.

Object Fractions

Although the skyfraction will be reduced to, at most, Snell’s cone, this maximum can be further reduced by neighboring geometry. This is illustrated in Figure 6.21. If a pixel




Figure 6.19: Comparison of the computation of sensor reaching radiance in the presence of an interface using GRS and SRS. For GRS, a ray is generated at a detector element during the second pass, which gives rise to secondary rays that are generated from the initial surface hit location (arrows leaving the hit point and exiting the interface). Samples beyond Snell’s cone are internally reflected back toward the submerged surface. For SRS, rays are generated at sources during the first pass. These rays pass through the interface (arrows entering the interface), undergo refraction, and in the absence of a scattering volume, are absorbed or reflected from the submerged surface. During the second pass, rays are generated at detector elements and photons on the surface are used to determine the sensor reaching radiance.




Figure 6.20: Skyfraction as a function of refractive boundary effects for primary hits near the vertical obstruction (a.) and at increasing distances from the vertical obstruction (b.) and (c.).

“sees” neighboring geometry with an angular height that is less than $[90^\circ - \text{critical angle}]$, then the skyfraction spans a zenith angle of twice the critical angle for all azimuthal angles (or a solid angle equal to Snell’s cone). This is illustrated in Figure 6.21 (right). Conversely, if a pixel is very near neighboring geometry, then it may “see” an obstruction with an

angular height [see Figure 6.21 (left) where the curved arrow shows the angular height of the obstruction] that is greater than $[90^\circ - \text{critical angle}]$. If so, the skyfraction above the obstruction corresponds to the *expected skyfraction in the absence of an interface*, while the skyfraction for all other azimuthal angles corresponds to twice the critical angle or Snell's cone.




Figure 6.21: Illustration of reduced skyfraction in the presence (left) and absence (right) of object effects. The curved arrow illustrates the angular height of the obstruction.

The Fate of Rays that Undergo Total Internal Reflection

The qualitative predictions concerning bounce counts between 0 and 3 (in the absence of a water volume) are still valid after submersion, with the additional consideration that rays beyond Snell's cone, which undergo total internal reflection, can contribute to sensor reaching radiance through multiple interface reflections as illustrated in Figure 6.22, assuming the GRS model. Although this is true, additional bounce contributions will show diminishing returns, and as such will not be explored in further detail here.




Figure 6.22: Rays that experience total internal reflection can contribute to sensor reaching radiance, although with diminishing returns.

Analytical Prediction of Interface Effects

The interface effects can be modeled analytically (and in DIRSIG) by forcing the in-water optical properties to be equal to that of air (see Table 6.2), except for a variation in refractive index. For comparative purposes Figure 6.23 illustrates the difference in sensor reaching radiance for a ground plane and 20m tall box ($r_g = 1.0$, $r_b = 0.7$) in the absence of water, and when submerged in a 30m deep water volume. The plot is a function of all aforementioned refractive induced contributors to sensor reaching radiance.

Table 6.2: Skyfraction and background reflection following submersion.

Object Geometry	Material Properties	Illumination	Atmosphere	Image Capture
Ground Plane, Box	$r_g = 1.0$	Uniform Downwelling	Simple Atmosphere. Transmission = 1.0	Vis-Panchromatic 0.4-0.6 μm Bounce Count = 1, 2, 3
Water Depth =	$r_b =$, variable			
30 or 7 m,	$n_a = 1.003$			
Box Height =	$n_w = 1.340$			
20 or 5 m	$\sigma_s = \sigma_a = 0.000$			

r_g is the ground reflectance, r_b is the box reflectance.

n_a is the refractive index of the air, n_w is the refractive index of the water volume.

Radiometry Solvers: Fresnel, Generic, SurfacePM and MediumPM.

The analytically predicted sensor reaching radiance as a function of the refractive interface and skyfraction effects was predicted using $L_{transmitted} + L_{reflected}$ where the *reflected* radiance is shown in Eq. (4.4), and the *transmitted* radiance is equal to $L_{skyfraction}$ as shown in Eq. (6.2), modified by Fresnel transmittance, and where the skyfraction shown in Eq. (6.1) is modified to represent (at most) Snell's cone.

Similarly, the analytical prediction for $r_b = 0.0$ and $r_b = 0.3$ is shown in Figure 6.24 ($r_g = 1.0$) for a 30m water volume with a 20m tall box (top). Since coastal waters are shallow, the scenario was repeated for a 7m deep water volume with a 5m high box [Figure 6.24 (bottom)].






Figure 6.23: Predicted sensor reaching radiance [$W m^{-2} sr^{-1}$] as a function of skyfraction in the absence (dotted line) and presence (dashed line) of a transmissive interface with $n_w = 1.34$ and $r_b = 0.7$ for a 30m deep water volume and a 20m box height.



(a.) Deep water volume.



(b.) Shallow water volume.

Figure 6.24: Analytically predicted sensor reaching radiance [$W m^{-2} sr^{-1}$] as a function of skyfraction for $r_b = 0.0$ (squares), $r_b = 0.3$ (dotted) and $r_g = 1.0$. The top plot (a.) shows a water depth of 30m and a box with a height of 20m. The bottom plot (b.) shows a water depth of 7m, and a box with a height of 5m.

6.3.2 Results

The analytical prediction of sensor reaching radiance using Eqs. (4.4), (6.2) (modified to account for total internal reflection) and (6.3) was compared to the sensor reaching radiance modeled by SRS and GRS, for $r_g = 1.0$ (30m deep water volume) and $r_b = 0.7$ (20m tall box) with $n_{water} = 1.34$ and $n_{air} = 1.003$. The comparison is plotted in Figure 6.25. The dashed line represents the analytical prediction, the diamonds depict the GRS results, while both the dotted and solid lines depict SRS results, but for a search radius of 1.0 and 0.3 units, respectively. Although the larger search radius shows less variance in the data (dotted line), the small search radius (solid line) shows reduced bias at all locations where horizontal and vertical facets meet (approximate pixel values of 35 and 85).




Figure 6.25: Analytically predicted (dashed line) sensor reaching radiance as a function of skyfraction after submersion for $r_b = 0.7$ and $r_g = 1.0$ versus SRS (SSR=0.3) (solid line), SRS (SSR = 1.0) (dotted line) and GRS (triangles).

The analytical prediction appears to overestimate the radiance return from the vertical facet of the box, although both GRS and SRS results converge. It would seem that the analytical prediction fails to model one or more additional numerical factors that contribute to the observed radiance reflected from the side of the box. In truth, there are any number of reasonable explanations for this discrepancy. One possibility is offered in Figure 6.26. This illustration shows the impact of using a bounded water volume. Some rays (assuming GRS) may navigate the vertical edges of the water volume's geometry and strike the ground plane. If so, these rays will not contribute to sensor reaching radiance

for this restricted bounce count, and one might expect that the numerically computed radiance would be less than that assumed analytically.




Figure 6.26: One possible explanation for the observed discrepancy between the analytical prediction of sensor reaching radiance, and that computed analytically, as shown in Figure 6.25.

The radiance comparison was repeated but for a bounce count of 2, allowing for background reflections onto the target of interest, in the form of reflections of the skyfraction of neighboring geometry. The SRS versus GRS results are shown in Figure 6.27. Again, SRS fails to model an increase in sensor reaching radiance on the ground plane, although a background reflected contribution on the vertical facets of the box is apparent. The analytical prediction now underestimates the observed radiance reflected from the side of the box, although the GRS and SRS results converge. Clearly, the numerical results include something not considered within the analytical prediction. This is not considered a failure, but rather the difficulty associated with modeling a numerical process using an analytical prediction. Instead, the focus is placed on the convergence between GRS and SRS, which shows that two fundamentally different numerical methods are capable of producing nearly consistent results (barring the bias introduced by SRS's search radius).




Figure 6.27: Analytically predicted sensor reaching radiance after submersion (dashed line), versus GRS (triangles) and SRS (dotted line) for skyfraction and background reflections (up to a total of two bounces) for $r_b = 0.7$ and $r_g = 1.0$.

To summarize the results of this chapter thus far, the GRS and SRS results are in fairly good agreement with each other, and the analytical prediction for (1.) skyfraction with and without a boundary interface, and (2.) for background reflections as modeled by GRS for bounce counts up to the maximum tested (bounce = 3). The SRS results (circa 2009) lack a background reflected contribution from the box and onto the shadowed ground plane. However, an unknown factor has caused recent SRS solutions to diverge more strongly from both the analytical prediction and the GRS results concerning background reflected contributions to sensor reaching radiance, as shown in Figure 6.16. Since the remainder of this chapter consists of research conducted in 2009, the results do not reflect the more recent SRS solution. Instead, the very promising results at the time of the initial investigation were considered sufficient to warrant continued study concerning in-water attenuation and scattering, as discussed in the following sections.

6.4 Comparison of GRS, SRS and Analytical Predictions for Skyfraction & Attenuation in the Presence of a Water Volume

6.4.1 Attenuation as a Function of Absorption

The attenuation of radiance during transit through the water volume is computed by the MediumPM Radiometry Solver. Regardless of whether the SRS or GRS Radiometry Solver is employed, the attenuation between the *detector and primary hit location* is determined based on the computed optical depth during the traditional second pass raytracing phase. This can be expressed analytically using the trigonometric path between the detector element and the hit location, after passage through the transmissive interface.

For GRS, the BRDF of the material intersected by the primary hit will be importance mapped, thereby generating importance samples. DIRSIG proceeds to numerically sample both the optical depth along various slanted paths between the primary hit location and sources within the scene (determined by the importance mapping of the material’s BRDF), and the Fresnel transmittance of the ray, based on the angle of incidence between the importance sample and the water-air interface. To mimic this analytically, a “universal ray” will be computed that represents the combined average extinction along different path lengths within the water volume and the Fresnel transmittance encountered at the boundary interface (see Figure 6.28).

With SRS, the attenuation between the primary hit location and sources within the scene has already been accounted for during the first pass photon mapping stage of rendering. During the first pass, a photon’s fate is based on the probability of absorption and scattering. If absorbed, the terminated photon history is usually followed by generation of a new photon for continued photon tracing. Since the generation of new photons increments the photon count, and since the flux associated with the photons in a map is determined based on the number of photons generated at sources during the first pass, then absorbed photons serve to decrease the flux associated with all photons in a given map.

6.4.2 Analytical Prediction

The analytical results can be computed using Eq. (6.4) where E_h is the total downwelling illumination, r_f is the Fresnel reflectance, T_f is the Fresnel transmittance, r is the object reflectance, F is the skyfraction (exact calculation but shown as F for brevity in expression), σ_a is the absorption coefficient, z is the water depth, θ_i is the angle of incidence at the air-water interface, and τ_u is the “universal ray” that accounts for Fresnel transmittance and attenuation during the first leg of the trip through the water volume, or from sources into the scene.




Figure 6.28: Universal average ray used to predict the attenuation and Fresnel transmittance effects along the first leg of transit within the water volume, or from sources into the scene.

A more generalized expression can be obtained if σ_a in Eq. (6.4) is replaced by the extinction coefficient, or $\sigma = \sigma_a + \sigma_s$. Using the total extinction coefficient in the analytical approximation, the predicted sensor reaching radiance as a function of refraction at the interface and attenuation in the water volume (total depth of 7m) is shown in Figure 6.29 for $r_g = 1.0$ and box reflectance (vertical height of 5m), and IOPs as follows: $r_b = 0.0/\sigma_s = 0.01/\sigma_a = 0.01$, $r_b = 0.3/\sigma_s = 0.03/\sigma_a = 0.01$, $r_b = 0.3/\sigma_s = 0.1/\sigma_a = 0.01$ and $r_b = 0.3/\sigma_s = 0.2/\sigma_a = 0.01$.

$$\begin{aligned}
 L &= L_{transmitted} + L_{reflected} \\
 L_{reflected} &= \frac{E_h}{\pi} r_f \\
 L_{transmitted} &= F E_h \frac{r}{\pi} \underbrace{T_f \exp \left[-\sigma_a \left(\frac{z}{\cos(\theta_i)} \right) \right]}_{\text{leg two}} \underbrace{T_{f_{ave}} \exp \left[-\sigma_a \left(\frac{z}{\cos(\theta_i)} \right) \right]_{ave}}_{\tau_u = \text{leg one}}
 \end{aligned}$$

Absorption Only

$$L_{transmitted} = F E_h \frac{r}{\pi} \underbrace{T_f \exp \left[-\sigma \left(\frac{z}{\cos(\theta_i)} \right) \right]}_{\text{leg two}} \underbrace{T_{f_{ave}} \exp \left[-\sigma \left(\frac{z}{\cos(\theta_i)} \right) \right]}_{\tau_u = \text{leg one}} \Big|_{ave} \quad (6.4)$$




Figure 6.29: Predicted sensor reaching radiance [$W m^{-2} sr^{-1}$] as a function of refraction/skyfraction and attenuation for $r_g = 1.0$, a 7m water volume and a 5m box height, with box reflectance/IOPS as follows: $r_b = 0.0/\sigma_s = 0.01/\sigma_a = 0.01$ (dashed-dotted), $r_b = 0.3/\sigma_s = 0.03/\sigma_a = 0.01$ (dashed), $r_b = 0.3/\sigma_s = 0.1/\sigma_a = 0.01$ (squares) and $r_b = 0.3/\sigma_s = 0.2/\sigma_a = 0.01$ (dotted).

6.4.3 Results

Figure 6.30 compares the analytical prediction of sensor reaching radiance using Eq. (6.4) for $\sigma_a = 0.1$ versus that modeled using GRS and SRS. Not only do the GRS and SRS results converge, they are also a close match with the analytical approximation. The small dip in DIRSIG modeled radiance near pixel location 300 has been raised with the DIRSIG development team. It was first believed to be a product of the computation of Fresnel reflection for normal incidence (which differs from that used at all other angles), but no apparent explanation for the deviation could be found based upon the DIRSIG computation. A member of the DIRSIG team now believes the error is part of the skydome configuration, and not related to any aspect of in-water radiative transfer.




Figure 6.30: Analytically predicted (dashed line), SRS modeled (dotted line) and GRS modeled (diamonds) sensor reaching radiance as a function of refraction/skyfraction and attenuation for $r_g = 1.0$, a 7m water volume and a 5m box height with $r_b = 0.3$, $\sigma_s = 0.0$, and $\sigma_a = 0.1$.

Since this section considers extinction that is solely a function of absorption, the next section will consider extinction that results from the combination of both absorption and scattering.

6.4.4 Attenuation as a Function of Absorption & Scattering - Using the Full Scattering Coefficient

If σ_a is replaced by $\sigma = \sigma_a + \sigma_s$, then a GRS attenuation-only simulation should continue to show good agreement with the analytical approximation shown in Eq. (6.4). Conversely, GRS and SRS (and therefore the SRS and analytical prediction) will begin to diverge. This is because the surface photon map cannot be reduced to description as an attenuation-only run when both scattering and absorption within the volume are probable. The computational mechanism within SRS not only precludes this possibility, but instead tries to account for the direction in which the energy has been scattered, rather than assuming that all scattered radiance is equivalent to lost radiance.

Referring to Figure 6.31, the top row represents SRS processes, while the bottom row represents GRS processes. The left-most column represents the first pass photon propagation step, which is absent when using GRS for an attenuation-only run (void of volumetric in-scattering computations). Therefore, GRS in the *absence of the volumetric photon map* generates sensor reaching radiance as a product solely of the second pass, where important samples are *attenuated due to absorption and scattering while in the volume*. Note that attenuation, as computed when using GRS, is a function of the absorption and scattering coefficients, and that there is no regard for the scattering phase function. This means that if a photon is scattered, it is considered lost radiance, even if the photon is perfectly forward scattered and therefore reincorporated into the propagating beam. Since a disabled volume photon map, when used in conjunction with GRS, produces results that are a function of skyfraction, the interface and the volumetric absorption and scattering coefficients, without regard for the distribution of the scattered photons, GRS attenuation-only runs are amenable to description using Eq. (6.4).

The equivalent SRS simulation is shown in the top row of Figure 6.31. During the first pass photon propagation step the photon maps will be populated using volumetric scattering events and surface reflection events. However, the scattering coefficient, and therefore the probability of an actual scattering event, does not directly decrease the flux of photons in a map since ray re-direction (in and of itself) does not require the generation of new photons at sources within the scene.

[Recall: the flux associated with each photon is determined by taking the integrated flux of all sources, and dividing it by the number of photons needed to fill a map. If using a combined volume and surface map, the denominator (division by number of photons) can be different for each map. When a given map (volume or surface) is full, the number of photons shot (while filling both maps) is pulled and used to determine the flux associated with the map that has just been closed. If additional photons are required to fill another map (volume or surface), then the flux in the yet-to-be-filled map will be less due to division by a larger number of shot photons].

Therefore, provided the maps are not bounce limited, the presence or absence of a scattering event (which is a function of the scattering coefficient) does not have a direct impact on the flux of photons in the map. Instead, it is the combination of the *scattering coefficient and phase function* that determines how scattering will impact the flux associated with photons in the surface map. The SRS modeled results remain *dependent* on the volumetric absorption and scattering properties, including the coefficients and the scattering phase function, even when the computation of in-scattered volumetric radiance has been disabled. These additional considerations are not represented by Eq. (6.4), and therefore SRS modeled attenuation-only results will diverge from the analytical prediction if $\sigma_s \neq 0.0$.




Figure 6.31: Comparison of SurfacePM (top) and Generic (bottom) to compute sensor reaching radiance. The left-most column is the first pass, and the right-most column illustrates sensor reaching radiance as a function of attenuation while in the volume.

In fact, if Eq. (6.4) is augmented so that the second leg of the trip uses σ_a instead of the extinction coefficient (despite the fact that both absorption and scattering are taking place), then a modest analytical prediction that better describes SRS attenuation-only computations can be formulated. This is shown in Eq. (6.5).

$$\begin{aligned}
L &= L_{transmitted} + L_{reflected} \\
L_{reflected} &= \frac{E_h}{\pi} r_f \\
L_{transmitted} &= F E_h \frac{r}{\pi} \overbrace{T_f \exp \left[-\sigma \left(\frac{z}{\cos(\theta_i)} \right) \right]}^{\text{Absorption + Scattering}} \underbrace{T_{f_{ave}} \exp \left[-\sigma_a \left(\frac{z}{\cos(\theta_i)} \right) \right]}_{\text{leg one} \quad \text{leg two}} \bigg|_{ave} \quad (6.5)
\end{aligned}$$

Eq. (6.5) actually suggests that the first-pass events are characterized by an attenuation that is equal to $\exp[-\sigma_a z]$ rather than $\exp[-(\sigma_a + \sigma_s)z]$. This implies that none of the scattered radiance during the first pass is equated as lost radiance, which of course, is an extreme prediction, and yet reasonable for highly forward peaked phase functions and short optical depths. However, it is conceivable that at least some of the scattered radiance will be lost for all other phase functions, although the degree to which is clearly not amenable to analytical approximation.

Of greater interest is that SRS attenuation-only simulations generate two major postulates concerning the use of SRS. First, they show that scattering decrements the radiance in an indirect manner based on radiance re-distribution. This also means that the scattering phase function *can* impact the flux associated with photons in a given surface map, and therefore the sensor reaching radiance. To illustrate, consider a single ray (*e.g.* laser or solar beam) illuminating a submerged surface. If the volume is a perfect forward scatterer, then scattering events between source and scene end up being re-incorporated back into the beam rather than lost during the scattering process. The surface photon map should fill relatively efficiently, and therefore each photon will be associated with greater flux. It is further postulated that a perfect forward scattering phase function will give rise to SRS results that are *primarily a function of absorption events during all first pass activity*, since all scattering events are directly re-incorporated into the propagating beam. Conversely, if the medium can be described by a phase function that perfectly backscatters, then scattering events between source and scene are all re-directed back toward the source. This decreases the flux associated with the photons on the surface, since more need to be cast out from sources in order to fill the surface map, as illustrated in Figure 6.32.

For the more moderate case of a uniform scattering phase function, the downwelling photons within the beam will be evenly scattered in all directions and eventually accumu-




Figure 6.32: Illustration of how the scattering phase function impacts the flux associated with photons in surface maps. Solid lines depict source-scene photons that are absorbed, dotted lines depict source-scene photons that are scattered, and dashed lines depict direct source-scene photons. Sensor-scene rays are not illustrated.

late at many submerged locations within reasonable proximity to the ray. Furthermore, photons that enter the interface within the Snell window can, in fact, be scattered out of the Snell cone, decreasing the sensor reaching radiance within the Snell cone, but increasing it beyond the edges of the cone. All other things being equal, average cosines in between extreme forward and backward peaked scattering should, intuitively, generate maps with intermediate flux characteristics.

This means that an invariant scattering (and absorption) coefficient *can* produce *variable* sensor reaching radiance as computed by SRS, when associated with different scattering phase functions, provided the map is not bounce-limited. The degree to which this is true will be a function of the extinction coefficient and optical depth, since dependency will be lost (or minimal) if the photon map is dominated by direct radiance.

6.4.5 Results

To test the hypotheses presented in the preceding section, the GRS and SRS models will be compared for attenuation-only simulations. In addition, the SRS model will be compared with a constant scattering and absorption coefficient, but different phase functions while the results for more moderate (*e.g.* less peaked and/or more uniform) phase functions can then be reasonably deduced.

Figure 6.33 compares GRS sensor reaching radiance for an attenuation-only simulation versus the analytical prediction shown in Eq. (6.4). Except at high extinction, the predicted results are slightly greater than that generated by DIRSIG. This deviation is believed to be largely the result of using an analytical “universal ray” to predict attenuation along the second leg of transit through the water volume.




Figure 6.33: Analytically predicted versus GRS generated sensor reaching radiance as a function of refractive induced skyfraction and attenuation for $r_b = 0$, $r_g = 1.0$, a water depth of 30m, and a box height of 20m. Legend: $\sigma_s = 0.01/\sigma_a = 0.01$ (dashed-dotted), $\sigma_s = 0.02/\sigma_a = 0.01$ (dashed), $\sigma_s = 0.03/\sigma_a = 0.01$ (squares), and $\sigma_s = 0.05/\sigma_a = 0.05$ (dotted). The GRS generated results for the equivalent IOPs are plotted with asterisks. *The dip in radiance at pixel location 300 is believed to be a function of the simple atmosphere used during this simulation.*

For a shallower scenario (7m deep water volume and 5m box height), Figure 6.29 illustrates the predicted sensor reaching radiance as a function of refractive and attenuation effects. The analytical prediction versus GRS generated sensor reaching radiance is shown in Figure 6.34 for $r_g = 1.0$ and box reflectance/IOPs as follows: $r_b = 0.0/\sigma_s = 0.01/\sigma_a = 0.01$, $r_b = 0.3/\sigma_s = 0.03/\sigma_a = 0.01$, $r_b = 0.3/\sigma_s = 0.1/\sigma_a = 0.01$ and $r_b = 0.3/\sigma_s = 0.2/\sigma_a = 0.01$.




Figure 6.34: Predicted sensor reaching radiance versus GRS as a function of refraction/skyfraction and attenuation for $r_g = 1.0$, a 7m water volume and a 5m box height, with box reflectance/IOPS as follows: $r_b = 0.0/\sigma_s = 0.01/\sigma_a = 0.01$ (dashed-dotted), $r_b = 0.3/\sigma_s = 0.03/\sigma_a = 0.01$ (dashed), $r_b = 0.3/\sigma_s = 0.1/\sigma_a = 0.01$ (squares) and $r_b = 0.3/\sigma_s = 0.2/\sigma_a = 0.01$ (dotted). GRS results are shown as asterisks. *The dip in radiance at pixel location 300 is believed to be a function of the simple atmosphere used during this simulation.*

Figure 6.35 plots the analytical prediction for water leaving radiance for $\sigma_s = 0.02/\sigma_a = 0.01$ and $\sigma_s = 0.03/\sigma_a = 0.01$ as compared to GRS and SRS for a 30m deep water volume and 20m box height. The sensor reaching radiance computed using SRS is substantially higher than that obtained using GRS for both sets of IOPs. As hypothesized, the SRS results diverge from those modeled by GRS, and that predicted analytically using Eq. (6.4). This is not to suggest that the SRS modeled results are incorrect, but rather illustrate that the computational mechanism used by SRS is not amenable to description using the analytical approximation shown in Eq. (6.4), or by GRS, since both equate scattered radiance to lost radiance, without regard for the scattered distribution of the radiance. However, the user can force SRS to converge with the analytical prediction (and GRS) if the photon mapping process is bounce limited. For example, if the maximum number of possible scattering events per shot photon is reduced to one, then each scattering event does indeed equate with lost radiance since a new photon's history must be traced for each scattering event. This is illustrated in Figure 6.36 which shows the GRS results (dotted line) and the SRS results (solid line, max bounce of 1). The results diverge for pixels on the box only by virtue of reflection (*e.g.* the GRS simulation used $r_b = 0.3$ while the SRS simulation used $r_b = 0.0$).




Figure 6.35: Analytical (solid line), GRS (squares), SRS (dotted line), $r_b = 0.0$, $r_g = 1.0$, 30m deep water volume and 20m box height. For each set of two (*e.g.* two solid lines, two dotted lines and two lines denoted by squares), the lower represents IOPS of $\sigma_s = 0.03/\sigma_a = 0.01$, and the higher represents IOPs of $\sigma_s = 0.02/\sigma_a = 0.01$. SRS variables: $g = 0$, SSR = 1.0, 8M photons, max bounce count = 50. GRS variables: $\mu = 10$, $\phi = 20$, importance = 900, decay = 5, bounce count (irrelevant beyond 1), quick samples = 1. *The dip in radiance at pixel location 300 is believed to be a function of the simple atmosphere used during this simulation.*




Figure 6.36: GRS (dotted line) $r_b = 0.3$, $r_g = 1.0$, 7m deep water volume and 5m box height, $\sigma_s = 0.1/\sigma_a = 0.01$. SRS (solid line) $r_b = 0.0$, $r_g = 1.0$, $\sigma_s = 0.1/\sigma_a = 0.01$ and bounce limited (max bounce of 1). *The dip in radiance at pixel location 300 is believed to be a function of the simple atmosphere used during this simulation.*

Since the results shown in Figure 6.35 were generated using an average cosine of $g = 0.0$ (max of 50 bounces), or uniform scattering, selected results for $r_b = 0.3$ and $\sigma_s = 0.2/\sigma_a = 0.01$ were repeated using forward and backscattering phase functions ($g = 0.99$, $g = 0.80$, $g = -0.80$ and $g = -0.99$) (max of 50 bounces), on the assumption that it is not truly possible to evaluate SRS for attenuation only effects in the absence of phase function effects when the simulation is not bounce limited. Figure 6.37 depicts the results of SRS for $r_b = 0.3$ and $\sigma_s = 0.2/\sigma_a = 0.01$ when $g = 0.0$ (triangles), $g = 0.99$ and 0.80 (squares) and $g = -0.99$ and $g = -0.80$ (dotted line) (Figure 6.38 shows the scattered photons (in the volume) and those reflected from the surface for this particular scenario. It is evident that *more than enough* photons have been generated to populate both maps). The plots indicate that the scattering phase function *can* have a major impact on the radiance distributed on submerged surfaces. This impact is greatest for the extreme phase function of $g = 0.99$. For the curves that hover near $0.01 \text{ W/m}^2 \text{ sr}$, the impact of the phase function is much more subtle, suggesting that the dominant contributor is direct radiance for the associated coefficients. [For reference, Figure 6.39 shows the Henyey-Greenstein phase function with variable average cosines.]




Figure 6.37: SRS $g = -0.80$ and -0.99 (dotted line), SRS $g = 0.0$ (triangles) and SRS $g = 0.80$ and 0.99 (squares) for $r_b = 0.3$, $r_g = 1.0$, 7m deep water volume and 5m box height, $\sigma_s = 0.2/\sigma_a = 0.01$. SRS variables: SSR = 1.0, 8M photons, max bounce count = 50.

If the optical depth or coefficients are increased, the impact of the phase function is more dramatic. This is illustrated in Figure 6.40 for backscattering $g = -0.80$ (dotted




Figure 6.38: Right: scattered photons in the volume. Left: Photons reflected from the surface. These plots show that *more than enough* photons have been used to describe the radiance distribution for this scene, and that deviations between the DIRSIG model and our expectations are not a function of a photon-starved simulation.




Figure 6.39: Henyey-Greenstein phase function with variable average cosines.

line), uniform $g = 0.0$ (triangles) and forward scattering $g = 0.80$ (square) phase functions with $\sigma_s = 0.8$ ($\sigma_a = 0.0$) (note that the plotted results do not reflect extinction during the




Figure 6.40: SRS $g = -0.80$ (dotted line), SRS $g = 0.0$ (triangles) and SRS $g = 0.80$ (squares) for $r_b = 0.3$, $r_g = 1.0$, 7m deep water volume and 5m box height, $\sigma_s = 0.8/\sigma_a = 0.0$. SRS variables: SSR = 1.0, 8M photons, max bounce count = 50.

second pass, which would lower the magnitude of all curves).

In summary, GRS solutions in the absence of volumetric in-scattering can be easily modeled using an analytical prediction that includes total extinction as a function of optical depth. This is not true with the SRS model since first pass volumetric scattering events directly influence the sensor reaching radiance, even if volumetric in-scattering in the second pass is excluded in the final computation. Provided the bounce count is not restricted, the SRS computed sensor reaching radiance reflected from surfaces submerged in attenuation-only volumes will therefore always predict a greater radiance solution than that modeled using GRS. Moreover, the SRS result can show a dependence on the scattering phase function, which is more pronounced at longer optical depths. The next section extends these results to include in-scattering within the volume.

6.4.6 In-Scattering

Since an attenuation-only simulation does not equate with real-world physics, the divergence of GRS and SRS for attenuation-only simulations is not a cause for alarm, *provided the computational models re-convergence when volumetric scattering is enabled*. To test this, GRS and SRS results were compared with volumetric scattering disabled and enabled, using $\sigma_s = 0.1/\sigma_a = 0.01$. The results are shown in Figure 6.41, where the closely paired lines show the difference between volumetric scattering disabled and enabled (*e.g.* solid line depicts SRS with volumetric scattering enabled, and squares show SRS with volumetric scattering disabled, while the dotted line depicts GRS with volumetric scattering enabled, and the asterisks show GRS with volumetric scattering disabled).




Figure 6.41: SRS without in-scattered radiance (squares), with in-scattered radiance (solid line), GRS without in-scattered radiance (asterisks), and with in-scattered radiance (dotted line). Other variables: $r_b = 0.3$, $r_g = 1.0$, 7m deep water volume and 5m box height, $\sigma_s = 0.1/\sigma_a = 0.01$, SSR = 1.0, max bounce count = 50, 8 million photons, 900 importance samples, $\mu = 10$, $\phi = 20$, bounce count = 3, decay rate = 5, quick samples = 1.

The continued divergence of GRS and SRS with volumetric scattering enabled is troubling. At least one Radiometry Solver, if not both computational mechanisms, fails to adequately describe *scattering in the presence of a highly reflective ground plane*. The significance of this cannot be over-stated. DIRSIG was formerly validated (as per the canonical problems described by (Mobley et al. 1993)) and shown to be in very good

agreement with other numerical and plane parallel models of radiative transfer. However, the canonical problems are dominated by infinitely deep volumes, which disregard the contribution from in-water reflections. Since DIRSIG is to be used in the future to simulate coastal ocean water systems, characterized by shallow water volumes and significant returns from reflection from in-water objects, validation of the combination of volumetric scattering and significant in-water reflective contributions is vital to future modeling interests. Therefore, it is important to investigate which radiometric computational model (GRS and SRS) is most likely deficient in its description of the physics, and why.

One sampling problem that is shared by both SRS and GRS is the sensor-target ray's ability to accurately compute backscattered radiance (*e.g.* radiance scattered toward the detector). To illustrate, consider the intuitive situation of a high degree of backscattering. One would expect this to be associated with a significant increase in sensor reaching radiance as compared to other average cosines (all things being considered, but ignoring the interplay of optical depth, turbidity of the water and the reflectance of submerged objects which can alter this assumption). However, a highly peaked phase function generates a difficult sampling problem since it is unlikely that a sensor-target ray will travel the exact same path as a source-target ray, therefore when the photons in the map are weighed by the scattering phase function, very few of the source-target photons will be perfectly scattered into the sensor's line of sight. This is illustrated in Figure 6.42. This means that the second pass computation between source and target can suffer from sampling problems. As such, the user may not obtain the intuitive result if modeling highly peaked phase functions, although more modest peaked phase functions should not suffer appreciably.




Figure 6.42: Illustration of second pass events coupled with a highly peaked (backward) phase function. The dotted lines are traces of rays used to generate the map during first pass events (or source-scene rays). The solid line is the sensor-scene ray. When the photons are weighed by an extremely peaked scattering phase function, few to none will be backscattered within the line of sight of the sensor.

Further inspection of the computational mechanism of GRS suggests that the same type of sampling problem will persist for all of the importance samples used by the Generic Radiometry Solver. Figure 6.43 is an extension of Figure 6.31. If the volumetric map *is enabled* during the second pass (right-most column), then the sensor reaching radiance as computed by SRS (top row) is a function of the photon population *on the target surface of interest, which is a function of absorption and scattering during the first pass, attenuation by absorption and scattering during the second pass and the volumetric in-scattered radiance between the target and detector*. If the volumetric map *is enabled* when using GRS (bottom row), the sensor reaching radiance is a function of *attenuation by absorption and scattering, and the addition of volumetric in-scattered radiance, between the target and detector and during all second pass importance sampling events*.

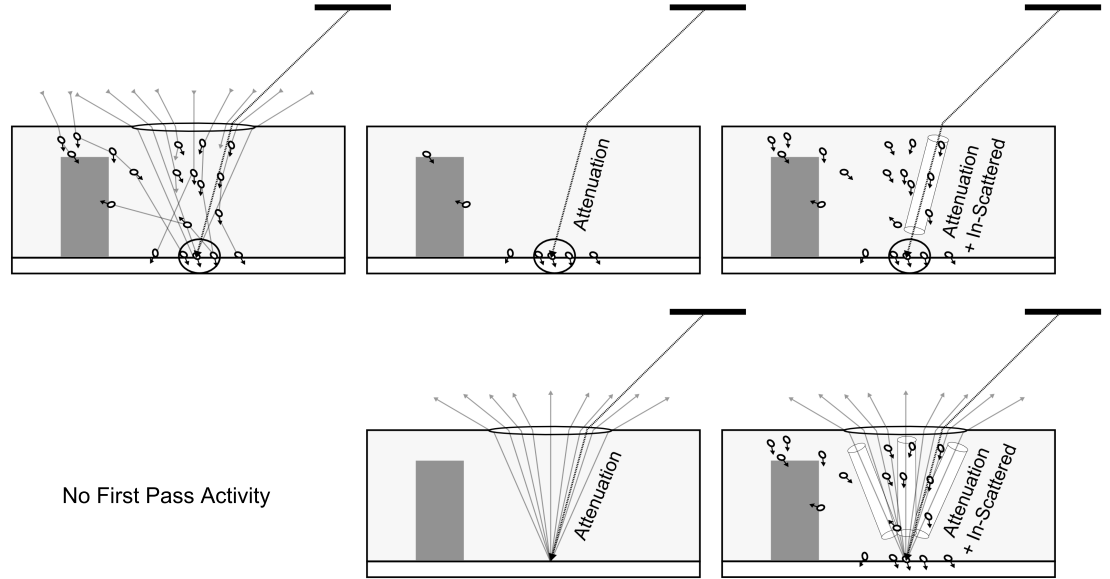


Figure 6.43: Comparison of SurfacePM (top) and Generic (bottom) to compute sensor reaching radiance. The left-most column is the first pass. The middle column illustrates sensor reaching radiance as a function of attenuation-only while in the volume, and the right-most column illustrates sensor reaching radiance as a function of attenuation and in-scattering.

Therefore, if SRS and GRS are to mimic real-world physics, the second pass radiance collection between targets and the detector must overcome the sampling problem associated with high backscatter (although it is acknowledged that water is dominated by a high degree of forward scatter), and for SRS and GRS to re-converge, the importance samples

used by GRS must replace first pass photon mapping events. For water, the latter issue is much more important since the problem depicted in Figure 6.42 for backscattering will be equally problematic when GRS carries out importance sampling with a highly forward peaked phase function. The more extreme the forward scattering, the less likely it is that a scene-to-source importance sample will follow a path in the volume that perfectly matches a source-to-scene first pass photon mapping path. Therefore, each importance sample will be decremented by the extinction coefficient, but may fail to collect the copious amounts of forward scattered downwelling irradiance (unless an infinite number of importance samples are collected). This is illustrated in Figure 6.44. The result is sensor reaching radiance that is lower than expected.




Figure 6.44: Illustration of first and second pass events coupled with a highly forward peaked phase function. The dotted lines are traces of rays used to generate the volume map during the first pass (or source-scene rays). The solid line is the sensor-scene ray and the corresponding importance samples. When the photons are weighed by an extremely peaked scattering phase function, fewer than expected will be collected as a result of discrete sampling.

A possible solution to this sampling problem is to recast highly forward scattered radiance as direct radiance. The computational details of this solution and its impact on sensor reaching radiance for GRS and SRS are discussed in the next section.

6.4.7 The Effective Scattering Coefficient

In order to assist the Generic Radiometry Solver when rendering participating media with volumetric scattering, the following changes were made (Goodenough 2009).

- All second pass events will be decremented by an extinction coefficient that is equal to the sum of the absorption coefficient and an effective scattering coefficient or $\sigma = \sigma_a + \sigma'_s$. The effective scattering coefficient is equal to the full scattering coefficient, less the portion that is forward scattered into a small solid angle about a ray that propagates between source and target. The result is to make highly forward scattered radiance equivalent to direct radiance.
- If σ'_s takes care of the forward scattered radiance, then to ensure that it is not doubly counted, the *volume* photon map that is generated during the first pass will only store scattering events that are scattered beyond a small solid angle about the propagating ray. The premise is that σ'_s reallocates small angle scattering events to direct radiance, and therefore these events need not be stored by the *volume* photon map.
- Regardless of the type of *volume* scattering event (small or large angle), the *surface* photon map records all photons that eventually reach the surface of interest.

The impact of these changes is several fold:

- Large angle scattering events from the first pass are stored in the volume photon map.
- During the second pass the MediumPM Radiometry Solver will compute volumetric in-scattering from multiply scattered photons that are scattered outside of the “forward” peak.
- The sensor-target second pass rays will be attenuated by an optical depth determined using an effective scattering coefficient for both SRS and GRS.
- GRS importance samples will be attenuated by an optical depth determined using an effective scattering coefficient.
- The overall result will be an increase in sensor reaching radiance as computed by GRS when combined with a scattering medium.
- SRS results should show little change since the surface map stores all photons (scattered through large or small angles), and the second pass restores the uncollected small scattering events within the volume via the effective scattering coefficient.

6.4.8 Results: Attenuation - SRS versus GRS

The first step is to determine how the effective scattering coefficient impacts attenuation-only simulations. Figure 6.45 compares sensor reaching radiance using an effective scattering coefficient for an attenuation-only run as computed by GRS and SRS. The SRS results are shown for $g = 0.99$ and $g = 0.80$ (greatest and lowest magnitude asterisk curves, respectively), $g = 0.00$ (triangles) and $g = -0.99$ (squares). These are in relative agreement with the pre-code-change sensor reaching radiance shown in Figure 6.37. The dotted lines depict GRS for $g = 0.99$, $g = 0.80$ and $g = 0.0$ (greatest and two lowest magnitude dotted lines, respectively, where the $g = 0.80$ and $g = 0.0$ curves overlap).

Although it is apparent that the SRS results are always greater in magnitude than the GRS solution, they both exhibit the same relative change and dependency on the scattering phase function. This is very promising since without an effective scattering coefficient, the GRS solution would be unable to model the impact of the scattering phase function on sensor reaching radiance in an attenuation-only simulation. The next section extends this comparison to include in-scattered radiance.




Figure 6.45: GRS $g = 0.99$, $g = 0.80$ and $g = 0.0$ (overlaps with $g = 0.80$) (dotted lines), SRS $g = 0.99$ and $g = 0.80$ (asterisks), SRS $g = 0.00$ (triangles) and SRS $g = -0.99$ (squares) for $r_b = 0.3$, $r_g = 1.0$, 7m deep water volume and 5m box height, $\sigma_s = 0.2/\sigma_a = 0.01$. SRS variables: SSR = 1.0, 8M photons, max bounce count = 50. GRS variables: $\mu = 10$, $\phi = 20$, importance = 900, decay = 5, bounce count 3, quick samples = 1.

6.4.9 Results: In-Scattering - SRS versus GRS

As noted in the previous section, the SRS and GRS solutions follow the same relative trend and scattering phase function dependency for an attenuation-only simulation. The divergence in actual magnitude shown in Figure 6.45 can be explained by the computational mechanisms used by each Radiometry Solver. The next question is whether or not GRS and SRS reconverge when in-scattered radiance is computed. This is illustrated in Figure 6.46 which compares GRS and SRS computed sensor reaching radiance as a function of bottom effects, skyfraction, background reflection, attenuation and volumetric in-scattering for an average cosine of 0.80 (Henye-Greenstein phase function), an absorption coefficient of $\sigma_a = 0.01$, and a scattering coefficient of $\sigma_s = 0.2$. The solid line in Figure 6.46 depicts the SRS results (MediumPM Radiometry Solver (MRS) segment length of 0.1 units), while the squares and triangles show GRS results with MRS segment lengths of 10.0 and 5.0 units, respectively. The results are in good agreement, especially for the smaller segment step-size of 0.5 units.




Figure 6.46: Comparison of sensor reaching radiance due to bottom effects, skyfraction, background reflection, attenuation and in-scattering for GRS (squares; MRS segment length = 10.0), (diamonds; MRS segment length = 5.0) and SRS (solid line; MRS segment length 0.1). GRS: 900 importance, $\mu = 10$, $\phi = 20$, decay = 5, bounce = 3, SRS: SSR = 1.0, MRS: $\sigma_a = 0.01$, $\sigma_s = 0.2$, Henye-Greenstein phase function ($g = 0.80$), core search fraction (CSF) = 1.0 = spectral search radius (SSR).

Figure 6.47 compares GRS and SRS modeled sensor reaching radiance for a forward ($g = 0.80$), backward ($g = -0.80$) and uniform ($g = 0.0$) phase function, with $\sigma_s = 0.2$ and




Figure 6.47: SRS $g = -0.8$ (solid), GRS $g = -0.80$ (dotted), SRS $g = 0.80$ (dashed), GRS $g = 0.80$ (diamonds), SRS $g = 0.0$ (squares), GRS $g = 0.0$ (x's). MRS: 8M photons, SSR = CSR = 1.0, segment length = 0.1 (SRS) and 5.0 (GRS), Henyey-Greenstein phase function, $\sigma_a = 0.01$, $\sigma_s = 0.2$. SRS: SSR = 1.0, 8M photons. GRS: Importance = 900, decay = 5, $\mu = 10$, $\phi = 20$, quick samples = 1.

$\sigma_a = 0.01$. The ground plane has been given a high reflectance of $r_g = 1.0$ to emphasize the reflected bottom effects, although typical sandy bottoms would have reflectance values between 0.2 and 0.6 as per Carder et al. (2003).

The convergence of SRS and GRS in Figures 6.46 and 6.47 is extremely satisfying and serves to validate the results of both computational mechanisms. Although initial comparisons of the raytraced and surface photon mapped solutions highlighted a computational difficulty not formerly considered, the divergence of the GRS and SRS solution has since been remedied by recasting small angle scattering events as direct radiance. Indeed, it is very encouraging to present results that confirm the agreement between the mechanistically different computational modes of a ray-traced and surface photon mapped solution of in-water objects. When taken in combination, the progressively more complicated scenarios that were compared in this chapter serve to validate contributions to sensor reaching radiance that are a function of skyfraction, background reflection, submersion in a water volume, attenuation, and volumetric in-scattering. The following two chapters will compare the internally consistent GRS and SRS solutions against the two-stream analytical prediction, published numerical results, and the 3D-taxing scenario of ship-shadow.

Chapter 7

Absolute Sensor Reaching Radiance (& Irradiance) - Direct Reduced Plus In-Scattered

Goal: Evaluate absolute radiometric fidelity by comparing DIRSIG modeled results for direct reduced and in-scattered radiance versus the analytical two-stream approximation, independent numerical plane-parallel transfer code, and field observations.

Source Function	✓		
Non-Source Function	✓	Model-vs-Model Comparison	✓
Submerged or Floating Objects		Model-vs-Analytical Prediction	✓
Boundary Interface	✓	Model-vs-Observation Comparison	✓
3D Geometry			

This chapter evaluates absolute sensor reaching radiance (and irradiance) including in-scattered contributions as modeled by the MediumPM Radiometry Solver (MRS). The validation compares DIRSIG modeled results to (1.) an analytical prediction, (2.) published numerical results, and (3.) field data (see Table 7.1 for simulation details). The remainder of this chapter is divided into three sections. The first section combines the direct reduced radiance described in the preceding chapter with an estimate of the in-scattered radiance as modeled by the two-stream approximation (for a simple scenario with uniform scattering and downwelling irradiance). The second section compares DIRSIG modeled results to numerical predictions published by Mobley et al. (1993) (or Canonical Problem #6). Finally, the DIRSIG data is compared to data profiles from a field collect.

Table 7.1: Solutions to the full radiative transfer equation.

Analytical Model*	Phase Function	Illumination Conditions	Bottom Boundary	Additional Details
Two-Stream	Uniform	Uniform	Lambertian	Compare to DIRSIG Generic and SurfacePMRadSolver.
Numerical Model	Phase Function	Illumination Conditions	Bottom Boundary	Additional Details
Numerical	Petzold	60° Zenith Sun	Lambertian	Canonical Problem #6 from Mobley et al. (1993).
Field Profiles	Phase Function	Illumination Conditions	Bottom Boundary	Additional Details
HydroRad Data	Petzold & Henyey-Greenstein	60° Zenith Sun	Lambertian	Using the SurfacePM & MediumPMRadSolvers.

7.1 The Two-Stream Approximation

In certain scenarios the two-stream approximation can provide an accurate description of the radiative transfer process and will therefore be used as a “sanity-check” against which DIRSIG’s numerical code can be compared. This is especially true in the presence of uniform skylight (*e.g.* constant downwelling irradiance over the hemisphere such as the influence of an optically thick overlaying cloud), reflective in-water objects, and for simple phase functions (which are two scenarios underrepresented in the aforementioned canonical problems (Mobley et al. 1993), (Goodenough 2007)).

The two-stream approximation is a plane-parallel model based on the RTE derived in Chapter 3. The full RTE describes the light field in terms of specific intensity or radiance $L_v(\vec{k}, \Omega, t, v)$ that is a function of position, direction, time and frequency. However, in plane-parallel problems the specific intensity is reduced to a function of depth and angle, and the angularly dependent half range specific intensities are replaced by their average over each hemisphere. Under this assumption Ishimaru (1978) reports that the two-stream method obtains reasonable agreement with experimental data, and has found extensive use. Despite this, it requires some empirical estimates, and the range of its validity and theoretical basis is not well established (Ishimaru 1978).

The plane-parallel model standard theory is fully explained by Thomas and Stamnes (1999), and is based on the assumption that specific intensity varies with vertical depth, but is invariant over the horizontal plane of interest. This allows the directional vector \vec{k} (see Section 3.2) to be re-expressed by z or depth. Since radiance or specific intensity is only a function of vertical position and direction, and since optical depth increases monotonically with z , τ can be used in place of the vertical coordinate z . Moreover, for time-independent processes $L_v(\vec{k}, \Omega, t, v)$ can be re-expressed as $L_v(z, \Omega)$, where the

frequency dependence (v) is not explicitly stated as a variable. The angular direction can then be redefined using an azimuthal angle ϕ , and a polar angle θ (between Ω and the surface normal \vec{n} of the coordinate system) leading to Eqs. (7.1) and (7.2), where $\mu = \cos \theta$ and \pm signifies direction (Thomas and Stamnes 1999):

$$\mu \frac{dL^+(\tau, \mu, \phi)}{d\tau} = L^+(\tau, \mu, \phi) - S^+(\tau, \mu, \phi) \quad (7.1)$$

$$-\mu \frac{dL^-(\tau, \mu, \phi)}{d\tau} = L^-(\tau, \mu, \phi) - S^-(\tau, \mu, \phi) \quad (7.2)$$

In generalized media the source function includes a *diffuse* term which describes the diffusely or multiply scattered radiance, the possibility of thermal emissions, and a first-order scattering term (single scattering or solar scattering) if external illumination is present. The complete source function including thermal, first order and multiple scattering terms is shown in Eq. (7.3) (reverting to the use of solid angle instead of θ and ϕ for brevity of expression), and using the scattering phase function to describe how radiation traveling in direction Ω' is scattered into direction Ω (e.g. the notation $p(\Omega', \Omega)$ refers to photons moving in direction Ω' and scattered into direction Ω):

$$\begin{aligned} S_v(\tau, \Omega) = & \underbrace{[1 - w_o(v)]B_v(T)}_{\text{thermal emission}} \\ & + \underbrace{S^*(\tau, \Omega)}_{\text{first-order scattering}} \\ & + \underbrace{\frac{w_o}{4\pi} \int_{4\pi} p(\Omega', \Omega) L_d(\Omega') d\Omega'}_{\text{multiple scattering}} \end{aligned} \quad (7.3)$$

The first order scattering term S^* specifically describes the initial scattering event for the incident external source (Thomas and Stamnes 1999). The generalized (full) *diffuse* transfer equation of interest is shown below where the subscript “d” refers to the diffuse field (as opposed to direct reduced radiance).

$$\mu \frac{dL_d(\tau, \mu)}{d\tau} = L_d(\tau, \mu) - \frac{w_o}{4\pi} \int_0^{2\pi} d\phi \int_{-1}^{+1} p(\mu', \mu) L_d(\tau, \mu') d\mu' - S^*(\tau, \mu) \quad (7.4)$$

Assuming the radiative field is independent of azimuth, integration over $\int_0^{2\pi} d\phi$ leaves the following:

$$\mu \frac{dL_d(\tau, \mu)}{d\tau} = L_d(\tau, \mu) - \frac{w_o}{2} \int_{-1}^{+1} p(\mu', \mu) L_d(\tau, \mu') d\mu' - S^*(\tau, \mu) \quad (7.5)$$

Ignoring thermal emission the half range specific intensities are:

$$\begin{aligned} -\mu \frac{dL_d^-(\tau, \mu)}{d\tau} &= L_d^-(\tau, \mu) - \frac{w_o}{2} \int_0^{+1} p(-\mu', -\mu) L_d^-(\tau, \mu') d\mu' \\ &\quad - \frac{w_o}{2} \int_0^{+1} p(\mu', -\mu) L_d^+(\tau, \mu') d\mu' \\ &\quad - \frac{w_o}{4\pi} E_s p(-\mu_o, -\mu) e^{-\frac{\tau}{\mu_o}} \end{aligned} \quad (7.6)$$

$$\begin{aligned} \mu \frac{dL_d^+(\tau, \mu)}{d\tau} &= L_d^+(\tau, \mu) - \frac{w_o}{2} \int_0^{+1} p(-\mu', \mu) L_d^-(\tau, \mu') d\mu' \\ &\quad - \frac{w_o}{2} \int_0^{+1} p(\mu', \mu) L_d^+(\tau, \mu') d\mu' \\ &\quad - \frac{w_o}{4\pi} E_s p(-\mu_o, \mu) e^{-\frac{\tau}{\mu_o}} \end{aligned} \quad (7.7)$$

If the angularly dependent specific intensities $[L_d^\pm(\tau, \mu)]$ are replaced with mean values over each hemisphere $[L_{d_{ave}}^\pm(\tau)]$ then the following result is obtained:

$$\begin{aligned} -\mu \frac{dL_{d_{ave}}^-(\tau)}{d\tau} &= L_{d_{ave}}^-(\tau) - \frac{w_o}{2} \int_0^{+1} p(-\mu', -\mu) L_{d_{ave}}^-(\tau) d\mu' \\ &\quad - \int_0^{+1} p(\mu', -\mu) L_{d_{ave}}^+(\tau) d\mu' \\ &\quad - \frac{w_o}{4\pi} E_s p(-\mu_o, -\mu) e^{-\frac{\tau}{\mu_o}} \end{aligned} \quad (7.8)$$

$$\begin{aligned} \mu \frac{dL_{d_{ave}}^+(\tau, \mu)}{d\tau} &= L_{d_{ave}}^+(\tau) - \frac{w_o}{2} \int_0^{+1} p(-\mu', \mu) L_{d_{ave}}^-(\tau) d\mu' \\ &\quad - \int_0^{+1} p(\mu', \mu) L_{d_{ave}}^+(\tau) d\mu' \\ &\quad - \frac{w_o}{4\pi} E_s p(-\mu_o, \mu) e^{-\frac{\tau}{\mu_o}} \end{aligned} \quad (7.9)$$

Since the average values are no longer a function of μ they can be removed from the integrals. Furthermore, the *backscattering coefficient* can be introduced as follows, using

the reciprocity of $p(-\mu', \mu) = p(\mu, -\mu')$ and $p(-\mu', -\mu) = p(\mu, \mu')$ and where $(1 - b)$ is equivalent to the *forward scattering coefficient* (Thomas and Stamnes 1999):

$$\begin{aligned} b(\mu) &= \frac{1}{2} \int_0^{+1} p(-\mu', \mu) d\mu' = \frac{1}{2} \int_0^{+1} p(\mu', -\mu) d\mu' \\ b &= \int_0^{+1} b(\mu) d\mu = \frac{1}{2} \int_0^{+1} d\mu \int_0^{+1} p(-\mu', \mu) d\mu' \\ &= \frac{1}{2} \int_0^{+1} d\mu \int_0^{+1} p(\mu', -\mu) d\mu' \end{aligned} \quad (7.10)$$

$$1 - b = \frac{1}{2} \int_0^{+1} d\mu \int_0^{+1} p(\mu', \mu) d\mu' = \frac{1}{2} \int_0^{+1} d\mu \int_0^{+1} p(-\mu', -\mu) d\mu' \quad (7.11)$$

This allows the two-stream equations [in the absence of thermal emission or when $B(T) = 0$] to be written as:

$$-\bar{\mu}^- \frac{dL_{dave}^-(\tau)}{d\tau} = L_{dave}^-(\tau) - w_o(1 - b)L_{dave}^-(\tau) - w_obL_{dave}^+(\tau) - S^{*-}(\tau) \quad (7.12)$$

$$\bar{\mu}^+ \frac{dL_{dave}^+(\tau)}{d\tau} = L_{dave}^+(\tau) - w_o(1 - b)L_{dave}^+(\tau) - w_obL_{dave}^-(\tau) - S^{*+}(\tau) \quad (7.13)$$

where $\bar{\mu}$ is the cosine of the average polar angle $\bar{\theta}$ (when the polar axis coincides with the z-axis), and where the first-order scattering terms are shown below:

$$S^{*-}(\tau) = \frac{w_o E_s}{2\pi} [1 - b(\mu)] e^{-\frac{\tau}{\bar{\mu}}} = X^- e^{-\frac{\tau}{\bar{\mu}}} \quad (7.14)$$

$$S^{*+}(\tau) = \frac{w_o E_s}{2\pi} b(\mu) e^{-\frac{\tau}{\bar{\mu}}} = X^+ e^{-\frac{\tau}{\bar{\mu}}} \quad (7.15)$$

Since the average polar angle will likely vary between the the upwelling and downwelling hemispheres for anisotropic scattering, choosing a constant value for both hemispheres is clearly an approximation, but in the absence of additional information, and/or for instances of isotropic scattering, a constant value provides a reasonable first order approximation. In general, the choice of the cosine for the average polar angle varies between 0.5 and 0.71 (Thomas and Stamnes 1999). Theoretically, it is equal to the specific-intensity weighted angular mean according to Eq. (7.16).

$$\bar{\mu}^{\pm} = \langle \mu \rangle^{\pm} \frac{2\pi \int_0^{+1} L_d^{\pm}(\tau, \mu) \mu d\mu}{2\pi \int_0^{+1} L_d^{\pm}(\tau, \mu) d\mu} = \frac{E_d^{\pm}}{2\pi L_d^{\pm}} \quad (7.16)$$


The reader is referred to Appendix A for additional insight into the average cosine of the radiance field, and therefore its use in the two-stream approximation.

Predicted Results


Development of the two-stream approximation leaves the reader with two expressions as shown in Eqs. (7.12) and (7.13). Neither equation has any direct practical use. Two specific solutions are sought. The first imposes uniform illumination, isotropic scattering and a non-reflecting bottom boundary layer (see Appendix Section A.2). The second extends the first solution to include a Lambertian reflective boundary (see Appendix Section A.2.1).

Armed with the two-stream solution, and after consideration of the factors that influence the average cosine of the radiance field [Section A.1], the two-stream approximation can be used to solve for sensor reaching radiance. For the reflective Lambertian water-ground boundary, the results must also include all aforementioned analytical predictions described in Chapter 6. Figure 7.1 depicts an example scenario for simulation, and the predicted sensor reaching radiance as a combination of the direct reduced radiance described in Chapter 6 and using the two-stream approximation to compute the in-scattered radiance.

That being said, it is acknowledged that the two-stream solution is a good indicator of sensor reaching radiance, but for scenarios that are unlikely to exist in real water bodies. Therefore, the two-stream approximation will not be the main focus of this validation. Instead, these bench-mark solutions allow the results from more realistic phase functions to be inferred. Furthermore, since the same DIRSIG code, sampling methodologies and RT model must be employed, regardless of the phase function and type of illumination, even simple scenarios amenable to solution by the two-stream approximation are useful indicators of the accuracy of the photon mapped solution.



(a.) Scenario under investigation.



(b.) Predicted result.

Figure 7.1: (a.) Illustration of scenario under investigation, consisting of a 5m tall reflective “box” $r_b = 0.3$ above a ground plane $r_g = 1.0$, submerged in a water volume with a 7m depth. (b.) Predicted sensor reaching radiance [$W m^{-2} sr^{-1}$] as a combination of direct reduced sensor reaching radiance as described in Chapter 6, and using the two-stream approximation to describe the in-scattered radiance.

7.1.1 Results

Figure 7.2 compares the two-stream analytical prediction of sensor reaching radiance to the GRS and SRS generated results for a fairly shallow water volume. The average cosine that provides a reasonable match for radiance reflected from the ground plane is equal to an angle of approximately 57 degrees, while a slightly higher angle of 66 is a better fit for the top of the box. This makes intuitive sense since the top of the box is in 2m of water, and the light field would be more peaked in 2m of water, as opposed to more diffuse and uniform in 7m of water (total depth of the water volume.) Since the results are provided as a sanity-check, a rigorous fit of the average cosine was not sought, and instead was quickly chosen as a ball-park value.




Figure 7.2: SRS (asterisks) versus GRS (squares) versus analytical two-stream prediction (dotted line) of scattered radiance for a 7m deep water volume and a 5m tall box.

Nonetheless, there is satisfactory agreement between the two DIRSIG numerical solutions for sensor reaching radiance and that predicted by the two-stream approximation. The only major deviation is near the boundary between the vertical obstruction and the ground plane (near pixel 100), which is clearly not amenable to solution by the plane parallel two-stream approximation. Of greater interest is the expected DIRSIG results which show a higher return due to in-scattering very near this vertical boundary (*e.g.* pixel 50), and a softening of the demarcation of the shadowed region of the ground plane as the radiance is diffused (*e.g.* progressing from pixel 100 toward pixel 150). The next section will compare the SRS and GRS absolute solution of sensor reaching radiance to that predicted by independent numerical models.

7.2 Canonical Problem # 6

The canonical problems from Mobley et al. (1993) (summarized in Figure 7.3) are dominated by non-reflecting boundaries and solar illumination in the absence of a diffuse sky. Only problem #6 considers a reflective ocean floor which more closely emulates shallow waters with significant contributions from reflective in-water objects. Although problem #6 was evaluated by Goodenough (2007), it was restricted to GRS, and performed prior to recent code updates that incorporate an effective scattering coefficient (as described in the preceding chapter). Therefore, problem #6 will be re-evaluated here, and extended to include SRS. The combined results of the canonical validation performed by Goodenough (2007), the proposed use of the two-stream approximation described above, and re-evaluation of canonical problem #6 encompasses the extreme validation scenarios of (1.) *solar illumination and infinitely deep volumes*, (2.) *solar illumination with bottom effects*, and (3.) *uniform skylight combined with diffusely reflective in-water objects*.

Parameter	Problem						
	1 Easy Problem	2 Base Problem	3 Stratified Water	4 Atmospheric Effects	5 Windblown Surface	6 Bottom Effects	7 Raman Scattering
Albedo, ω_0	0.9, 0.2	0.9, 0.2	Depth dependent	0.9	0.9	0.2	0.29 at 417 nm 0.15 at 486 nm Eqs. (3) and (9)
Phase function	Rayleigh Eq. (3)	Particle Table 2	Depth dependent	Particle Table 2	Particle Table 2	Particle Table 2	
Air–water surface	Flat	Flat	Flat	Flat	Capillary waves	Flat	Flat
Diffuse sky radiance	0	0	0	Various models	0	0	0
Internal sources	0	0	0	0	0	0	Various models
Bottom boundary	Infinitely deep	Infinitely deep	Infinitely deep	Infinitely deep	Infinitely deep	Lambertian at $\tau = 5$	Infinitely deep

Figure 7.3: Summary of the canonical problems reproduced from Mobley et al. (1993).

7.2.1 Results

The predicted result for problem #6 is shown in Figure 7.4, where L_u is the upwelling radiance, E_{ou} is the upwelling scalar irradiance and E_d is the downwelling irradiance. Nearly equivalent DIRSIG simulations employing GRS and SRS were performed, using a geometric water depth of 5m, a scattering coefficient of 0.2, an absorption coefficient of 0.8, a Ward Lambertian BRDF with reflectance of 0.5, and DIRSIG’s inherent Petzold scattering phase function.

The GRS and SRS modeled results for problem #6 upwelling radiance, scalar irradiance and downwelling irradiance are shown in Figure 7.5. The symbols (triangles,




Figure 7.4: Predicted results for canonical problem #6 reproduced from Mobley et al. (1993).

diamonds and asterisks) are the mean predicted values provided by Mobley et al. (1993) (which were limited to two values (provided in tabular form) and where the standard deviation of the predicted values are actually less than the size of the symbols used to plot the mean values in Figure 7.5).

Overall, both GRS and SRS conform to the predicted results. The upwelling radiance as modeled by GRS is noisier than that predicted by SRS, which is a function of the second pass importance sampling. With SRS, the second pass need only sample the vertical path between target and detector, thereby effectively limiting the noise in the solution, while GRS must also sample between target and sources.

The agreement between GRS and SRS with both the two-stream approximate solution, and the numerical results predicted by Mobley et al. (1993), combined with the validation of direct reduced radiance shown in the preceding chapter, confirms DIRSIG's ability to model bottom effects in and around 3D in-water objects.

Before concluding this Section it should be noted that canonical problem #6 allowed for the validation of volumetric boundary compensation. In fact, it was determined that the volumetric boundary compensation was struggling to converge, as evidenced by the variations in the solution for upwelling radiance as a function of the search radius assigned within MRS. Naturally, the solution noise will vary with the size of the search radius, but a bias was also evident, being more pronounced at greater depths when slant paths (taken by GRS importance samples) would require effective boundary compensation [to negate both (a.) the foreshortening of the computed (and variable) search area, and (b.) attempting to sample photons on and below the ground plane (*e.g.* “inside” non-transmissive geometry)]. The result of the boundary bias error is illustrated in Figure

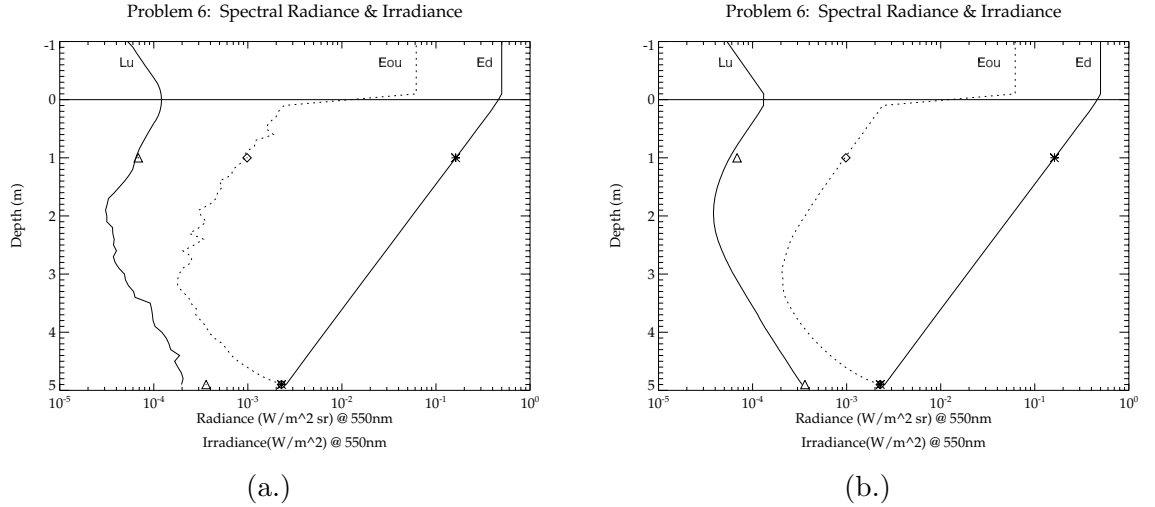


Figure 7.5: GRS (a.) and SRS (b.) modeled results for canonical problem #6. Each curve is labeled as either scalar upwelling irradiance [E_{ou} , dotted line], downwelling irradiance [E_d , solid line], or upwelling radiance [L_u , solid line]. The symbols (triangles, diamonds and asterisks) are the mean predicted values provided by Mobley et al. (1993) at the ground plane and 1m below the air-water interface. Note that both the spectral search radius and the cross sectional search fraction (in MRS) were set equal to 1.0 units (using 1 million photons). Additional GRS parameters include: 300 importance samples, quick samples = 1, $\mu = 10$, $\phi = 20$, max bounce = 1, and decay rate = 5. Additional SRS parameters include: SSR = 1.0 units and 1 million photons.

7.6, (top row) which depicts the GRS-MRS modeled results with search radii of 0.1 (a.), 4.0 (b.), 8.0 (c.) and 10.0m (d.). The larger the search radius, the smaller, and/or more erroneous the computed upwelling radiance near boundaries (e.g. depth = 5.0m). This volumetric boundary compensation difficulty has since been remedied (compare the more recent results shown in the bottom row of Figure 7.6 to those shown in the top row). Note that the search radii of 8.0 and 10.0m are extremely large, to the point of being somewhat unreasonable for this simulation, yet the new boundary compensation mechanism stills does a reasonable job of converging to the expected solution.

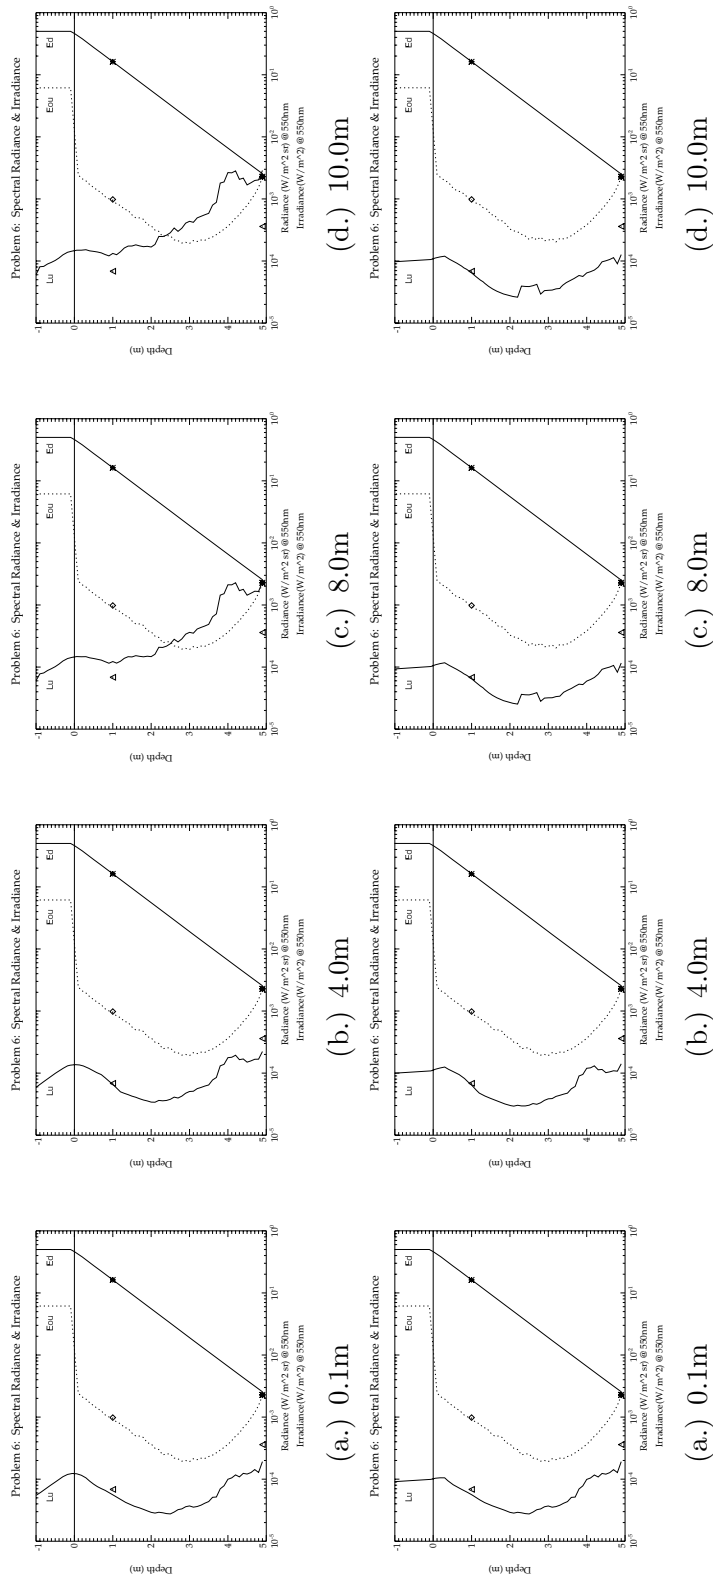


Figure 7.6: GRS results using MRS search radii of 0.1, 4.0, 8.0 and 10.0m. The Generic Radiometry Solver has been configured to request 300 importance samples per primary hit. Boundary compensation must be employed for each sample ray and sequentially in increments along the length of each importance sample. For the extreme search radii of 8.0 and 10.0m, the DIRSIG generated results begin to diverge from the expected solution (top row). The bottom row depicts an improved version of boundary compensation. Using this newer solution, the GRS results are consistent, regardless of the MRS assigned search radius, and converge very near the expected solution.

7.3 Field Profiles

Inspired by canonical problem #6, a small field collect was carried out (on October 12th 2009) that examined the upwelling and downwelling irradiance in coastal waters off Durand Beach in Rochester, New York. The data was collected using the HydroRad field instrument, which was first calibrated using a standard of spectral irradiance, as shown in Figure 7.7 and 7.8, for two of the three HydroRad receivers. In each plot, the curves at left show the original HydroRad response (solid line) to a spectral irradiance standard at fixed source-receiver distances, versus the expected spectral irradiance (dotted line). The plot at right shows an example of the receiver response after calibration, where the crosses depict the HydroRad receiver response at the shortest source-receiver distance, the diamonds show the expected response, and the solid line shows the calibrated (adjusted) HydroRad response, which now matches the expected response based on the standard of spectral irradiance.




Figure 7.7: HydroRad calibration examples for planar irradiance using channel D. The plot at left (a.) shows the HydroRad response (solid line) and the expected planar irradiance (dotted line, as a function of the standard of spectral irradiance (F864) at 5 distances [50.0, 67.5, 89.5, 116.0, 142.0] cm). The plot at right (b.) illustrates the curve for the standard of spectral irradiance at 50.0cm (diamonds), the HydroRad response (crosses), and the corrected and calibrated HydroRad response (solid line).

The in-water benthic bottom was assumed to be sandy, and upwelling and downwelling planar irradiance measurements from the beach sand served as the model for planar irradiance reflectance. The smoothed result is shown in Figure 7.9. Figure 7.10 depicts the laboratory analysis for the beach water (including the results from analyses of water samples taken at additional locations, although not discussed here). Table 7.2 reports the spectral absorption coefficients for the water sample at 550nm. Scattering coefficients were obtained by assuming that the water samples conformed to the specific coefficients




Figure 7.8: HydroRad calibration examples for planar irradiance using channel C. The plot at left (a.) shows the HydroRad response (solid line) and the expected planar irradiance (dotted line, as a function of the standard of spectral irradiance (F864) at 5 distances [49.75, 59.0, 71.2, 99.3, 142.0] cm). The plot at right (b.) illustrates the curve for the standard of spectral irradiance at 49.75cm (diamonds), the HydroRad response (crosses), and the corrected and calibrated HydroRad response (solid line).




Figure 7.9: Smoothed beach sand planar irradiance reflectance.


reported by Raqueno (1999), and shown in Table 7.3. Finally, Table 7.4 shows the DIRSIG IOP configuration.

Table 7.2: IOPs at 550 nm.


Description	TSS (gm^{-3})	Particle (m^{-1})	Chlorophyll (m^{-1})	CDOM (m^{-1})
Beach	0.7843	0.0612	0.0978	0.1058

Table 7.3: Specific coefficients at 550 nm, taken from generalized Lake Ontario data (Raqueno 1999).


TSS Specific Scattering Coefficient (m^2mg^{-1})	TSS Specific Absorption Coefficient (m^2mg^{-1})	Chlorophyll Specific Scattering Coefficient (m^2mg^{-1})	Chlorophyll Specific Absorption Coefficient (m^2mg^{-1})
0.5345	0.0221	0.1190	0.011 - 0.025



(a.) Particle absorption $\frac{1}{m}$.



(b.) Chlorophyll absorption $\frac{1}{m}$.



(c.) CDOM absorption $\frac{1}{m}$.

Figure 7.10: Absorption coefficients for various field collected water samples. Of specific interest here is the results for the Durand beach water sample shown as asterisks.

Table 7.4: DIRSIG water IOPs.

<pre> ADD_SCATTERING_MODEL { TYPE = constant ID = water SCAT = 0.00218 } </pre>	<pre> ADD_PHASE_FUNCTION_MODEL { TYPE = rayleigh ID = rayleigh SCATID = water K1 = 0.3911 K2 = 0.326 } </pre>	<pre> REFRACTIVE_INDEX_MODEL { TYPE = constant IOR = 1.34 } </pre>
<pre> ADD_ABSORPTION_MODEL { TYPE = constant ID = water_abs ABS = 0.05172 } </pre>	<pre> ADD_SCATTERING_MODEL { TYPE = constant ID = particle SCAT = 0.419252 } </pre>	<pre> ADD_ABSORPTION_MODEL { TYPE = constant ID = particle_ab ABS = 0.06116 } </pre>
<pre> ADD_PHASE_FUNCTION_MODEL { TYPE = petzold SCATID = particle } </pre>	<pre> ADD_SCATTERING_MODEL { TYPE = constant ID = chloro SCAT = 0.4659 } </pre>	<pre> ADD_ABSORPTION_MODEL { TYPE = constant ID = chloro_ab ABS = 0.0904701 } </pre>
<pre> ADD_PHASE_FUNCTION_MODEL { TYPE = petzold SCATID = chloro } </pre>	<pre> ADD_ABSORPTION_MODEL { TYPE = constant ID = cdom ABS = 0.1058 } </pre>	

7.3.1 Results

The water profile (collected using the HydroRad) was compared to that generated by DIRSIG, assuming the simulation coefficients reported in Table 7.4 and a sand reflectance of 0.1 at 550nm (rather than 0.14 due to a slight improvement in the fit). The comparison is plotted in Figure 7.11, where the *lines* show the field data profiles, and the *lines + symbols* show the equivalent radiometric variable generated by DIRSIG.

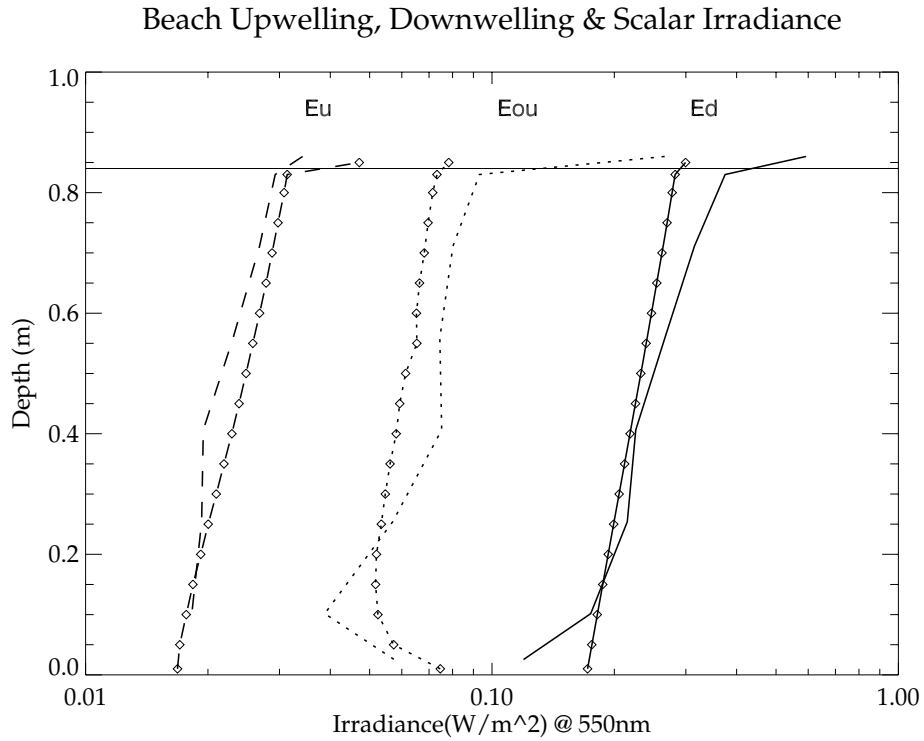


Figure 7.11: DIRSIG simulated results (symbols + lines) versus field water profile data (lines) for downwelling irradiance, upwelling irradiance, and upwelling scalar irradiance (using parameters listed in Table 7.4).

Although the results are very promising, Figure 7.12 compares the water profile data to DIRSIG simulations with slight variations in the IOPs. These results indicate that the field profile data is somewhat resistant to wide changes in IOP assignment. Although this does not negate DIRSIG's success at modeling the field results, it does indicate that the field data is a less than rigorous test-case. Some of this insufficiency is related to the shallow depth (approximately 1m) over which the field profiles were collected (which was

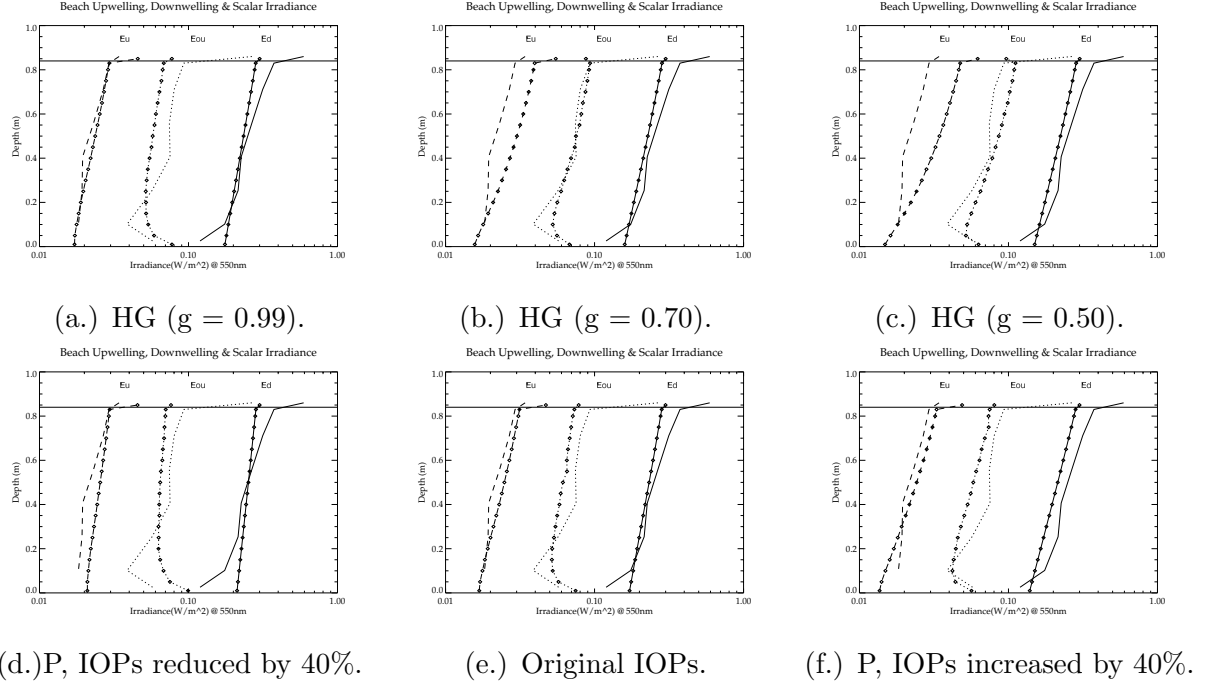


Figure 7.12: Variations: Henyey-Greenstein phase function with $g = 0.99$ (a.), $g = 0.70$ (b.), $g = 0.50$ (c.) and all other parameters the same as that listed in table 7.4. Petzold phase function with IOPs 40% lower (d.), the original IOPs (e.) and with IOPs 40% higher (f.).

limited by field instrumentation (*e.g.* approximately 100 feet of HydroRad cable, deployed from shore), and time of day/year (*e.g.* shallow water levels and low tide).

However, when the solution for the field profile is taken in combination with DIRSIG's ability to generate sensor reaching radiance that is in agreement with both analytical predictions and numerical solutions, this chapter as a whole nicely demonstrates DIRSIG's ability to successfully model radiative transfer in coastal ocean-water systems. The only caveat is that the radiative transfer problems presented in this chapter can be solved equally well using plane-parallel assumptions. Despite the limited 2D variability in the test-cases presented here, DIRSIG utilizes the same 3D tools, functions and code regardless of whether the user presents a 2D or a 3D scenario for solution. Thus, DIRSIG's successful solution of the 2D test scenarios shown here still lends confidence in the computational model's ability to handle 3D scenarios of interest. This is explicitly tested in the following chapter which compares DIRSIG's solution of 3D ship-shadow to that predicted by an independently developed Monte Carlo code.

Chapter 8

Ship-Shadow and Self-Shading

One dimensional RT solutions that treat the medium as a horizontally infinite homogeneous slab fail to predict accurate radiance distributions when used with geometries that violate 1D assumptions. An example of this is modeling the radiance distribution adjacent to and beneath a floating object. This implies that 2D and 3D scenarios, which leverage the benefits of the MC code, are rarely amenable to validation by other numerical or approximate methods. If possible, the MC code can be compared against physical measurements, but such detailed laboratory or field experiments are not always feasible, and even if possible, are rarely pursued.

The process of describing the light field for such a complex scene not only illustrates the utility of the MC code, but has several practical applications. For example, a detailed study of the radiance distribution around and beneath an obstruction, from the surface to various depths, is important in describing the dynamic range required for an underwater sensor. Reinersman and Carder (2004) found that the light field beneath a floating platform can exhibit a variation of at least 3 orders of magnitude, impacting the desired dynamic range and gain necessary for remotely operated in-water sensors.

An additional utility of light field prediction for self-shadowing scenes is the need to model the interfering effects expected when measuring radiance distributions from research platforms (Spinrad and Widder 1992). The presence of a *finite deployment structure* in close proximity to the location at which radiance measurements are taken “introduces abrupt medium changes, breaks the translation-invariant symmetry within a plane-parallel atmosphere-ocean system, and casts shadow perturbations on the surrounding radiant energy field” which culminates in 3D light field inhomogeneities (Doyle and Zibordi 2002). Therefore, a sensitivity study can be undertaken to evaluate the perturbation’s dependence on the optical properties of the water, the geometric relationship between illumination and obstruction, and the location at which the research field measurements are conducted.

Even if the instrument platform (*e.g.* detector) can be moved sufficiently far from the moored research vessel or tower to negate radiometric measurements that are influenced

by the ship's shadow, the shelf-shadowing by the field instrument will also impact the observations. If a point-detector is modeled on an instrument platform with radius R and height h , the measured upwelling radiance is impacted by the platform dimensions. If the platform is modeled in the absence of height (*e.g.* assumed to be a flat disk) then the measured upwelling irradiance will be a function of the extinction coefficient and the radius of the instrument through the product σR (Gordon and Ding 1992). The impact is lessened as the single scattering albedo increases (*e.g.* increasing w_o results in increased multiply scattered light, which fills in the self-shadow depicted in Figure 8.1, and where w_o is the probability that a photon will be scattered given an extinction event (and $[1 - w_o]$ is the probability of an absorption event per extinction event, or the co-albedo) (Thomas and Stamnes 1999)). Therefore, for a given σR , the error decreases with increasing w_o , but this dependence weakens as the solar zenith angle is increased (Gordon and Ding 1992) since the nadir-view of the shadow column through the water volume decreases as shown in Figure 8.1 (b.). If instrument height is considered in the model, then the results indicate that in low absorption water and with small instrument size, a flat disk is a reasonable approximation of the self-shading effects, but with larger instruments in turbid waters, the upwelling radiance will be impacted by the height dimension of the platform, underestimating the light field even more (Piskozub 2004a).




Figure 8.1: Illustration of the shadow beneath a flat field instrument. The nadir field of view of the point-detector is shown in (a.). As the zenith angle of the sun is increased (b.), the shadow column is less visible in a nadir measurement. As the single scattering albedo increases (c.) the shadow becomes more pyramidal as light diffuses in, thereby decreasing the self-shading impact.

The impact of the ship-shadow on radiometric measurements has been investigated using Monte Carlo numerical simulations. In these studies, the in-water light field perturbations can be divided into those caused by the ship's shadow, and those caused by interaction of the upwelling photons with the underside of the ship's hull. Since the up-

welling irradiance is typically less than 10% of the downwelling irradiance, and since the albedo or reflectance of the ship's hull is usually very small (less than 10%), the perturbation due to upwelling photons interacting with the hull of the ship and contributing to the downwelling light field is at most 1% (or 10% of 10% of the maximum downwelling irradiance) (Gordon 1985). For the downwelling irradiance, if the single scattering albedo is zero, then no perturbation error should be apparent, provided the measurements are not taken within the ship shadow region. If the upwelling irradiance is modeled, the nature of the scattering indicatrix will determine how much of the upwelling irradiance is a function of pure back-scattering, and how much is a function of forward scattered radiance that is reflected from the bottom Lambertian surface. However, transit through the water means that the upwelling irradiance that reaches the air-water interface will be fairly diffuse, and at least much more diffuse than that found for the downwelling irradiance. Gordon (1985) reports that if the light field is evaluated in the presence and absence of the obstruction, the perturbation of the upwelling radiance and irradiance when the obstruction is present depends only weakly on the scattering phase function (at least when the size of the obstruction is less than the scattering length), which makes sense since the illumination has undergone multiple scattering events.

The illumination conditions naturally play a large role in the observed perturbation. Figure 8.2 illustrates a typical observation scenario, with the sun and instrument platform (on a boom) off the stern of the ship. This ensures that the ship's shadow trails away and beneath the field instrument (Weir et al. 1994). If the solar zenith angle is increased to 90 degrees, it is reasonable to expect the perturbation to decrease. This should be true for both upwelling and downwelling light fields. Conversely, a uniform downwelling sky (versus a direct solar beam) results in greater downwelling perturbations than those observed for all reasonable sun zenith and solar angles.

A field study conducted by Weir et al. (1994), using a scene similar to that depicted in Figure 8.2 was carried out, but using two measurement instruments, one held by the boom between 1 and 6m from the stern of the ship, and a second with buoyant fins at a distance of at least 20m from the research vessel. Simultaneous depth profiles were evaluated using both detectors (*e.g.* let detector A be the instrument fixed by the boom, and let detector B represent the device "fished out" to at least 20m). A total of 15 paired depth evaluations were conducted; 7 with A at 1m and 8 with A at 6m. For *downwelling irradiance*, Weir et al. (1994) did not find a *statistical difference* between detector A and detector B, at either 1 or 6m, although at 1m the measured downwelling irradiance is always underestimated, suggesting the influence of the ship shadow. In contrast, the *upwelling radiance* for detector A at 1m was statistically different from that evaluated by detector B, with the mean difference increasing toward the surface (where the ship shadow would be more dominant). For detector A at 6m, a statistical difference was not found, suggesting that, for this illumination geometry and conditions, measurements taken at 20m from the research vessel are equally unaffected (or affected) by the ship as




Figure 8.2: Illustration of solar zenith and azimuth.

those taken at 6m. The conclusions put forth by Weir et al. (1994) describe the lack of variability in the downwelling irradiance, and the presence of variability in the upwelling radiance at deployment distances less than 3m and in the upper 20m of the water volume, with the greatest variability near the sea surface. Unfortunately, the noise in the data is rather large, and the authors clearly acknowledge the variability induced by wave glint, small clouds, ship roll, random errors, etc.

Complementary to the field research conducted by Weir et al. (1994), Gordon (1985) studied ship-shadow perturbations using a Monte Carlo code simulation. The results amounted to a sensitivity study that evaluated the perturbation error as a function of illumination conditions and optical properties. Gordon (1985) summarizes that the perturbation induced errors for downwelling irradiance in the presence of a direct solar beam is on the order of $\approx 1 - 3\%$ when the zenith angle is greater than 20° , and the azimuth angle is between $-45^\circ < \phi_o < 45^\circ$ (see Figure 8.2). He goes on to state that this error “increases with increasing c (extinction coefficient) and w_o but decreases with increasing θ_o ”. Inspection of the results plotted by Gordon (1985) visually indicate that this is true for an increase in the extinction coefficient, although only two trials at 0.1 and 0.3 m^{-1} are shown (additional plots include a coefficient of 0.5 m^{-1} , but include additional parameter variations). Gordon (1985) also plots the downwelling irradiance error for three single scattering albedos ($w_o = 0.5, 0.7$ and 0.9). At the surface, the error is clearly greater for the smallest albedo of 0.5 . As the depth increases, the impact of w_o is much less apparent, but increasing w_o contributes a slight decrease in the irradiance error for collimated

illumination. In contrast to the direct illumination conditions, the perturbation error for diffuse illumination is largely independent of the IOPs. Perhaps quite intuitively, it is more dependent on skyfraction effects (Gordon 1985). The upwelling irradiance shows greater error than the downwelling (*e.g.* measurements $\approx 1\text{ m}$ from the side of the ship, in water with extinction coefficients between 0.1 and 0.3 m^{-1} will be between 10 and 30% less than expected when measured over zero to 50m in depth). Naturally, all errors decrease as the sensing platform is moved further from the obstruction.

Since Gordon (1985) used between 15,000 and 10,000 photon histories to compute the perturbations, Piskozub (2004b) reinvestigated the effects of ship shadows using 400 times as many photons, including a numerical representation of surface waves. The work by Piskozub (2004b) included a non-reflecting ship, described as a black rectangular box, completely lacking an “above-water” structure. The results indicate that the wind direction (*e.g.* wind-to-sun azimuth) is irrelevant to the induced error for both upwelling and downwelling irradiance, as is the wind velocity impact on upwelling irradiance when varied between 0 and 15 m s^{-1} [although an ambiguity in the caption provided by Piskozub (2004b) makes it unclear if the results are for a depth of $z = 10\text{m}$, or rather three separate measurements at 1, 10 and 20m]. However, assuming the validity of these conclusions, then additional trials need not consider the impact of the wind velocity in future ship-shadow studies.

Reinersman and Carder (2004) used another MC code and undertook the more global task of characterizing the irradiance distribution expected beneath a floating object by taking several measurements to prepare contour plots between 0 and 15m on either side of the mid-point of the floating obstruction (see Figure 8.3). The problem under consideration consisted of a floating barge in a uniform water column illuminated with a 30 degree zenith angle direct beam (79% of downwelling irradiance) and a diffuse uniform sky (21% of downwelling irradiance).

The floating object was given a Lambertian reflectance of 0.1, and the bottom plane a Lambertian reflectance of 0.15 for clear water, and 0.05 for a turbid estuarine water trial. The optical parameters are shown in Table 8.1, although the authors fail to specify the exact scattering indicatrix employed.

The MC code described by Reinersman and Carder (2004) uses repetitive volumes to describe a scene in an effort to decrease the number of photon histories that must be traced in order to build up the scattered light distribution in a scene. The photon history is reduced since a solution for a given volume can be repeated elsewhere for the same geometry and optical properties. Technically, this means that the results presented by Reinersman and Carder (2004) are for the scene described in Figure 8.3, but with the left and right sides of the 2D scene wrapped to create a scene with a series of long parallel “barges”, each 10 m wide, spaced 20 m apart. This means that the upwelling and downwelling irradiance distributions may be affected by neighboring “barges”, depending upon the resolution of the repetitive “volumes”.




Figure 8.3: 2D illustration of ship-shadow scene. Figure adapted from Reinersman and Carder (2004).

Table 8.1: Table adapted from Reinersman and Carder (2004) to describe the optical properties for the 2D ship-shadow model.

Property	Value	Value
Chlorophyll Concentration	$\frac{0.3 \text{ mg}}{\text{m}^3}$	$\frac{5 \text{ mg}}{\text{m}^3}$
σ_a (water)	0.05172 m^{-1}	0.05172 m^{-1}
σ_s (water)	0.00218 m^{-1}	0.00218 m^{-1}
σ_a (particles and CDOM*)	0.01304 m^{-1}	0.1150 m^{-1}
σ_s (particles)	0.1448 m^{-1}	0.8391 m^{-1}
Bottom Reflectance	0.15	0.05
Ship Reflectance	0.10	0.10

*CDOM = colored dissolved organic matter.

Based on the scenario shown in Figure 8.3, including observations by Gordon (1985) and Weir et al. (1994), several interesting phenomenon can be simulated and evaluated. First, the ship shadow perturbations should be most significant near the air-water interface for downwelling irradiance, and most significant near the reflective bottom boundary for upwelling irradiance. The latter is true, until the radiance becomes so diffuse that it lacks any perturbation due to the floating obstruction above. Second, if the ship reflectance is varied from zero to increasing values, the radiance distribution *reflected* from the floating object (non-shadowed side) should increase (*e.g.* upwelling-and-reflected contributions to downwelling irradiance). Third, the surface shadow will be extended as the zenith angle

is increased, but the in-water shadow column will bend further behind the ship and out of the field of view of the measuring instrument for a typical observation scenario (see Figure 8.1 (c.)) for a comparable illustration). Fourth, in the absence of in-scattered radiance contributions (*e.g.* only attenuation and out-scattering), the shadow of the object on the Lambertian ground plane and through the water column should be relatively sharp. When the in-scattered radiance is computed, the downward irradiance will diffuse into the shadow. This will cause the edges of the shadow column to blur, effectively shrinking the radius of the shadow column (see Figure 8.1 (c.)) for a comparable illustration). If sufficient light reaches the bottom reflecting plane, this will likewise be evident in the shadow projected onto the bottom Lambertian surface. Naturally, this change in the shadow column will be a function of the scattering indicatrix (*e.g.* predominant forward and/or back scattering will blur and shrink the shadow column less than isotropic scattering). Moreover, the modeling of such a scene is clearly a 3D problem, and nicely utilizes the benefits of a photon mapped Monte Carlo code to describe in-water radiative transfer.

8.1 Proposed Methodology - Simulated Ship-Shadow

Goal: Evaluate the normalized upwelling and downwelling irradiance in the presence of ship-shadow.

Source Function	✓		
Non-Source Function		Model-vs-Model Comparison	✓
Submerged or Floating Objects	✓	Model-vs-Analytical Prediction	
Boundary Interface		Model-vs-Observation Comparison	✓
3D Geometry	✓		

Figure 8.3 describes a specific 3D radiance scene to be modeled by DIRSIG. The results will be qualitatively evaluated based upon the physics of the scene, and compared to published numerical and field data by Reinersman and Carder (2004). When appropriate, a quantitative comparison will also be undertaken. The scene to be simulated consists of an interesting and realistic 3D radiative transfer problem, with obvious validation appeal. The IOPs will include the extremes of clear and turbid water, as per the proposed scenarios A and B shown in Table 8.2. DIRSIG will be used to generate upwelling and downwelling irradiance in the presence of a floating obstruction. The results will be analyzed qualitatively and quantitatively, and compared against the results shown in Figure 8.6 and 8.7.

The experimental parameters provided by Reinersman and Carder (2004) (and summarized in Table 8.2) are limited to a single spectral wavelength (532 nm). Likewise, all DIRSIG experimental results are referenced to this single wavelength. In addition to the lack of additional spectral data, the experimental parameters do not include a scattering phase function for either pure water, chlorophyll or particulate matter. Therefore,

Table 8.2: Proposed ship-shadow scenarios.

Property	Scenario A	Scenario B
Chlorophyll Concentration	$\frac{0.3 \text{ mg}}{m^3}$	$\frac{5 \text{ mg}}{m^3}$
σ_a (water)	0.05172 m^{-1}	0.05172 m^{-1}
σ_s (water)	0.00218 m^{-1}	0.00218 m^{-1}
σ_a (particles and CDOM*)	0.01304 m^{-1}	0.1150 m^{-1}
σ_s (particles)	0.1448 m^{-1}	0.8391 m^{-1}
Bottom Reflectance	0.15	0.05
Ship Reflectance	0.10	0.10
Scattering Phase Function	Pure Water + Petzold	Pure Water + Petzold
Illumination	**	**

All variables are reported at 532 nm. * = colored dissolved organic matter. ** = 30° zenith angle direct beam (79% of downwelling irradiance) and diffuse uniform sky (21% of downwelling irradiance). Scenarios A and B are adapted from Reinersman and Carder (2004).

reasonable assumptions will be made when generating comparable DIRSIG simulations. The pure water scattering coefficient can be mimicked in DIRSIG using the Buiteveld based scattering model with coefficient K and exponent q . The original model predicts a coefficient of 0.0016 m^{-1} at 532 nm, as per Eq. (8.1) (Goodenough 2007).

$$\sigma_s(\lambda) = 3.4123664 e^8 [\lambda(nm)^{-4.1537676}] \quad (8.1)$$

If the exponent is augmented to $q = -4.1067360$ the coefficient will be $\sigma_s = 0.00218 \text{ m}^{-1}$ at 532 nm. The augmented model will therefore follow the spectral scattering coefficient curve shown in Figure 8.4.

```
ADD_SCATTERING_MODEL {
TYPE = buiteveld
ID = buiteveld
K = 3.4123664
Q = -4.1067360
}
```

In the absence of any additional information, the scattering phase function for pure water will be assumed to follow the Rayleigh form provided by Goodenough (2007) as an intrinsic DIRSIG IOP.

```
ADD_PHASE_FUNCTION_MODEL {
ID = rayleigh
}
```




Figure 8.4: Augmented Buiteveld scattering coefficient.

```

TYPE = rayleigh
SCATID = buiteveld
K1 = 0.391128
K2 = 0.835000
}

```

The pure water absorption coefficient can likewise be met using DIRSIG’s intrinsic Hybrid absorption model, which includes the value of $\sigma_a(530\text{ nm}) = 0.0505$ and $\sigma_a(535\text{ nm}) = 0.0527$. Therefore, the interpolated value at 532 nm can be expected to reasonably satisfy the value of 0.05172 m^{-1} used by Reinersman and Carder (2004).

```

ADD_ABSORPTION_MODEL {
TYPE =sbpf
ID = sbpf
}

```

Reinersman and Carder (2004) provide chlorophyll concentration, as well as the particle and CDOM-based absorption coefficients, and the particle-based scattering coefficient expected for each “water type”. The CDOM and particle coefficients can be used directly, while the chlorophyll scattering and absorption coefficients can be determined by assuming specific coefficients of $0.0146\text{ } \frac{\text{m}^2}{\text{mg}}$ (absorption) and $0.1162\text{ } \frac{\text{m}^2}{\text{mg}}$ (scattering) (source: WIKI

(Vodacek) file dated May 20, 1999 derived from Lake Ontario field collect data). Since a specific scattering phase function is excluded in the experimental parameters provided by Reinersman and Carder (2004), a Petzold scattering phase function will be implemented. The remaining details related to the DIRSIG configuration for scenario A are shown below.

```
ADD_PHASE_FUNCTION_MODEL {
  TYPE = petzold
  SCATID = particle
}
```

```
ADD_SCATTERING_MODEL {
  TYPE = constant
  SCAT = 0.1448
  ID = particle
}
```

```
ADD_ABSORPTION_MODEL {
  TYPE = constant
  ABS = 0.01304
  ID = particle_CDOM
}
```

```
ADD_ABSORPTION_MODEL {
  TYPE = constant
  ID = chloro
  ABS = 0.00439734
}
```

```
ADD_PHASE_FUNCTION_MODEL {
  TYPE = petzold
  SCATID = chloro
}
```

```
ADD_SCATTERING_MODEL {
  TYPE = constant
  ID = chloro
  SCAT = 0.03486
}
```

Results

The results were generated using monitor disks (virtual DIRSIG detectors) that can record upwelling and downwelling irradiance. The monitor disks were situated at 1m increments

both vertically and horizontally along a single transect through the simulation as illustrated in Figure 8.5. In addition, disks were placed just above the bottom boundary, just below the obstruction, and just above and below the air/water interface. The DIRSIG results (which are a function of the MediumPM Radiometry Solver) were then compared to that predicted by Reinersman and Carder (2004).




Figure 8.5: X,Y,Z center location of monitors to collect upwelling and downwelling irradiance. Note that all monitors have a 1-unit radius unless otherwise specified.

Figure 8.6 and 8.7 depict the downwelling and upwelling irradiance contours (respectively) that result from the aforementioned simulation parameters. In each Figure, the top row shows the predicted results from Reinersman and Carder (2004) for the clear (left) and turbid (right) water scenarios. The bottom row illustrates the DIRSIG generated results. In each Figure, horizontal contours highlight vertical gradients, and vertical contours highlight horizontal gradients. In the plots of downwelling irradiance, the horizontal contours indicate radiance that is downwelling from the atmosphere (or solar disk), and radiance reflected from the underside of the barge. Beneath the barge, the vertical contours highlight the edges of the volumetric shadow, and indicate diffuse volumetric inscattering. In the plots of upwelling irradiance, the horizontal contours show reflected radiance, and/or radiance that has been backscattered. Again, the vertical contours near the barge show volumetric inscattering.


In Figure 8.6, the DIRSIG results are a good match to that predicted by Reinersman and Carder (2004) for the clear water trial. The greatest perturbations are (and should be) near the sea surface. The vertical contours show downwelling irradiance that is diffusing

into the volumetric shadow cast by the barge, and the horizontal contours beneath the floating obstruction show diffuse radiance that has been reflected by the underside of the barge.


The downwelling irradiance from the turbid water trial mimics that of the clear water scenario, but with a few exceptions. First, the contours are more closely spaced due to the increased attenuation. Second, the results are less noisy. This is believed to be a function of three things. To begin with, the radiance is still fairly forward peaked for the clear water scenario, which means less probable scattering events can generate higher variance. This is much less likely with the turbid water scenario where the distribution has become much more diffuse. In addition, the turbid water scenario has a lower single scattering albedo (approximately 0.88 versus 0.91 for the clear water scenario), and a much higher total extinction coefficient (approximately 1.6 versus 0.2 for the clear water scenario), which collectively is offered as an explanation for the decrease in noise. Finally, the DIRSIG results near the ground plane deviate from those shown by Reinersman and Carder (2004). The DIRSIG simulation reveals horizontal contours near the ground plane and beneath the barge, which are absent in the prediction. This is believed to be a function of a higher degree of scattering (either overall, or from a less forward peaked phase function), which has filled in the volumetric shadow to a greater degree than that predicted by Reinersman and Carder (2004). Given the uncertainty in the phase function employed in the prediction, this deviation is not believed to be a flaw in the DIRSIG solution, but rather a nice prediction of the results that can be expected with a more diffuse water scenario. Moreover, it is reasonable to assert that as turbidity increases, the downwelling irradiance decreases, and at some point there is a shift wherein in-scattered radiance, as opposed to reflected radiance, dominates the irradiance distribution beneath the obstruction.

The upwelling irradiance depicted in Figure 8.7 for the clear water trial is actually believed to be a better representation of reality than that provided by the prediction. Both show that for this scenario, a great deal of the irradiance is a function of former reflection from the ground plane. This means that the upwelling irradiance is less uniform near the ground plane, but becomes more uniform closer to the barge (in other words, the greatest perturbations are near highly reflective objects or the ground plane). However, the contours far removed from the volumetric shadow are still fairly vertical in the prediction, but nearly horizontal in the DIRSIG simulation. Reinersman and Carder (2004) acknowledge that their vertical contours may be an artifact of their simulation model (see above for Reinersman and Carder (2004) model details). The artifacts are absent from the DIRSIG simulated results by virtue of the photon mapping methodology.


The upwelling irradiance contours for the turbid water scenario share characteristics with the downwelling irradiance contour plots, in that the contours are more closely spaced, and less noisy, than their clearer water companions. The steep vertical contours that emanate from the bottom of the floating object suggest that the area just beneath the




(a.) Clear water.



(b.) Turbid water.




(c.) Clear water.




(d.) Turbid water.

Figure 8.6: (a.) Predicted and (c.) DIRSIG modeled normalized \log_{10} downwelling irradiance at 532 nm as a function of depth and distance from a floating obstruction for chlorophyll = 0.3 mg m^{-3} . Predicted (b.) and DIRSIG modeled (d.), but with chlorophyll = 5.0 mg m^{-3} . The predicted results are adapted from Reinersman and Carder (2004). The DIRSIG simulations use SRS, with more than 100 photons per cubic meter in the volume, and more than 240 photons per square meter on the ground surface, with an SSR of 0.8m.


barge is illuminated solely by scattered radiance rather than reflection from the bottom plane (Reinersman and Carder 2004). At greater depths, the DIRSIG results deviate from the predictions. Again, this is believed to be a function of greater scattering, either overall, or from a less forward peaked phase function, ultimately generating a nearly diffuse irradiance field near the ground plane. Note that the DIRSIG simulation results have not been smoothed in order to accentuate the variability associated with the IOP assignment.




(a.) Clear water.



(b.) Turbid water.



(c.) Clear water.



(d.) Turbid water.

Figure 8.7: (a.) Predicted and (c.) DIRSIG modeled normalized \log_{10} upwelling irradiance at 532 nm as a function of depth and distance from a floating obstruction for chlorophyll = 0.3 mg m^{-3} . Predicted (b.) and DIRSIG modeled (d.), but with chlorophyll = 5.0 mg m^{-3} . The predicted results are adapted from Reinersman and Carder (2004). The DIRSIG simulations use SRS, with more than 100 photons per cubic meter in the volume, and more than 240 photons per square meter on the ground surface, with an SSR of 0.8m.

8.2 Validation of a 3D Radiative Problem: Ship-Shadow

If any conclusions can be drawn from the results presented in this chapter, and illustrated visually via Figures 8.6 and 8.7, it is that the contour plots (which are a function of the MediumPM Radiometry Solver) nicely illustrate the power and fidelity of a photon mapped solution. Furthermore, it is rare to have sufficient knowledge of the simulation parameters of a Monte Carlo method so that a photon mapped solution can be directly compared to that generated by another MC model, especially for a such a complex and interesting 3D

scenario. Although not intended as a one-to-one comparison with the results predicted by Reinersman and Carder (2004), the degree of similarity between DIRSIG and these predictions is considered very promising.

Chapter 9

Image Quality

Although this chapter will not expound upon the computational details of the photon mapped RT solution, a few words will be required in order to set the stage concerning image quality. This chapter will briefly discuss the technique, and its associated image quality concerns, before concluding with a small trade study that illustrates the expected signal to noise ratio (SNR), variance and bias for three different turbidity scenarios.

Photon mapping is a two-pass rendering technique. Each pass uses a different algorithm to compute a portion of the radiative transport process. The first pass is sometimes termed particle tracing owing to its similarity to raytracing. The particle tracing process corresponds to tracing light paths from sources through the scene. Along the path, the particle can be scattered or absorbed. If scattered, the light path takes on a new direction and the incident and outgoing paths form a vertex. Each path vertex is treated as a particle, therefore N light paths can give rise to more than N particles (Suykens-De Laet 2002). The particles are usually referred to as photons.

The fate of the ray depends on probabilistic descriptions of the scene, including source sampling, extinction events, the probability of scattering versus absorption for a given extinction event, and if scattered, the probable new direction given the medium's scattering phase function. Adding additional objects to the medium creates the possibility of new and different absorption and reflection events based on the optical properties of the submerged materials.

The second pass is a radiance reconstruction process that uses the stored photon map. Rays are traced from detector elements into the scene, and if inside the photon mapped participating medium, the photon map is queried to solve for in-scattered radiance contributions. The reconstruction process is very similar to a density estimation problem (Suykens-De Laet 2002). Density estimation problems seek to solve for a given probability density function using a limited number of observed sample points. In radiance reconstruction, the photons in the map adjacent to the reconstruction point represent the discrete points used in the reconstruction of the in-scattered radiance.

9.1 Overview of Density Estimation

As noted above, the first pass produces a photon map with a volume density. The total volume density is a function of the user-assigned variable *max photons* and the size of the bounding box. Locally (*e.g.* per pixel), the volume density is a function of the max photons, the illumination conditions, the scattering characteristics of the medium, and the user-assigned *search radius* variable (see Figure 9.1 which illustrates variable search radii (*h*) and max photon sample number (*N*)). Loosely, the photons in the map can be viewed as a set of observed data points, sampled from an unknown radiance function. This scattered radiance function must be reconstructed at each pixel location using a limited number of photons or samples. The process is very much like a traditional density estimation problem, although there are subtle differences which will be explained along the way. To begin, it is important to understand the basic tenets of density estimation. A simple and useful introduction to this field can be found in Silverman (1998). Much of this discussion follows directly from this text.

9.1.1 Univariate Density Estimation

There are two approaches to density estimation. The first is to assume that the observed data points have been drawn from a parametric distribution, such as a normal distribution which can be described by a mean and variance. If the observed data samples can be used to estimate the mean and variance, then the entire underlying density function can be conveniently reproduced. Except for simple scenarios, the photon mapped scattered radiance distribution is not constrained to fit a family of parametric curves. Therefore, nonparametric estimation techniques must (generally) be considered.

Silverman (1998) describes various nonparametric methods of density estimation, of which three univariate examples will be summarized below. The last (the kernel estimator) is used by DIRSIG, but its introduction is preceded by a discussion of the histogram and naive estimators since they each provide insight into the formulation of the kernel estimator.

The *histogram density estimator* places N observation points, drawn from X_1, \dots, X_N samples, into a specified number of bins:

$$\hat{f}(x) = \frac{1}{N} \frac{\text{number of } X_i \text{ in same bin as } x}{\text{bin width containing } x} \quad (9.1)$$

The user exercises control over the size of the bin width (which determines the amount of smoothing in the data), and the choice of bin origin (which can have a noticeable impact on the perceived distribution). Although the smoothing can be varied in a logical and productive way, the perceived nature of the distribution varies with the choice of bin origin in a non-instructive manner and must be considered when using histogram estimators. The *naive estimator* overcomes the bin origin problem by constructing a histogram where the

reconstruction point is the center of the sampling interval (*e.g.* $[x - h, x + h]$ where h is the half bin width). However, the user must still specify the parameter h and therefore the degree of smoothing.

$$\hat{f}(x) = \frac{1}{N} \frac{\text{number of } X_i \text{ in } [x - h, x + h]}{2h} \quad (9.2)$$

The naive estimator can also be expressed according to Eq. (9.3):

$$\hat{f}(x) = \frac{1}{N} \frac{1}{h} \sum_{i=1}^N w\left(\frac{x - X_i}{h}\right) \quad (9.3)$$

where the weight variable $w(t)$ is shown below and using the change of variable $t = \frac{x - X_i}{h}$:

$$w(t) = \begin{cases} \frac{1}{2} & \text{if } |t| < 1; \\ 0 & \text{otherwise.} \end{cases}$$

The *kernel estimator* is simply the naive estimator, but with the weight function $w(t)$ replaced by a *kernel* function where $K(t)$ is any function that satisfies:

$$\int_{-\infty}^{+\infty} K(t) dt = 1 \quad (9.4)$$

The kernel estimator is then:

$$\hat{f}(x) = \frac{1}{N} \frac{1}{h} \sum_{i=1}^N K\left(\frac{x - X_i}{h}\right) \quad (9.5)$$

where h is the window width parameter responsible for smoothing (and sometimes termed the bandwidth). In words, and by analogy to the naive estimator, imagine a kernel function that takes the form of a normal distribution. The kernel estimator is then the sum of Gaussian functions placed at observation points, and where h determines the width of each Gaussian function. As $h \rightarrow 0$, the result is a sum of Dirac delta functions at each observation point, and as h becomes excessively large, the details in the function $\hat{f}(x)$ are lost.

Silverman (1998) reports that one shortcoming of the kernel estimator is that the fixed window width causes noise to appear in long-tailed distributions, and that if smoothed out, the details in the primary portions of the distribution are lost (*e.g.* a fixed h can either under smooth the tails or over smooth the bulk of the distribution).

9.1.2 Mean Square Error

Using any one of the three density estimators described above results in both bias and variance, the degree of which depends on h , or both h and N . The mean square error in an estimate is (Silverman 1998):

$$\begin{aligned} MSE_x(\hat{f}) &= E[\hat{f}(x) - f(x)]^2 \\ &= [E[\hat{f}(x)] - f(x)]^2 + var \hat{f}(x) \\ &= [bias_h(x)]^2 + var \hat{f}(x) \end{aligned} \quad (9.6)$$

where E represents the expected value. This means that the MSE is the sum of the squared bias and the variance at x . The density estimate introduces a trade-off between the variance or noise in the data, versus systematic error or bias in the estimate (Schjoth et al. 2007). In photon mapping, the stochastic Monte Carlo sampling process causes the variance, which can be driven down by increasing the number of samples used to estimate the radiance. Conversely, the bias “blurs” the illumination which is problematic in shadows, caustics or anywhere where the radiance distribution is not smoothly varying. The bias does not depend directly on the sample size, but rather the smoothing function (*e.g.* the window width h in the estimator). However, if h is chosen as a function of N , then the bias will indirectly depend on sample size (Silverman 1998).

In order to intelligently reconstruct the radiance at an observation point, it is important to understand how the density estimation impacts both bias and variance. The expected value of $\hat{f}(x)$ can be determined as follows (Silverman 1998), (Walter et al. 1997):

$$E[\hat{f}(x)] = \int_{-\infty}^{+\infty} \frac{1}{h} K\left(\frac{x - X_i}{h}\right) f(X_i) dX_i \quad (9.7)$$

This means that the density estimate “converges, not to the correct function f , but rather to the correct function convolved with the kernel” (Walter et al. 1997). The expected value of \hat{f} is then a smoothed version of the true density (or a biased estimate of the function), unless (1.) the window width is zero, or (2.) the underlying function is uniform. Increasing the number of samples will not result in a correct estimate, and instead, the bandwidth or h must be decreased as the number of samples increases.

$$\begin{aligned} bias_h(x) &= E[\hat{f}(x)] - f(x) \\ &= \int_{-\infty}^{+\infty} \frac{1}{h} K\left(\frac{x - X_i}{h}\right) f(X_i) dX_i - f(x) \end{aligned} \quad (9.8)$$

To understand how the density estimate impacts the variance, first consider *approximate estimators* of both the bias and variance, which are shown below without derivation (see Silverman (1998) for additional details).

$$\begin{aligned}
bias_h(x) &\approx \frac{1}{2}h^2 f''(x) \int_{-\infty}^{+\infty} t^2 K(t) dt \\
var \hat{f}(x) &\approx \frac{1}{N} \frac{1}{h} f(x) \int_{-\infty}^{+\infty} K(t)^2 dt
\end{aligned} \tag{9.9}$$

The estimates show that the bias is proportional to h^2 and the second derivative of the density function (and yet not a function of N). Since the second derivative of $f(x)$ is a measure of the high frequency content in $f(x)$, the bias in a rapidly fluctuating density can be reduced by employing a smaller h (which of course matches intuition). It is also evident that the variance is proportional to $\frac{1}{Nh}$. Therefore, decreasing the noise requires a larger h and/or a larger N . The trade-off between the reduction in bias versus variance is now apparent; decreasing bias requires a small h , but decreasing variance requires a large h . This is illustrated in Figure 9.1. The left column shows a fixed, small h , while the right column shows a fixed, large h . For each fixed bin width, the top row shows a simulation using a higher N , while the bottom row shows the exact same simulation, but using a smaller N . On the top surface of the volume is an obstruction which generates a shadowed area within the volume. When combined with the larger bin width, the sampling area extends across the edge of the shadow, blurring the results and generating bias. This is not so when using the smaller bin width. However, when the number of samples is decreased (bottom row), the smaller bin width will show greater variance since the result will be based on much fewer samples.

9.1.3 Multivariate Density Estimation

As a corollary to the univariate kernel estimator discussed above, the multivariate density estimator for d -dimensional space is:

$$\hat{f}(\vec{x}) = \frac{1}{Nh^d} \sum_{i=1}^N K\left(\frac{\vec{x} - \vec{X}_i}{h^d}\right) \tag{9.10}$$

where K is usually defined as a radially symmetric unimodal probability density function. Examples include the multivariate normal density (Silverman 1998) and the 2-D Epanechnikov kernel (Walter et al. 1997). A common photon mapping kernel is also shown below:




Figure 9.1: Illustration of trade-off between bin width (h) and sample number (N).

Normal Density:

$$K(\vec{t}) = \frac{1}{2\pi} \exp \left[-\frac{1}{2} \vec{t}^T \vec{t} \right] \quad (9.11)$$

2-D Epanechnikov:

$$K(\vec{t}) = \begin{cases} \frac{2}{\pi} (1 - |\vec{t}|^2) & \text{if } |t| \leq 1; \\ 0 & \text{otherwise.} \end{cases}$$

Common photon mapping Kernel:

$$K(\vec{t}) = \begin{cases} \frac{1}{c_d} & \text{if } |t| \leq 1; \\ 0 & \text{otherwise.} \end{cases}$$

The bias and variance dependencies described for the univariate case likewise apply to the multivariate case, where bias is proportional to h^2 and the second derivative of the distribution, while the variance is proportional to $\frac{1}{Nh^d}$.

9.1.4 Radiance Reconstruction as a Density Estimation Problem

Radiance reconstruction using the photon map is similar, although not equivalent, to the basic density estimation problem. Eq. (9.13) shows DIRSIG's kernel estimator for computation of scattered radiance using the 2-D kernel function shown in Eq. (9.12), and where ϕ_i is photon flux and p is the scattering phase function:

$$K(\vec{x}) = \begin{cases} \frac{1}{c_d} & \text{if } |x| \leq 1; \\ 0 & \text{otherwise.} \end{cases}$$

$$\begin{aligned} c_1 &= 2 \\ c_2 &= \pi \\ c_3 &= \frac{4\pi}{3} \end{aligned} \tag{9.12}$$

$$\begin{aligned} \hat{f}(\vec{x}) &= \frac{1}{Nh^d} \sum_{i=1}^N K\left(\frac{\vec{x} - \vec{X}_i}{h^d}\right) \\ \hat{L}(\vec{x} \rightarrow \omega_o) &= \frac{1}{h^2} \sum_{i=1}^N \phi_i p(\vec{x}, \omega_i \rightarrow \omega_o) K\left(\frac{\vec{x} - \vec{X}_i}{h}\right) \\ \hat{L}(\vec{x} \rightarrow \omega_o) &= \frac{\sum_{i=1}^{k \in h} \phi_i p(\vec{x}, \omega_i \rightarrow \omega_o)}{\pi h^2} \end{aligned} \tag{9.13}$$

The variable N is hidden within the flux term since DIRSIG creates a photon map where each photon (in the map) has a constant flux [*e.g.* the flux is equal to the sum of the flux from all sources, divided by the number of photons shot into the scene (Goodenough 2008)]. However, Eq. (9.13) is dissimilar to the density estimation problem in three major ways. First, the reconstruction is not a probability density function, but instead a

radiance function. Second, the samples have unequal weight. This is because the radiance reconstruction at a point is actually the reconstruction of the *outgoing* radiance as a function of the *incoming* radiance weighed by the *scattering phase function*. The nature of the scattering phase function and the direction of each incoming photon means that each photon can have a variable weight contribution to the outgoing radiance. Finally, Eq. (9.13) is employed repeatedly within each ray sub-segment that comprises a full path from detector element to scene (dictated by reflections, refractions and transmission events that are not necessarily a function of the actual photon map). The sensor reaching radiance is then the result of repetitive in-scattered radiance estimates combined with traditional raytraced radiance computations.

Aside: Note that Eq. (9.13) illustrates that varying photon power, although not changing the expected value of the radiance, can increase variance (e.g. the variance in the estimator will be proportional to the variance in the photon power). This is the motivation behind importance sampling and equal photon flux.

To be a consistent estimator, both the variance and bias must approach zero as $N \rightarrow \infty$. This can be accomplished when the number of photons in the map are increased, and the search radius is decreased (but at a lower rate) (Suykens-De Laet 2002). However, the rate of convergence will depend on the high frequency details in the radiance distribution (e.g. shadows and caustics).

In remote sensing applications, the user must decide how much bias and variance is acceptable in the reconstruction, for some level of efficiency in scene simulation. This will vary depending upon the eventual application, the scene-geometry constraints, and the expected degree of high frequency detail in the radiance distribution of the simulated image. Three examples will be illustrated that show how the sensor, the application, and the inherent radiance distribution within the scene impact the expected and/or desired attributes of the photon mapped scene:

1. **Sensor:** The bias introduced by h (sometimes termed bin width or bandwidth, and defined as the spectral search “radius” in DIRSIG) is analogous to the blur induced by detector elements. If the detector element size and sensor-altitude are fixed, then no apparent gain will be obtained if the photon mapping bias is reduced below that generated by the size of the pixel projected onto the ground (water) plane of interest.
2. **Application:** The eventual application can dictate the level of bias and variance that is acceptable in a given simulation. For example, both the bias and variance can be problematic when attempting to extract target signatures for eventual target detection routines. This is especially true if statistical target detection routines will be used, which rely on the mean and covariance of a given target. Conversely, for water quality analysis, averaging over several pixels is acceptable, which reduces the variance, and for uniform environments, will not introduce unacceptable bias.

3. **Inherent Radiance Distribution:** The inherent radiance distribution in the scene will be a function of the overall illumination conditions, the medium IOPs, and the scene geometry. For example, a floating object will introduce a shadow. The nature of the shadow will be a function of the sun position, and the degree of downwelling illumination (e.g uniform downwelling illumination in a hazy scene will diminish sharp radiance features). For a slow and smoothly varying radiance distribution, bias will be a minimal problem, allowing variance reduction to take primary importance. Conversely, for a scene with high frequency radiance content, such as a caustic or the aforementioned shadow, the bias will be significant if the kernel blurs the radiance feature of interest.

For a given set of IOPs, illumination conditions, and scene geometry, the user can define the number of photons to be used in the map, and the width of the smoothing window h . For a given scene, the impact of these user-defined variables can be compared against the variance and bias in the simulated image, allowing a user to decide what is acceptable for a given application. The user defined variables will be discussed in greater detail in the next section, including illustrations of indirect ways a user can impact the photon map performance.

9.2 User-Defined Photon Mapping Variables

The user can directly and indirectly influence the photon mapping process through the user-defined variables of *max photons*, *bounding box*, *spectral search radius (SSR)*, and *core search fraction (CSF)*. The max photons and bounding box variables directly influence the first pass, and determine the total volume density of the map. Conversely, the SSR and CSF variables influence the second pass radiance reconstruction estimate. Indirectly, the user influences the process by virtue of creating geometry, and specifying illumination, material and optical properties for the scene. For a fixed scene, the quality of the simulated imagery hinges on the user’s ability to appropriately determine the number of photons to be traced, and the smoothing of the estimate around the reconstruction point.

Since it is assumed that an *application* will dictate all scene related variables (e.g. scene geometry, material and optical properties, sensor geometry, etc.), then the user’s understanding of the photon mapping process *within a given scene*, as well as the impact of user-defined variables, determines the overall image quality. Even for a simple scene, this requires a great deal of intuition.

Each pass in the photon mapping process impacts image quality, where quality will refer to noise (variance) and bias in the radiance reconstruction. The impact of the second pass and recreation of the in-scattered radiance distribution using discrete samples has already been discussed. However, the first pass must produce a physically-valid description of scattering events if there is any hope of a physically-valid reconstruction. Computationally

this depends on several factors, some of which the user has direct control over. (Note: the issue of appropriate sampling (*e.g.* of sources, the scattering phase function, etc.) has already been fully investigated and will not be re-evaluated here. Instead, the reader is referred to Goodenough (2007) (Chapter 4) for a detailed discussion). Some variables are obvious contributors to the reconstruction, but there are several less obvious ways a user can impact image quality. The following list provides definitions and basic insight into the indirect and user-defined variables that describe the photon mapping process.

1. Bounding Box

The bounding box defines the volume to be photon mapped. As a rule-of-thumb, it should be slightly larger than the actual volume to be mapped, and/or the projection of the sensor onto the ground plane or water depth that the user wishes to render. For a fixed number of photons, this ensures that the photons comprising the map are visible to the sensor. Photons mapped outside the sensor's field of view do not contribute to sensor reaching radiance, but consume computational resources and dilute the photon volume density for a fixed maximum number of photons. However, if the size of the bounding box is too restrictive, then the edges of the simulation will appear erroneous since the map will not include rays that were originally traced outside of the sensor's field of view, but which can be reasonably expected to scatter into the sensor's field of view.

2. Max Photons

The number of max photons determines the overall volume density. The larger the volume to be mapped, the greater the number of photon scattering events that are necessary in order to build up a sufficient volume density. The user defined variable of max photons does not define the number of uniquely source-generated rays, but rather the number of unique scattering events. For example, if a user sets this variable to 50,000 photons, the result is a map with 50,000 scattering events, although any one uniquely generated ray (traced from source to medium) can contribute any number of scattering events. The maximum number of scattering events per ray is also set by the user (max bounce count). Ideally, this is an infinite number, if not the most appropriate computational setting. This is not to suggest that the photon will absolutely undergo the maximum number of allowable scattering events, since it may otherwise be absorbed during its transit within the participating medium, or after interaction with other pieces of geometry.

The variable max photons is not to be confused with the variable *max source photons*. The latter describes the maximum number of rays to be generated at the source during the first-pass. It should be set much larger than the max photon variable. Note that if the geometry or inherent optical properties of the medium are such that the max source photons is reached before the map is complete, the photon map will

be truncated and contain fewer photons than expected. The user should examine the DIRSIG standard console output if in doubt since he or she might not otherwise be informed that the truncation has happened (Goodenough 2008).

3. Inherent Optical Properties (IOPs)

The IOPs are used to determine the probability of a scattering event versus an extinction event (as well as the probability of an absorption event per extinction event). A high single scattering albedo (w_o) means that the extinction process is dominated by scattering events. The greater the number of possible scattering events, the faster the photon map will be generated. For a high co-albedo, the opposite is true. The absolute value of the scattering and absorption coefficient is also important. A photon map that describes a large volume accompanied by a small scattering coefficient will be more time consuming to generate than a photon map that describes a small volume accompanied by a large scattering coefficient.

4. Source-Scene Geometry and the Bounding Box

The geometric design of the source-scene relationship can influence the importance sampling process. For example, if an off-nadir high radiance source is illuminating a scene consisting of various objects floating or otherwise occluding a water volume below, then many photons from the importance-sampled high-radiance source will not reach the water volume, and instead impact the floating objects. These photons, cast from the important source, will not contribute to the photon map, so that additional rays must be cast until a sufficient number of generated photons contribute to the photon mapping process. For a specific geometry, the user can influence this process when setting the bounding box that describes the portion of the medium to be photon mapped. If the bounding box includes very little unoccluded water, the first-pass may never terminate. If the user increases the size of the bounding box, the time required to build the map may improve, but for a fixed number of photons, the image quality will deteriorate due to noise.

5. A Photon's Sphere of Influence

During the second pass reconstruction step rays are traced from detector elements into the scene. If a ray moves inside the photon mapped participating medium, the photon map is queried to solve for the in-scattered radiance. DIRSIG searches within a cross sectional area around a propagating ray to collect photons and estimate the in-scattered radiance. The size of the cross sectional search area is determined by the user-defined variables of spectral search radius (SSR) and core search fraction (CSF). *The radius defines a photon's sphere of influence, as well as the radius around the propagating ray in which photons will be considered to contribute to in-scattered radiance.* If the photon's sphere of influence intersects the propagating ray, then the photon is within the ray's search volume of interest and is used in the in-scattered

radiance estimate. For the selected photons, the map stores the photon’s direction of travel (among other things). The photon’s direction of travel is weighed against the ray’s direction of propagation using the scattering phase function, which determines the magnitude of the photon’s flux that is scattered into the direction of interest.

The subtle difference between the SSR and CSF is important. The SSR is always greater than or equal to the size of the core search radius (*e.g.* the CSF is a *fractional* value of the SSR). Photon’s within a propagating ray’s SSR are used to *update* the spectral character of the photons found in the CSF. An example will illustrate this more clearly. Consider a small CSF within which 100 photons are found. The spectral character of the 100 photons will depend on the radiance distribution of the sources, as well as the spectral properties of the medium and any other objects within the scene. We will assume that this spectral character is a slowly varying function across the scene. As the number of spectral bands increases, the fraction of photons within any spectral band may decrease. This means that the spectral distribution of the photons in the core may not adequately represent the spectral characteristics of the scene. On the other hand, the spectral distribution estimated from a larger search radius is likely to be more representative of the scene (at least for a slowly varying radiance distribution). Therefore, the spectral content of the photons in the SSR is used to update the spectral characteristics of the core. For example, of the photons in the CSF, all 100 may be “red”. However, if the SSR contains 200 photons, divided among 100 “red”, 50 “green” and 50 “blue”, then the spectral character of the core photons will be updated to 50 “red”, 25 “green”, and 25 “blue” photons [or 50%:25%:25%].

For a scene with very slowly varying spectral character, the SSR can be made arbitrarily large, and in turn, the *updated* spectral content of photons in the CSF will be a better estimate of the true distribution. However, for a scene with rapidly varying spectral characteristics, a large SSR will average out the spectral quality of smaller details, and produce spurious results. Likewise, for a scene with very slowly varying spectral and (and panchromatic radiometric) character, the CSF can be made arbitrarily large, and in turn, the density estimate will depend on a greater number of samples, thereby decreasing the variance.

9.3 Evaluation of Image Quality

The preceding discussion illustrates how the addition of probabilistic in-scattered radiance to DIRSIG simulations adds a noise and bias component not paralleled in solutions that are based solely on raytracing. An *ideal* photon mapped simulation would trace an *infinite* number of photon histories and sample those histories using an *infinitesimally* small sampling interval. Clearly, this is unrealistic given efficiency considerations and computa-

tional constraints. As such, the criteria for a *successfully* photon mapped scene must be readjusted. A reasonable compromise is to demand that a *sufficient* number of photon histories are traced in order to accurately describe the underlying scattered radiance distribution, and that the sampling interval be *small enough* to negate the introduction of considerable bias in the result. Unfortunately, there is no mechanism by which a user can know *a priori* how many photon histories will be sufficient to describe a particular scene, and how those histories should be sampled when computing in-scattered radiance.

The Computer Graphics community has found several viable solutions to this problem using a host of adaptive sampling techniques that vary the number of samples, and/or the sampling interval, in an effort to reduce both bias and noise in the final simulation. Computationally, any one of these solutions would be of great benefit to a DIRSIG user that employs photon mapping during scene simulation. Unfortunately, these existing solutions can severely compromise simulation efficiency, especially if generating a scene with modest (or significant) spatial and spectral content.

One way to buffer the simulation efficiency is to place a limit on the number of adaptive samples that are evaluated for convergence. Instead of evaluating the sample size and sampling interval at all locations in the scene, the user can limit the convergence evaluation to pre-selected discrete locations of interest. This of course assumes that the global solution can be inferred by the solution that satisfies the discrete locations. When reasonable, this approach strikes a compromise by generating a first-order estimate of the photon mapping parameters, without the efficiency burden associated with scene-wide adaptive sampling techniques.

At these discrete locations, the convergence evaluation can take on many forms. One option is to continue expanding the search radius until the computed flux becomes constant, which is assumed to be the search area that captures the minimum number of photons required to obtain a reliable solution. The search radius is then further expanded until the flux again changes. This larger radius captures a variation in the underlying scattered radiance distribution which should not be averaged *across*, and the optimal radius is somewhere between these two extremes. Another option is to fix the search radius, and then continue to store photons until the computed flux becomes stable. This assumes that the pre-selected search radius will introduce minimal bias, and instead, the adaptive process focuses on driving down the variance in the solution.

Although any convergence evaluation will be based on assumptions and associated strengths and weaknesses, the convergence evaluation proposed here is based on determining the scattering point spread function (PSF) and signal to noise ratio (SNR) associated with an edge spread function (ESF) that is strategically positioned within the scene. Since the medium's PSF represents the diameter over which a point source would be dispersed due to scattering within the volume, it is reasonable to assert that this diameter will be related to the search radius over which samples can be averaged, and yet not introduce a significant increase in additional bias. This is true, provided that the bias in the search

radius *outweighs* the bias associated with the pixel’s ground instantaneous field of view (GIFOV), and that the bias is both *detectable* and *significant* to the solution.

Based on this assertion, the minimum SR of interest would be equal to the size of the pixel, projected to the depth of interest in the scene. This represents the minimum diameter over which the detector will blur a point source (assuming a detector-limited system modeled by pixel dimensions, although if the system were optics-limited, and/or the user was aware of an underlying system PSF, then these metrics could easily be substituted). This means that the minimum measurable PSF is that induced by the detector, and therefore scattering PSFs less than this will not be separable from the detector induced bias. The search radius can then be expanded and the PSF recomputed. If the PSF does not change, then the expanded SR does not introduce additional bias into the medium’s scattering PSF, and the SR can be further expanded until it interferes and degrades the scattering PSF.

Another point to consider is the Radiometry Solver associated with the actual surface of the step function. If this solution is based on ray-tracing, then it is assumed that the surface Radiometry Solver will introduce minimal *additional* variance into the estimate. Conversely, if this solver also employs a photon map, then the surface photon map’s search radius must also be considered as a contributor to both variance and the ESF bias.

9.4 Expected Trends

PSF

It is proposed that the PSF of a scattering medium will be a function of the optical depth (τ), the single scattering albedo (ω_o) and the average cosine (g) of the scattering phase function (SPF), as illustrated in Table 9.1. As noted above, if the PSF is computed using a fixed sampling interval, a fixed search radius for in-scattered radiance and a fixed number of photon histories, then the PSF of the scattering medium can be biased due to sampling and noise-induced uncertainty. Since the computational contributions cannot be avoided, the goal here is to determine the magnitude of the impact that these variables have on the PSF (assuming that photon mapping parameters that generate an unbiased PSF will likewise generate unbiased imagery.)

The scattering PSF of water has been shown to be a good match to the small angle scattering approximation (SAA). This approximation assumes that the majority of the radiation passing through a participating medium continues to travel parallel to the incident beam. This is true of media that have highly forward-peaked scattering phase functions, such as water, and is especially important for underwater image transfer (which assumes reasonable optical depths within which image transfer will occur). Since the standard SAA theory is fully described by Wells (1969) and Zege et al. (1991), it will not be described in

Table 9.1: Theoretical influence of IOPs and sampling on the medium’s PSF.

Parameter	Influence on PSF
Optical Depth	At short optical depths the PSF can be dominated by single scattering events and mimic the SPF. However, as the optical depth increases, multiple scattering ensues and the scattered radiance will become more diffuse, leading to a broader PSF.
Single Scattering Albedo	A fixed extinction coefficient can be reached by many combinations of scattering and absorption coefficients. For a given optical depth, the single scattering albedo determines the rate at which the medium becomes diffuse. A higher degree of diffusivity can be reached at a faster rate when coupled with a higher single scattering albedo.
Average Cosine of SPF	A more peaked phase function must be coupled with a greater number of multiple scattering events before the medium becomes diffuse. Therefore, the average cosine of the SPF, combined with the optical depth and single scattering albedo will determine the rate at which the PSF broadens.
Sampling Interval	For a detector-limited system, the minimum detector bias is the projection of the pixel onto the ground plane. A medium PSF smaller than this bias cannot be extracted from the bias induced by the detector element. However, a larger PSF can be coupled with a larger SR without the introduction of considerable bias.
Search Radius	The search radius used to compute in-scattered radiance contributions generates bias in a manner that is similar to the averaging that occurs as a function of detector sampling.
Photon Map Density	An adequate number of photon histories must be traced in order to accurately reproduce the medium’s PSF. If too few photons are used, then the PSF can appear noisy and unstable.

detail here, and instead, we will only be concerned with practical implications of the SAA. Although the SAA can be used to solve a full RT problem, it is nearly exclusively used for PSF and modulation transfer function (MTF) evaluation as it relates to image transfer in scattering media. The MTF is often recast into a convenient range-independent representation (Wells 1969) termed the decay function (Mertens and Replogle 1977) and/or the decay transfer function (DTF) (Hou et al. 2007b). The DTF assumes that the fractional rate of decay of a frequency is proportional to the distance over which the light propagates within the medium (Jaffe 1995). This means that the DTF is a description of scattering at a given range, which can be used to determine the MTF at any arbitrary range. Experimental field observations versus predicted DTF range independence were evaluated by Voss (1991b) using PSF measurements at various depths in the Sargasso Sea during December of 1989. Sensor-source ranges of 14.5m, 21.8m, 25.4m, 61.7m, 74.6m, and 87.7m were examined. The results indicate good agreement with the SAA predictions (both in range-independence and predictive behavior – *e.g.* the DTF at 87.7m can be used to predict PSFs at the aforementioned sensor-source ranges, noting that the Sargasso Sea is very clear with a beam attenuation that averages between 0.082 and 0.074 m^{-1} from the surface down to a depth of 80m). Field/tank experimentation also supports the assertion that if the PSF at several attenuation lengths can be determined for a given water type, then an empirically derived model can be used to calculate the PSF at other attenuation lengths, which confirms the existence of a range-independent DTF (Voss 1991a).

Of additional interest is how well a given empirical model follows the small angle scattering approximation for ocean waters, and how this relates to numerical Monte Carlo models. Hou et al. (2008) conducted such a comparison by examining various PSF models for natural waters. The models will hereafter be referred to as **Duntley**, **Dolin**, **Wells**, **Hou** and **Monte Carlo**, noting that the **Wells** method is essentially the SAA method as described in Hou et al. (2007c). The **Monte Carlo** method uses physically measured IOPs obtained from observations in a experimental tank (pool) scenario. The experiment used tap water, Maalox and absorption dye. IOPs and the volume scattering function were given as input to the **Monte Carlo** model which produced numerically evaluated beam spread functions (BSFs) and PSFs.

Of specific interest to the research presented here is that Hou et al. (2008) reports that the **Wells** model matches the **Monte Carlo** results very well, and that the results are a better match than the empirical **Duntley** model. This includes the maximum measured single scattering albedo of 0.89, at an optical depth of 7.3, and out to angular measurements of nearly 60 degrees.

Therefore, the SAA appears to be a valid description of numerical Monte Carlo and empirically derived PSFs (at least over reasonable optical depths with $\tau \approx 7.3$ (Hou et al. 2008) or $\tau \approx 7.6$ (McLean and Voss 1991) in which image transfer will be of significant interest). Moreover, the results indicate that the scattering PSF is fairly narrow for optical depths in which the SAA is a good predictor of behavior (and within which one can expect significant reflected contributions to sensor reaching radiance). These results, along with intuition, suggest six major PSF considerations:

1. The scattering PSF *for geometrically or optically shallow depths* will be smaller than most remote sensing GIFOVs.
2. Since water's scattering phase function is very forward peaked, the backscattered contributions to sensor reaching radiance will be minor at short optical depths when reflected contributions from in-water objects can be expected to dominate.
3. The SAA theory suggests that small angle scattering events can be recast as direct radiance, so that only "large" angle scattering events contribute to a widening PSF. This implies a range in which the central peak of the PSF will remain relatively constant, while the tails of the PSF widen.
4. DIRSIG has adopted the SAA mentality in that small angle scattering events are also recast as direct radiance (see Section 6.4.7). The motivation for this construct within DIRSIG was the integration/compatibility of photon mapped and ray-traced solutions (*e.g.* using a photon mapped volume with a ray-traced surface Radiometry Solver like GRS). This suggests that the main peak of the PSF will remain relatively constant over reasonable optical depths, while the tails will begin to widen with additional scattering.

5. All proceeding points suggest (that at optical depths of interest) the central peak of the PSF will be small, as will the backscattered contribution to sensor reaching radiance. This means that the user can expect minimal bias to accompany increasing volumetric search radii. If both volumetric and surface photon mapping are employed, then the surface search radius is expected to induce greater bias than the volume search radius (*e.g.* when the reflected radiance dominates the sensor reaching radiance). To overcome this, the DIRSIG surface photon map uses an adaptive search algorithm that amounts to averaging *along* rapid changes in radiance, rather than *across* them, thereby reducing any bias that can be expected from the surface photon map Radiometry Solver.
6. If these postulates are found to be true, then the user can use a relatively large search radius that is several orders of magnitude greater than the GIFOV and still expect minimal to low bias. Moreover, the increasing search radius will drive down the variance in the solution and increase the SNR. The only caveat is that the solution will converge to the true distribution convolved with the kernel function. So although the SNR will increase, the data may show visual artifacts of the convolution; this should be lessened by increasing the total number of photons used to construct the map.

SNR

As with the PSF, it is believed that the SNR will be a function of the optical depth, the single scattering albedo and the average cosine of the scattering phase function, as illustrated in Table 9.2. Computationally, it will also be influenced by the sampling interval, the in-scattered radiance search radius, and the number of photon histories that have been traced. Moreover, the SNR and PSF influence each other. The smaller the PSF, the smaller the maximum allowable SR (assuming that a biased PSF significantly degrades the solution), and the greater the variance of the data for a fixed total number of photons. Conversely, the larger the PSF, the larger the maximum allowable SR, and the lower the variance of the data for a fixed total number of photons. Naturally, a minimum SNR must be obtained in order to accurately capture the characteristics of the PSF, and allow for a visually pleasing result. However, it is not clear if a threshold exists that represents the minimum SNR associated with a visually pleasing simulation, and how *measurable* changes in SNR impact *perceptual changes* in image quality.

Table 9.2: Theoretical influence of IOPs and sampling on SNR.

Parameter	Influence on SNR
Optical Depth	As the optical depth increases, the direct radiance decreases by virtue of the extinction coefficient, thereby driving down the signal. Depending on the magnitude of the single scattering albedo, the scattered radiance may increase with increasing optical depth, and/or decrease due to absorption. The balance of the signal versus the variance will determine if the SNR rises or falls. The result will also be influenced by the reflectance of the benthic plane.
Single Scattering Albedo	For a fixed optical depth, the single scattering albedo (and the co-albedo) influence the variance of the scattered signal, and the rate at which the direct signal decreases due to absorption. A fixed extinction coefficient can be reached by many combinations of scattering and absorption coefficients, but a higher scattered signal will be associated with a higher single scattering albedo. If we assume a fixed extinction coefficient (and fixed number of photons), then a low single scattering albedo will be associated with less scattered signal, although the instantaneous signal at any given location can vary drastically from its neighbors. Conversely, a high scattering albedo will be associated with higher scattered radiance that is more uniform from one location to the next. The net impact on SNR will be a function of the balance between the magnitude of the direct signal and the variance of the in-scattered signal.
Average Co-sine of SPF	A more peaked phase function can generate discrete areas with high and low SNR, depending on the trajectory of rays within the scene, the single scattering albedo and the optical depth.
Sampling Interval	The size of the pixel projected to a particular geometric depth will determine the minimum area over which the signal will be averaged.
Search Radius	The smaller the search radius used to compute in-scattered radiance contributions, the greater the variance in the signal.
Photon Map Density	Naturally, tracing an increasing number of photon histories will drive down the variance in the signal.

9.5 Scenes with Considerable Spectral Sampling

The user-set photon mapping parameters of SR and total number of photons are band independent. The mean number of photons per spectral band will be a function of importance sampling of the source and the IOPs of the water. Note that DIRSIG is clever and will use the user-set simulation bandpass to restrict the otherwise wasteful generation of spectral photons that are outside of the sensor’s bandpass (*e.g.* if you simulate an RGB image, infrared energy from the solar disk will not populate and dilute the photon map). Multiband simulations require that the user select a single SR and photon map to be used across all bands. However, it is unlikely that the optimal solution for SR and mean photon density is constant across all bands of interest. That being said, it is nearly impossible to surmise how the SNR will vary per band. For example, a band with high source emission will be spectrally well-represented in the photon map. If coupled with high medium scattering, there will be a reasonably large number of samples on which to base the in-scattered radiance computation, which implies low variance, but since the

direct signal will be increasingly reduced, the in-scattered samples start to become the greatest contributor to the signal, and the variance can indeed be high, making both low and high SNR possible. Conversely, a band with low source emission will contribute fewer spectral photons to the map. If this band has low medium scattering, there will be fewer spectral samples on which to base the in-scattered radiance, which suggests high variance. Conversely, if paired with high medium absorption, then the scattered signal will be very low, and having fewer photons on which to base this estimate may not be considerably detrimentally. All this is true provided DIRSIG uses the spectral tag associated with each photon when determining the number of photons on which to base the in-scattered computation. On the other hand, a variance-reduction technique has been implemented within DIRSIG, that makes use of all photons when computing the in-scattered radiance in each spectral band. Despite making the interplay of variance and optical properties less intuitive, this clearly drives down the variance in the solution, and the impact of this variance-reduction tool will be evaluated.

In addition, low scattering coefficients and short optical depths mean that the PSF is dominated by single scattering (and therefore it will have a smaller full-width-at-half-max [FWHM]) suggesting the necessity of a small SR, and therefore greater variance, again assuming that the biased PSF is detrimental to the solution. Conversely, large scattering coefficients and long optical depths will be marked by comparatively broad PSFs that are dominated by multiple scattering, which can be accompanied by larger SR and reduced variance. Therefore, an optimal band-specific solution will not be found if using a band-independent SR and SNR. Although this thought experiment seems theoretically sound, the small proportion of backscattered energy in any single band means that an optimal solution on a per-band-basis is likely not worth the user's effort, and instead, there is merit in arguing that the user should configure the simulation parameters to optimize the solution for the spectral band with the greatest scattering (and least absorption). That way, the band with the greatest in-scattered signal will be the least noisy, even if the Radiometry Solver configuration is potentially less-than-optimal for other bands.

9.6 Scenes with Large Spatial Extent

Naturally, the larger the spatial extent of the scene, the larger the total number of photons that will be required. Therefore, it seems logical to simulate large scenes from smaller tiles. If this is carried out, it is suggested that a single initial tile be used to determine the optimal photon mapping parameters, and that these first-order estimates be adopted when simulating all remaining tiles. That way, a large scene can be stitched together without the possibility of an obvious change in SNR and bias. Again, although this seems theoretically sound, it is hard to argue the utility of simulating an expansive scene using

a photon mapped solution, and instead, it is believed that the most reasonable use of the photon mapped solution will be over small scenes that require high resolution and high fidelity, which may very well negate tiling considerations completely.

9.7 Noise and Asymmetry in the PSF

The proposed algorithm assumes that the PSF will vary (as outlined in Table 9.1) depending on the IOPs and environmental variables of the simulation. It also assumes that this *true* PSF can be accurately determined from somewhat noisy data. Generating a reasonable PSF from a noisy ESF is not an easy task since even small fluctuations in the data are exacerbated when the derivative is taken. The problem can be partially mitigated by taking several transects across the edge and using the average, and/or assuming a parametric form for the data. This will be especially important for oblique views which can lead to asymmetric PSFs. In fact, it is suggested that the user place the edge function so that it is centered along or near the optical axis of the sensor to eliminate further complications.

9.8 Theoretical Proof of Concept

The trade study will determine the utility of the proposed metrics (bias and SNR), and generate secondary data that can be used to investigate the relationship between (1.) PSF, (2.) SNR, (3.) optical depth, (4.) single scattering albedo, and (5.) the average cosine of the scattering phase function. It is likely that proof of concept will be achieved before a comprehensive study of these relationships can occur, but when useful, attempts will be made to investigate any findings, trends and relationships that may surface regarding SNR, variance, and the FWHM of the PSF with the aforementioned environmental variables.

The trade study will have two specific goals. First, to determine the degree to which the PSF bias varies with the user-defined search radius. This assumes that the computational degradation of the scattering PSF will in fact translate into quantifiable and/or visually significant imagery bias (which is yet undetermined given the fact that the scattering PSF is fairly narrow at shallow optical depths of interest, and knowing that DIRSIG recasts small angle scattering events as direct radiance). If PSF bias turns out to be relatively negligible across several SR settings, then all future simulations can be rendered using a relatively large search radius (barring the generation of convolution artifacts). Although possible, it would be very inefficient to generate *truth imagery* for a large number of IOPs,

where *truth imagery* is defined as a simulation generated with nearly infinite samples and an infinitesimally small sampling interval so as to ensure zero computational induced bias and noise. Instead, a succession of images will be evaluated, with constant IOPs and geometry, but varying SR, allowing for *relative* changes in bias to be evaluated. These relative changes will be based on quantifiable changes in the ESF and PSF. This will determine the relative imagery bias and how it relates to the bias in the PSF.

The second goal is to empirically evaluate how computed SNR relates to visually pleasing simulations. This is not an easy task since it is not clear if a threshold exists that represents the minimum SNR associated with a visually pleasing simulation, and how *measurable* changes in SNR impact *perceptual changes* in image quality. Since a thorough psychophysical investigation of this issue is beyond the scope of this research, visual image quality will be assessed by a simple subjective ranking of images (of a single scene), and this ranking will be used to examine any possible trends. Of course, the user is encouraged to set mathematical bounds on the SNR given his or her goals, rather than a visual assessment of the imagery, and these goals may be based on any number of external factors, such as reliable target detection or detector noise considerations.

9.8.1 Simulation Trade Space

Spectral

The analysis will begin with a panchromatic image between 0.4 and 0.7 microns, and using IOPs that are constant across all wavelengths. If successful, wavelength dependent IOPs and specific band passes will be analyzed. A few bands between 0.4 and 0.7 microns will be considered sufficient to illustrate the application to multi- and/or hyperspectral data.

IOPs & Atmosphere

Three sets of IOPs will be considered; clear coastal, bay and extremely turbid water. Hou et al. (2007c), Jaffe (1992), Jaffe (1995) and Reinersman and Carder (2004) provide coefficients for various water types with $\sigma_a = 0.038, 0.013$ or 0.041 and $\sigma_s = 0.35, 0.1448$ or 0.211 for clear coastal water, $\sigma_a = 0.143$ and $\sigma_s = 0.180$ for bay water, and $\sigma_a = 0.27$ or 0.115 and $\sigma_s = 0.95$ or 0.839 for turbid water. Table 9.3 shows the water IOPs adopted for this trade study.

The somewhat realistic single parameter Henyey-Greenstein phase function will be employed, varying the average cosine between 0.85 and 0.97. This will be compared to with the Petzold scattering phase function. A simple atmosphere will be used, with a constant solar to skylight ratio.

Table 9.3: Absorption and scattering coefficients for trade study.

Water Type	Coefficients (m^{-1})		
	Absorption	Scattering	Total Extinction
Clear, Coastal (Group #1)	0.038	0.350	0.388
Bay (Group #3)	0.143	0.180	0.323
Turbid (Group #4b)	0.268	0.680	0.948

Air-Water Interface

Variations in the interface will be problematic. Sinusoidal and facetized waves will introduce refractive shifts in the ESF. Although the interface determines the refracted direction of photons inside the medium (and therefore caustics), if variations in direction occur over a sufficiently small spatial interval, the shifts introduce uncertainty in the actual location of the edge, and can degrade the user’s ability to evaluate the scattering PSF. If the edge is obvious, then shifts can be removed before PSF evaluation, but this would become ever more problematic with increasing turbidity and decreasing SNR.

An alternative solution would be to use the algorithm with the detector submerged below the interface, evaluate the scattering PSF and SNR, obtain the first-order estimates of the photon mapping parameters, and then reuse the map, along with the photon mapping parameters, when simulating the imagery with the detector above the interface. Since this will be inefficient to evaluate during proof of concept, it will be considered a logical extension of the proposed study, not requiring independent validation. However, this obstacle can (hopefully) be seamlessly overcome if the adaptive sampling is integrated within DIRSIG, making use of *observers* (or virtual detectors) that can be arbitrarily positioned within the scene, all the while positioning the true detector above the interface.

The user also has the alternative option of using procedural subpixel microfacet capillary waves. This type of interface is likely to increase the variance as a function of importance sampling, and further blur the ESF, generating a larger medium+interface PSF, rather than introducing shifts in the ESF. However, for reasonable wind speeds within which image transfer is expected (*e.g.* less than approximately $10 \frac{m}{s}$), the interface induced bias of the ESF is believed to be less than most remote sensing GIFOVs. In addition, simulations employing procedural capillary waves are extremely lengthy to generate, and require additional subpixel and adaptive sampling by the sensor. Therefore, in the interests of time, and based on intuition regarding the expected bias and GIFOV, subpixel capillary wave induced bias will not be investigated using DIRSIG.

9.9 Alternative Uses, Outcomes & Considerations

SNR Priority

If the user is more concerned with noise, and/or has *a priori* information that indicates the scene will be void of rapid fluctuations in in-scattered radiance, it may be useful to fix the SR and opt to maximize SNR determinations only.

In-Scattered Radiance is a Minor Contribution to Bias

If in-scattered radiance is found to be a minor contributor to scene bias (across the entire trade space described here), then any future adaptive sampling algorithm to determine search radius will be unnecessary. Instead, the SR can be fixed for all future simulations. Luckily, this possible SR-constancy will be apparent based on the methodology employed above. Therefore, failure (or disproof of concept) will still be deemed beneficial, and in fact, equally useful to future simulations. If true, then SNR considerations will dominate the photon map user-settings. The fact that increasing SR will generate a larger convolution kernel will not be investigated here. However, it is acknowledged that should increasing SR lead to minimal bias, future investigations may need to balance SNR with artifacts of the convolution.

Approximate Relations Derived from Density Estimation

The following recaps the relationships that were described in preceding sections regarding density estimation, as well as additional considerations that are relevant to the photon mapped solution. If the trade space data warrants, tentative trends in bias and variance will be provided.

- In traditional 2D density estimation the bias is proportional to the square of the radius and the second derivative of the density function. Since the radiance distribution is expected to vary with the character of the scattering phase function, the relationship between bias and water IOPs will likely include consideration of the average cosine of the phase function and the scattering coefficient (and or the single scattering albedo).
- In traditional 2D density estimation the variance is proportional to $\frac{1}{Nr^2}$. This assumes that each additional sample is equally weighted, which is not true in the case of photon mapping with a forward peaked phase function. Instead, the significance of new samples depends on both the scattering phase function and extinction (or

optical depth since a sample at depth A does not make an identical contribution to sensor reaching radiance as the same sample but at depth B).

9.10 Integration within DIRSIG: The Virtual Step-Function

The goal of this trade study is to determine how the user-set parameters of SR and mean photon density impact the perceived image quality in terms of bias and SNR. Depending on the results, it may be reasonable to propose that DIRSIG photon mapped solutions simulate a *virtual step function* within the user-set scene geometry, and utilize *observers* or virtual detectors capable of monitoring the ESF (and with arbitrary placement either above or below the water interface, and or having the ability to by-pass the interface).

In practice, the user would initiate the DIRSIG simulation, flagging the simulation to carry out adaptive photon map sampling over the virtual step function. The user would be required to provide a minimum SNR, and the geometric depth at which DIRSIG should generate the *virtual* ESF. Ideally, the user would also set a maximum for both the SR and total photon count, providing the algorithm with an exit strategy should something within the simulation be problematic. Once the simulation process is initiated, the algorithm would adaptively vary the photon mapping parameters until the bias and SNR associated with the step function converges to some user-supplied minimum threshold. It is then assumed that the parameters that meet the needs of the user, when evaluated for the step function, will likewise be reasonable first-order estimates of scene-wide parameters that can be used during the second pass of the simulation, which would proceed as usual, in the absence of the *virtual* ESF and observers, and making use of the precomputed photon map.

9.11 Results: SNR

Each water type described in Table 9.3 was paired with four different scattering phase functions (SPF). The first phase function was extremely forward peaked, with an average cosine of 0.97 (a backscatter fraction of 0.0064), and implemented using the one-parameter Henyey-Greenstein (HG) phase function. The second variant used the same analytical HG phase function form, but coupled with an average cosine of 0.92, and a backscattering fraction of 0.01795. The backscatter fraction of 0.01795 roughly matches that attributed to the Petzold (P) particle phase function (derived from measurements taken in the San Diego Harbor (Mobley et al. 2002)), which when integrated between 0 and 90° produces an average cosine of 0.98, and a backscatter fraction of 0.0183 (Mobley et al. 2002). Despite the fact that the backscatter metrics converge for the HG ($g = 0.92$) and P functions, these SPFs have different characteristics and where therefore compared, asserting that a convergent backscatter fraction is insufficient as a single metric to describe the SNR associated with a simulation. The final SPF variant was simulated using the HG phase function, but with a much larger backscatter fraction of 0.0361, and an average cosine of $g = 0.85$.

The SNR evaluation was initiated by first simulating each water type using approximately $1000 \frac{\text{photons}}{\text{m}^3}$. The first three Figures in section 9.15 depict the clear water simulations with a constant phase function type, and increasing backscatter fractions (average cosines of 0.97 (HG), 0.92 (HG), 0.85 (HG), in Figures 9.25, 9.26 and 9.27, respectively). The fourth Figure (9.28) depicts the simulation that employed the Petzold phase function ($g = 0.98$). This sequence of Figures is repeated for the bay (Figures 9.29 - 9.32) and turbid (Figures 9.33 - 9.45) water scenarios in the next set of 8 Figures. Inspection of the Figures shows that a *higher SNR* is associated with the *bay water coefficients* and *greater backscatter returns*, but that the characteristics of the phase function also contribute a large influence over the SNR. This is evident from the lower SNR and greater mottled appearance associated with the Petzold phase function when compared with the Henyey Greenstein with an average cosine of 0.92 (both with backscatter fractions of approximately 0.018), although it is acknowledged that the Petzold phase function (which is based on discrete tabulated values) appears to suffer interpolation or sampling artifacts.

Figure 9.2 plots SNR versus search radius and average cosine for the dark and bright portions of the step function. The SNR increases as the search radius increases, and the average cosine decreases (becomes less forward peaked). This is true for both the bright and dark portions of the step function, although the highest SNR is associated with the greatest geometric depth (5M) for the dark target, and the most shallow depth (1M) for the bright target (note that the SNR ratio is computed using the mean sensor reaching radiance between the target and sensor, versus the standard deviation of that signal, and not the mean signal versus the background or surround).

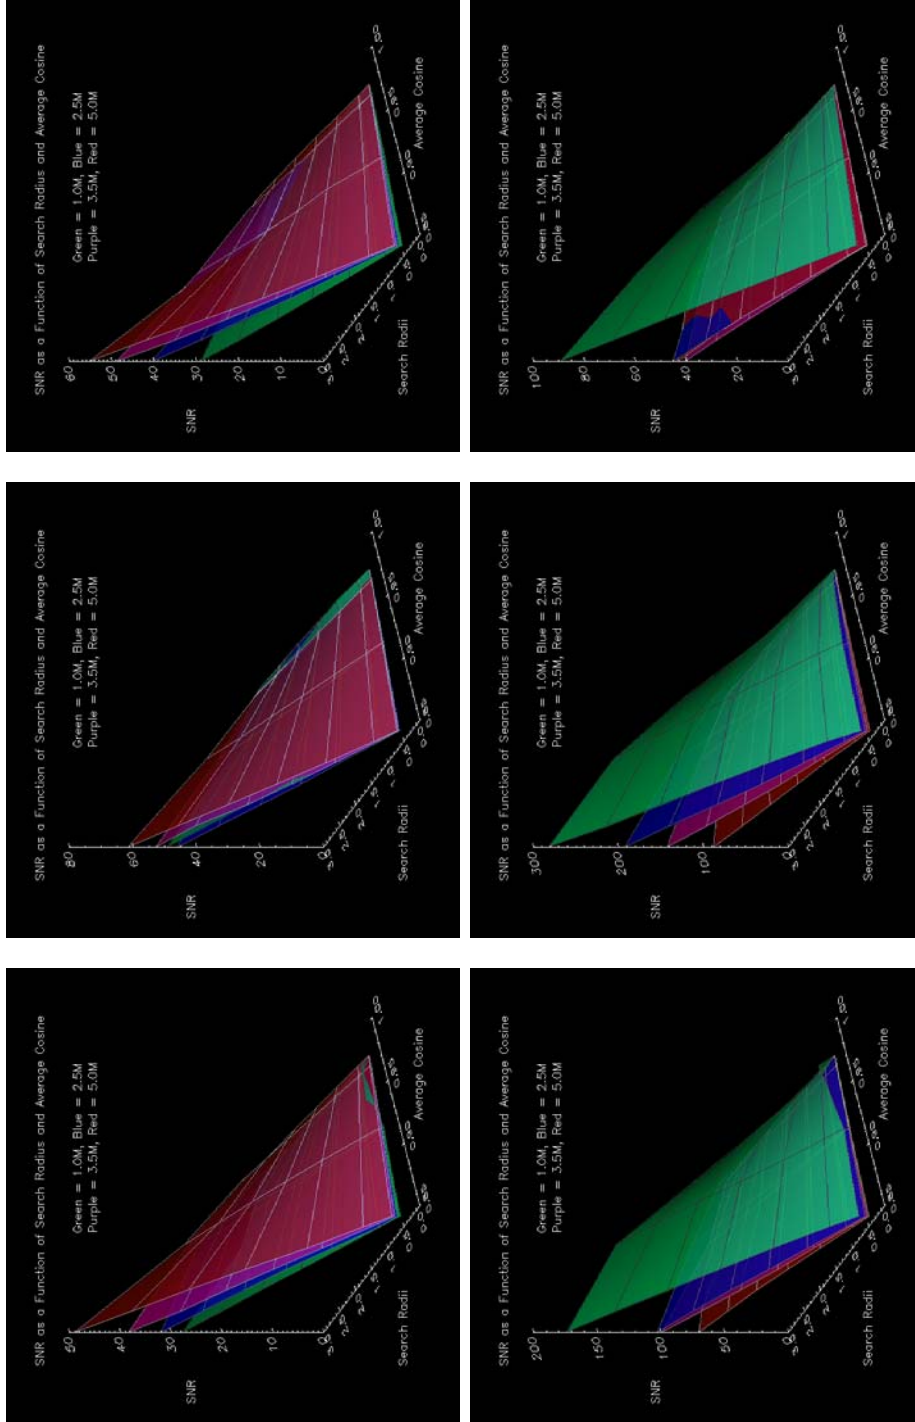


Figure 9.2: Plot of SNR (left to right: coastal, bay and turbid water types) versus search radius and average cosine for the dark (top row) and bright (bottom row) sides of the step function situated at 1.0, 2.5, 3.5 and 5.0m from the water surface (note that the water volume has a total depth of 8.0m between the air surface and ground plane). The plots illustrate increasing SNR with increasing search radius, and decreasing average cosine. Note that targets were placed at fixed geometric depths (rather than fixed optical depths), since it is assumed that a coastal scene (targets, etc.) will be dictated by fixed geometry, bathymetry, etc., but that the water constitutes for that coastal area can vary widely depending on environmental and weather conditions.

Since it is straightforward to assert that increasing SNR can be achieved by either increasing the mean signal, and/or decreasing the signal variance, it is reasonable to assume that the increasing SNR associated with increasing search radius is the product of decreasing variance. This is because the mean signal will be constant, regardless of the search radius, provided the simulation has a constant or smoothly varying radiance distribution, and sufficient photon histories have been traced (*e.g.* the simulation is not photon starved). Moreover, it is reasonable to assert that the increasing SNR for a fixed search radius and increasing backscatter fraction is a function of utilizing a greater number of (more highly probable) samples when computing the in-scattered radiance, coupled with an increasing mean signal.

The change in SNR with depth for bright and dark targets is slightly less intuitive. The backscattered signal will always be relatively small in forward scattering waters. It will also be largely independent of the direct, reduced and reflected signal. For a dark or absorptive object, the mean direct reflected signal is always zero. At shallow depths, the low backscattered signal and zero direct signal make for a low SNR. As the depth increases, the scattered radiance becomes more diffuse and the mean sensor reaching radiance between the target and sensor will increase as shown in Figure 9.3 (top row). The variance also increases, as can be inferred from inspection of the width of the frequency histogram (not shown) and Figure 9.4 (top row). However, the increasing mean dominates and generates higher SNR at greater depths.

For bright objects with high reflected signals (shallow depths), the mean direct signal will be high. When this high mean signal is compared to the low backscattered signal, the SNR will be high. At increasing depths, the direct signal will decrease, while the backscattered signal increases, while overall, the mean signal and variance decrease (as a function of extinction, and more diffuse scattering, respectively). This is apparent in Figure 9.3 and 9.4 (bottom row) which depicts the mean and variance associated with bright targets at variable depths, and summarized in Table 9.4. Collectively, these observations help to explain the change in SNR with depth for bright and dark targets.

Table 9.4: Change in mean and variance with increasing depth.

Target	Direct Signal	Backscattered Signal	Mean Signal	Frequency Histogram*	Variance	SNR
Bright	↓	↑ (\leq Direct)	↓	Contracts	↓	↓
Dark	Zero	↑ (\geq Direct)	↑	Widens	↑	↑

*The range between the maximum and minimum radiance value that accounts for 60% of the data, or the radiance range with a cumulative frequency equal to or greater than 20% and less than or equal to 80%. This Table summarizes the data shown in Figures 9.3 and 9.4.

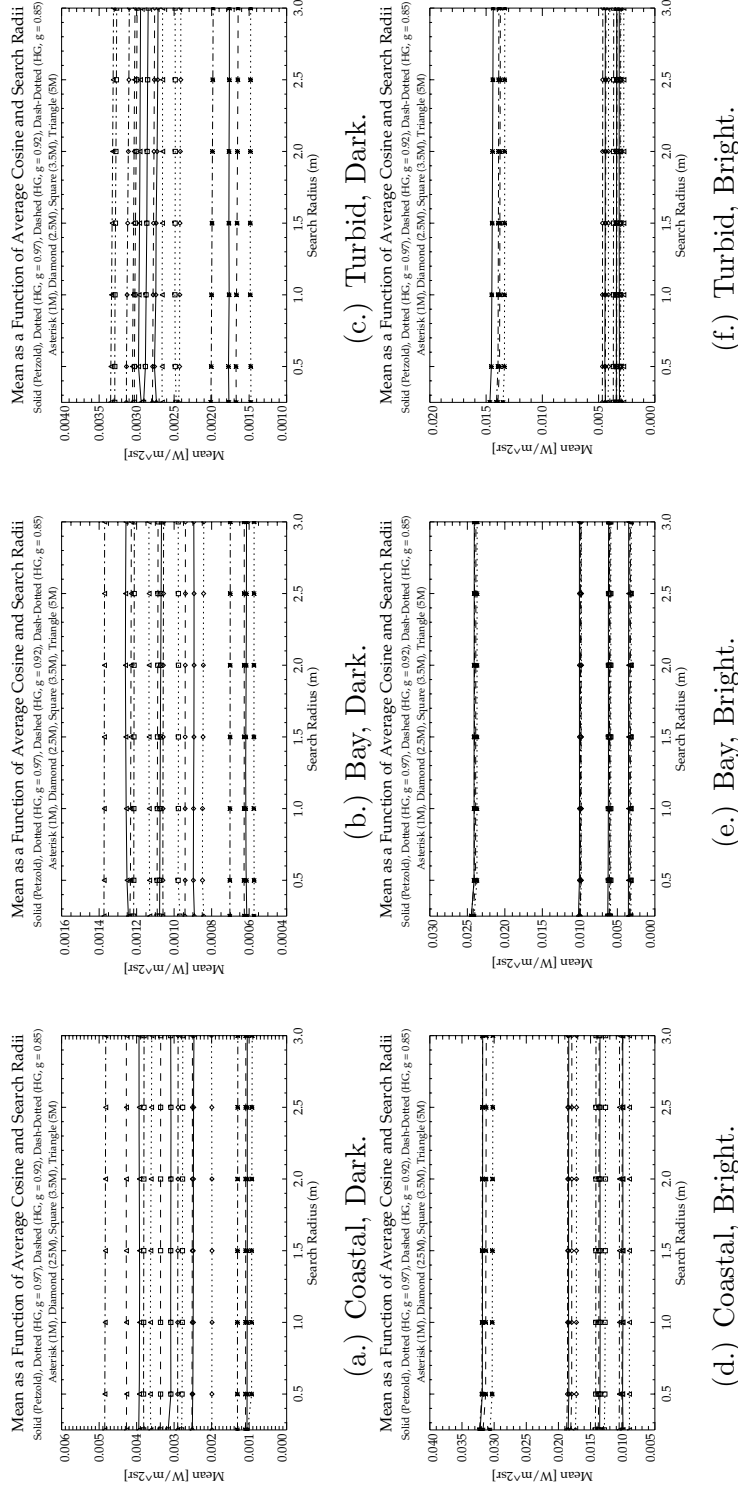


Figure 9.3: Mean sensor reaching radiance for dark targets (top row) and bright targets (bottom row). For both rows, the left most column [(a.) & (d.)] depicts the results of the coastal water scenario (Group #1), the middle [(b.) & (e.)], bay water (Group #3) and the far right [(c.) & (f.)], turbid water (Group #4b). All other variables are as labeled on plot.

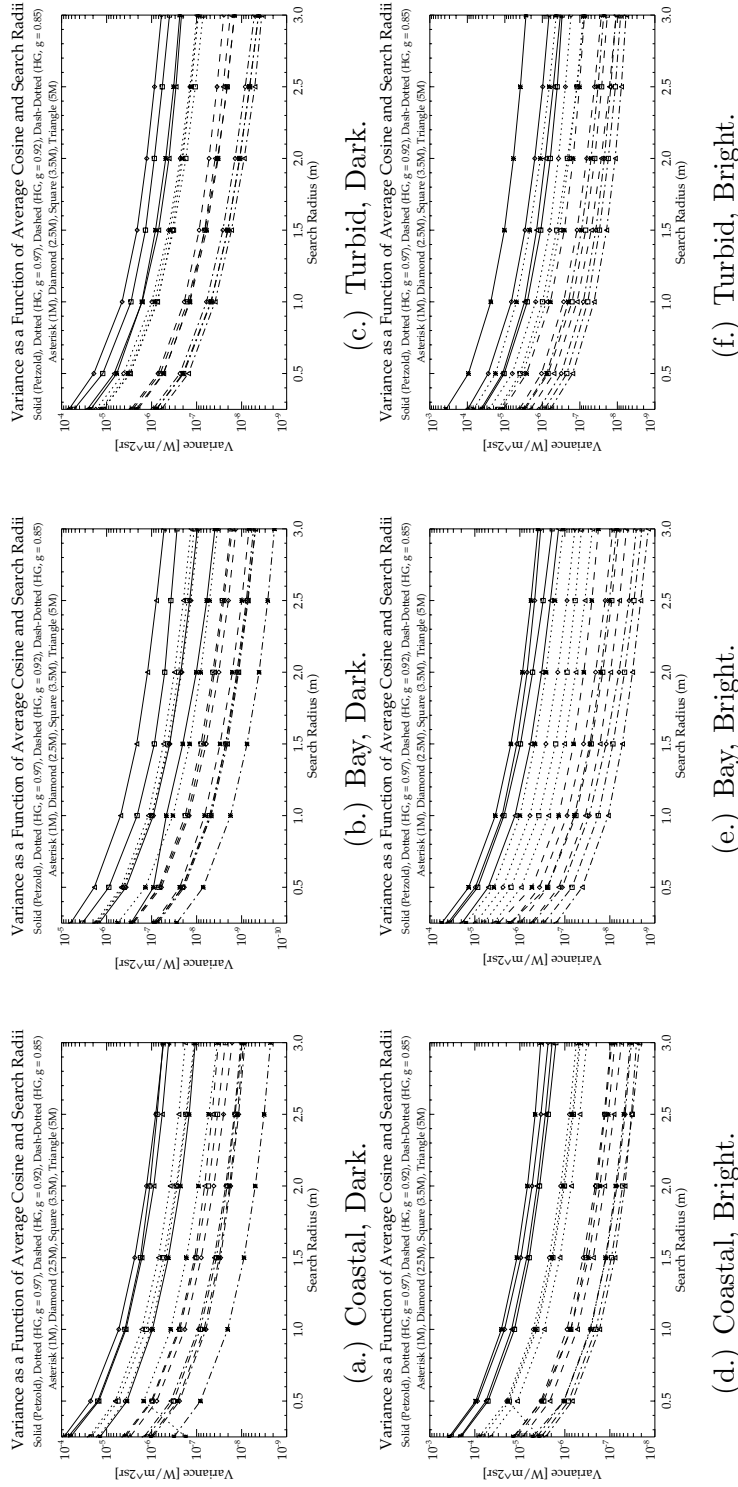



Figure 9.4: Variance in the sensor reaching radiance for dark targets (top row) and bright targets (bottom row). or both rows, the left most column [(a.) & (d.)] depicts the results of the coastal water scenario (Group #1), the middle [(b.) & (e.)], bay water (Group #3) and the far right [(c.) & (f.)], turbid water (Group #4b). All other variables are as labeled on plot.

The spread of the data in Figures 9.3 and 9.4 also provides some insight into the variation in sensor reaching radiance (both signal mean and variance) that can be expected for variations in the backscatter fraction. For dark objects, the highest mean (and lowest variance) is associated with the largest backscatter fraction of 0.035 (for SPF = HG and $g = 0.85$). This is also true for bright targets, except at very shallow depths when the Petzold phase function can generate equal or higher mean sensor reaching radiance, despite its smaller backscatter fraction (0.018). If the errors in mean associated with a larger backscatter fraction can be tolerated (the greater the target’s reflectance, the lower this error will be), it is clearly advisable to simulate data using a higher backscatter fraction and thereby enjoy a larger SNR. This is true regardless of the water’s extinction coefficient (and therefore turbidity), as shown in Figure 9.2.


Since the reconstruction of radiance using photon mapping has been compared to traditional density estimation, it would be interesting to examine the relationship between variance and the density estimation prediction described by Eq. (9.14) where N is the total number of photons used to generate the map, and h is the radius of the search area raised to the dimensionality d of the problem.

$$\sigma_{(N,h)}^2 \propto \frac{1}{N h^d} \quad (9.14)$$


If density estimation were to fully describe radiance reconstruction using photon mapping, then we would expect the variance to be inversely related to the total number of photon histories traced, and the search radius raised to the power (dimensionality) of 2 to 3. The ambiguity regarding the dimensionality of the problem is related to the fact that a dimension of 2 is used to search for photons (*e.g.* the user-assigned search radius is used to compute the search *area*), but this is done repeatedly throughout the volume, forming a “cylinder-like” final search volume, which ultimately leads to a 3D reconstruction search *volume*. Therefore, some power between 2 and 3 seems like a reasonable expectation. This was examined in Figure 9.5 which plots variance and trend lines for HG ($g = 0.92$), with mean photon densities of 9000 (diamonds), 3000 (squares) and 1000 (triangles) $\frac{\text{photons}}{\text{m}^3}$ for coastal, bay and turbid water coefficients with targets at 1.0m below the air-water interface. The dashed lines connect the observed data points, while the solid lines show the computed trend line. The results indicate that the variance is a function of the search radius raised to a power of approximately 2.4 to 3.0. This is an interesting observation, both because it persists across all water types and at all target depths (see additional plots in Appendix B), and because the results seem to suggest that the problem has a dimensionality somewhere between 2D and 3D, as expected. Of course, this is not presented as an absolute indication of a proportionality, since other factors may very well contribute to the observed trends, but the results can certainly be used as a starting point for future investigations.




(a.) Coastal, Dark.




(b.) Bay, Dark.




(c.) Turbid, Dark.



(d.) Coastal, Bright.



(e.) Bay, Bright.



(f.) Turbid, Bright.

Figure 9.5: Variance in radiance (W/m^2sr) and trend lines for HG ($g = 0.92$), with mean photon densities of 9000 (diamonds), 3000 (squares) and 1000 (triangles) $\frac{photons}{m^3}$. The dashed lines connect the observed data points, while the solid lines show the computed trend line.

Eq. (9.14) also suggests that if a ratio is computed that compares the variance at a fixed search radius h , but varying mean photon density N , then the results will be linear and proportional to the ratio of the mean photon density as shown in Eq. (9.15).

$$\begin{aligned}\sigma_{(N_1, h_1)}^2 &\propto \frac{1}{N_1 h_1^d} \\ \sigma_{(N_2, h_2)}^2 &\propto \frac{1}{N_2 h_2^d} \\ \frac{\sigma_{N_1}^2}{\sigma_{N_2}^2} &\propto \frac{N_2}{N_1} \quad \text{when } h_1 = h_2\end{aligned}\tag{9.15}$$

This is plotted in Figure 9.6 for coastal, bay and turbid water coefficients HG ($g = 0.92$), at 1.0 and 2.5m target depths (plots for targets at 3.5 and 5.0m depths can be found in Appendix B). The dashed lines correspond to dark targets, with ratios of 3000/9000 (diamonds), 1000/9000 (squares) and 1000/3000 (triangles). The solid lines are the counterpart for bright targets, with ratios of 3000/9000 (Xs), 1000/9000 (asterisks) and 1000/3000 (circles). The results are relatively linear (although not perfect), and often cluster near 3 (which is an exact solution of $\frac{N_2}{N_1}$ for ratios of 3000/9000 and 1000/3000). The data that clusters between approximately 5 and 9 corresponds to the ratio for 1000/9000, which has $\frac{N_2}{N_1} = 9$. These latter plots perhaps cluster near a multiple of the fraction $\frac{N_2}{N_1} = 9$, although they are less linear than the counterparts that cluster near 3, and may indicate additional contributions to the proportionality that have not been considered here. Again, an absolute indication of proportionality is not being presented, and instead, the results are meant as a starting point for future investigations. If such trends are examined in the future, it is important to consider that density estimation does not directly describe radiance reconstruction using photon mapping, since the latter uses samples with unequal weight. Three of these complicating factors are itemized below:

1. Sample density varies with optical depth (*e.g.* photons are eventually extinguished at sufficient optical depths, so a photon that is scattered at depth A does not make an equivalent contribution to sensor reaching radiance as a photon that is scattered at depth B).
2. Sample magnitude varies with the single scattering albedo (or with water turbidity *e.g.* absorptive events decrease the flux associated with all photons in the map).
3. Sample magnitude varies with the water's phase function for a fixed water type (*e.g.* photons make unequal contributions to the signal since their backscattered magnitude varies as a function of the photon's direction of travel, and probability of being backscattered into the sensor's line of sight).

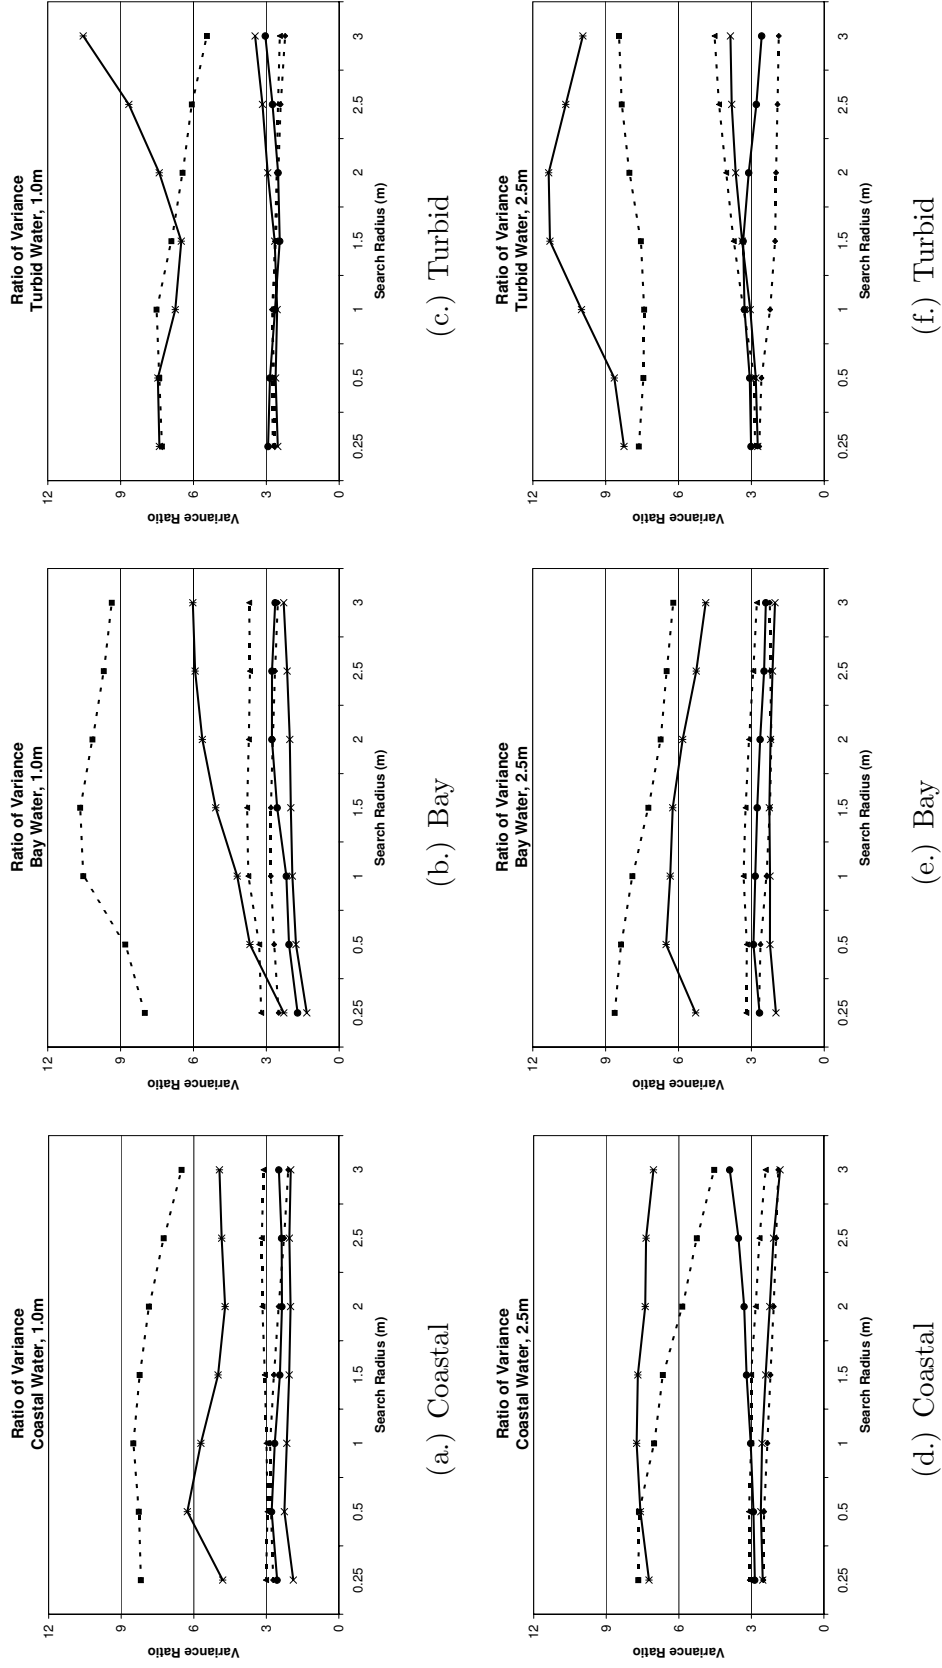


Figure 9.6: Ratio of variance (HG $g = 0.92$) as a function of mean photon density and search radius. Dark targets (dashed lines), with ratio 3000/9000 (diamonds), 1000/9000 (squares) and 1000/3000 (triangles). Bright targets (solid lines), with ratio 3000/9000 (Xs), 1000/9000 (asterisks) and 1000/3000 (circles). Traditional density estimation indicates that $\frac{\sigma_{N1}^2}{\sigma_{N2}^2} \propto \frac{N2}{N1}$.

Essentially, this means that all photons are not created equal, and that the reconstruction of outgoing radiance using incoming and scattered radiance is likely to increase the variance beyond that expected in traditional density estimation problems.

Since the results shown in plots 9.3 and 9.4 are based on targets with fixed geometric depths, selected water scenarios were re-simulated, but this time forcing a fixed optical depth, and using a single Petzold phase function. Figure 9.7 (left most column) plots mean radiance versus optical depth and the single scattering albedo associated with each water type. Since the direct signal is constant for a given optical depth, it is apparent that the mean signal decreases more rapidly when the optical depth of interest is reached via a smaller single scattering albedo (or greater absorption). This is true for both bright and dark objects (keeping in mind that the sensor reaching radiance between a dark (bright) target and the sensor is greater at longer (shorter) optical depths).

Since scattered radiance contributes to the variance in the photon mapped signal, it is logical to assert that a lower single scattering albedo will be associated with less variance. This appears to be roughly true given the trend shown in Figure 9.7 (middle column). However, the actual variance is a function of both the mean and instantaneous signal per pixel, and a higher single scattering albedo should be associated with a higher mean signal. It is therefore acknowledged that the current investigation does not sufficiently describe signal variance and its dependency on other parameters. Finally, Figure 9.7 (right most column) depicts SNR versus the scattering coefficient. The trend indicates that at long optical depths and for bright targets, the SNR varies little with the scattering coefficient, but that the SNR decreases with an increasing coefficient at shallow optical depths. For dark targets, regardless of optical depth, the SNR decreases with increasing coefficients.

Although Figure 9.7 provides some interesting insight into potential causal relationships, the data is in no way meant to replace an exhaustive sensitivity analysis of how water IOPs impact signal mean, variance and SNR. The results shown here, along with those described by Figures 9.6 and 9.5 (and additional Figures shown in Appendix B) are meant to show first-order estimates of trends, and the potential for future studies (either for pure academic interest, or for practical applications should such information be relevant).




Figure 9.7: Mean, variance and SNR versus optical depth and either the single scattering albedo, or the scattering coefficient. All simulations utilized a Petzold particle scattering phase function and a mean photon density of 3000 photons per m^3 . The top row depicts results for dark targets, and the bottom for bright targets. Note that the optical depth is highest nearest the z-axis for dark objects, and lowest near the z-axis for bright objects.

Since the majority of former plots shows variance and mean signal versus photon mapping and/or water IOP parameters, the following plots will actually consider the rate of change in SNR as a function of either the search radius, or the mean photon density (where mean photon density is computed from the total number of stored photons divided by the photon map's bounding box). SNR versus search radius is shown for simulations with mean densities of $1000 \frac{\text{photons}}{\text{m}^3}$ and $9000 \frac{\text{photons}}{\text{m}^3}$ for water which can be described by the Petzold phase function, and the Henyey-Greenstein phase function with $g = 0.92$ (see associated imagery in Figures 9.37 - 9.48). The results are plotted in Figure 9.8 and 9.9 for bright and dark targets, respectively. Although, it is evident that increasing SNR is associated with an increased search radius and mean photon density, this benefit is only reaped at larger search radii (*e.g.* tripling the mean photon density but keeping a small search radius provides very little gain in SNR). Interestingly, the highest SNR is nearly always associated with the bay-water IOPs. In all cases, the highest SNR, regardless of water type, is associated with the HG phase function with $g = 0.92$ (as compared to the Petzold phase function with a similar backscatter fraction). In fact, the Petzold phase function lowers the SNR, and flattens the rate of increase in SNR with search radius. Figure 9.9 shows that for dark targets in water volumes assigned the Petzold phase function, the water IOPs have very little influence on the SNR, although it is acknowledged that future work should examine the Petzold phase function for sampling and/or interpolation artifacts.

For a fixed search radius (1.0m and 3.5m), Figures 9.10 and 9.11 plot increasing SNR as a function of increasing mean photon density (see associated imagery in Figures 9.37 - 9.48). The rate of increase in SNR with increasing search radius and mean photon density is always greater for the HG ($g = 0.92$) phase function as compared to the Petzold phase function. Again, when the Petzold phase function is employed, the SNR varies little across the water IOPs examined here. Focusing on the HG scattering phase function, a comparison of Figure 9.10 with Figure 9.11, shows how the target's reflectance can influence the SNR. With a highly reflective target, the water-type (coefficients) influence the SNR, but for dark targets variations in SNR become increasingly impervious to variations in water coefficients at increasing geometric depths, even converging for different water types at fixed geometric depths (*e.g.* the SNR curves for targets at 3.5m converge for Groups #1 and #4b (coastal and turbid water, respectively) and at 5.0m for Groups #1 and #3 (coastal and bay water, respectively)).

Depth ↓

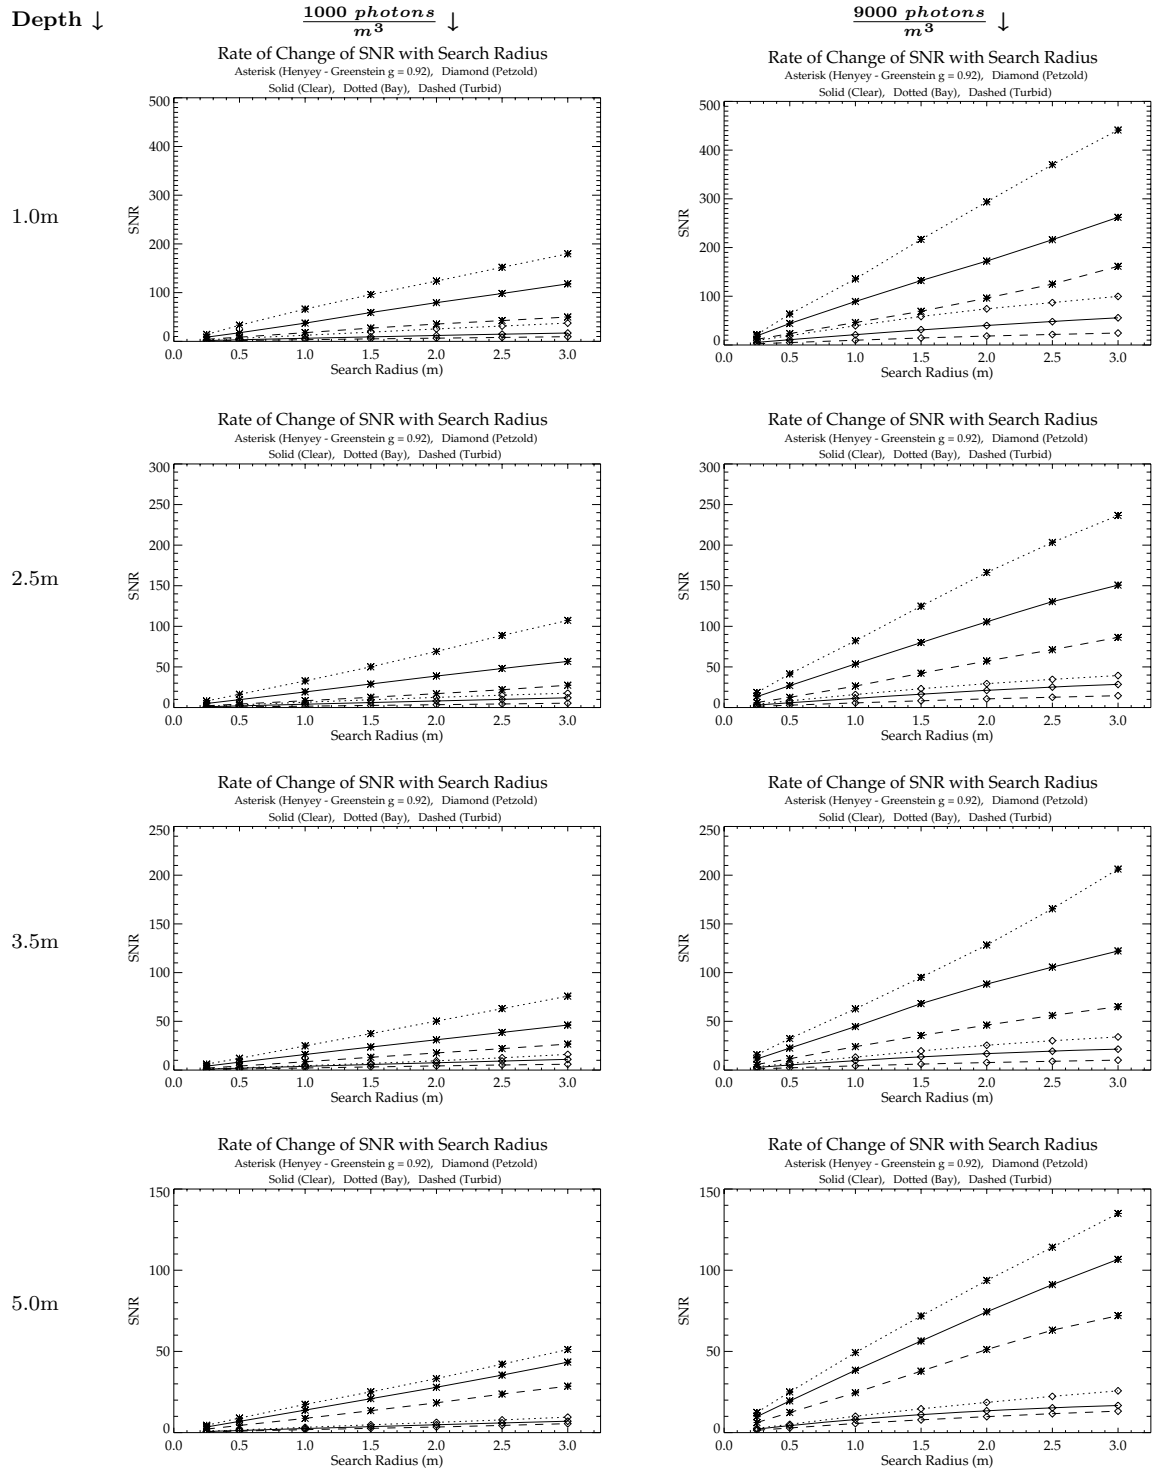


Figure 9.8: SNR of bright targets versus search radius for 2 mean photon densities and 4 target depths.

Depth ↓

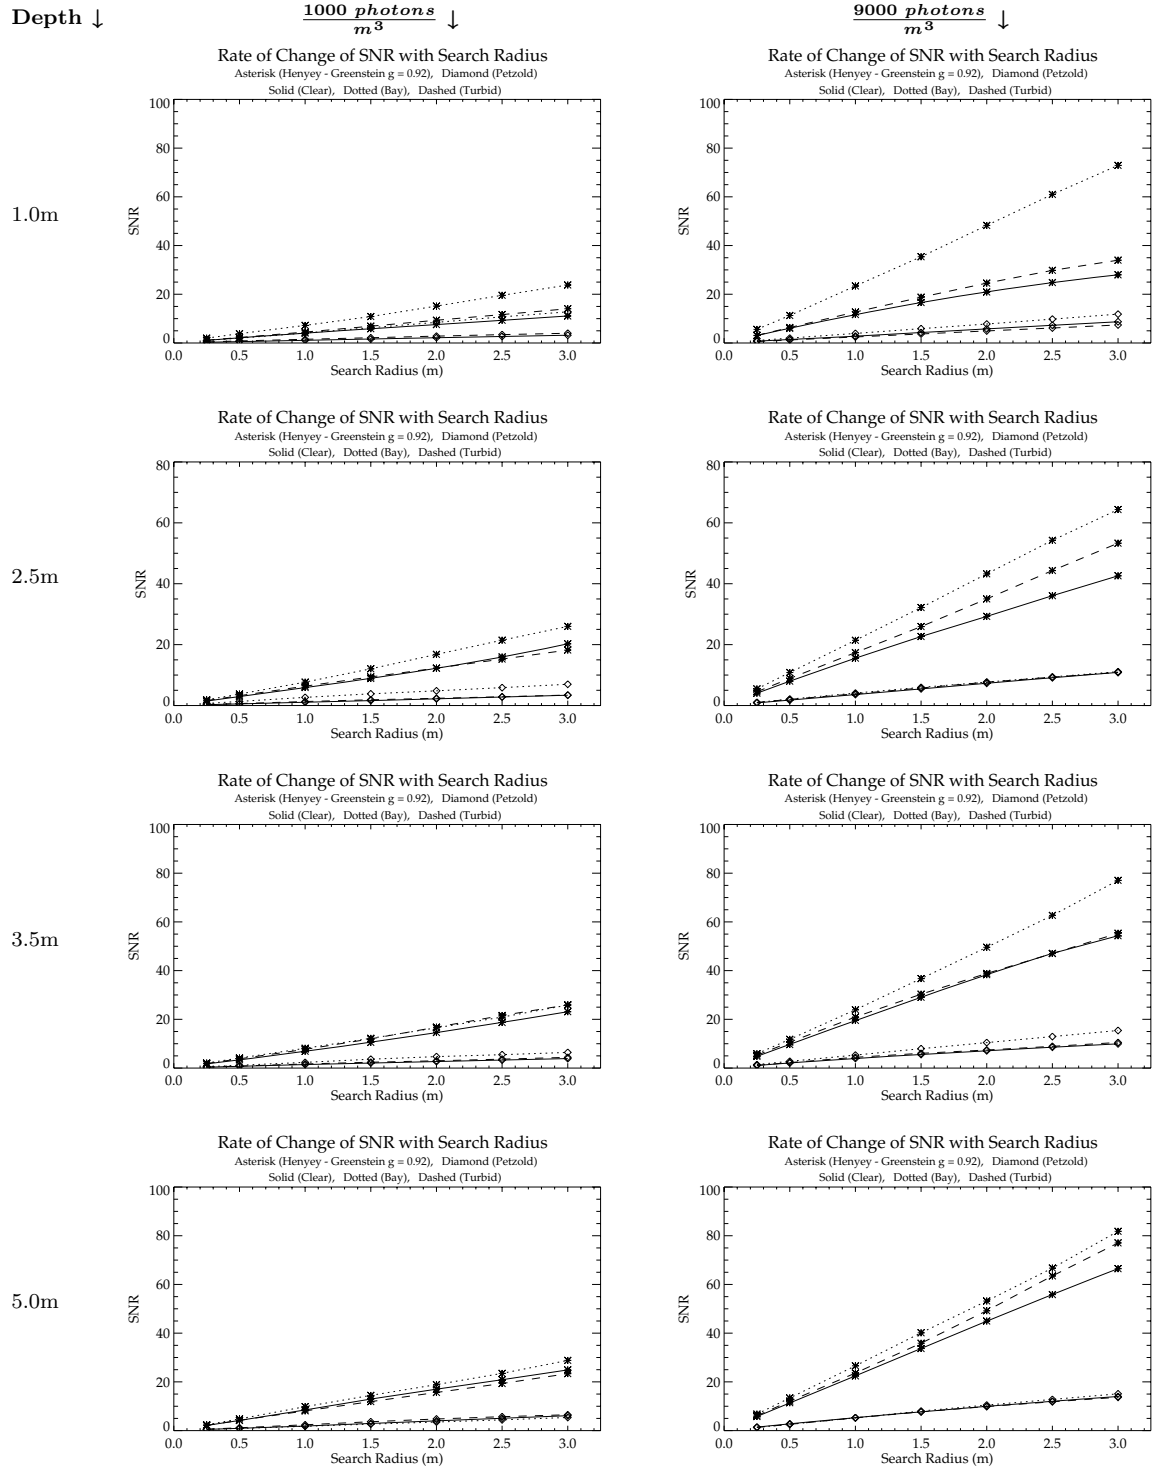




Figure 9.9: SNR of dark targets versus search radius for 2 mean photon densities and 4 target depths.

Depth ↓



Search Radius = 1.0m ↓

Search Radius = 3.0m ↓



1.0m



2.5m



3.5m



5.0m







Figure 9.10: SNR of bright target versus mean photon density for 2 search radii and at 4 target depths.

Depth ↓



Search Radius = 1.0m ↓

Search Radius = 3.0m ↓



1.0m



2.5m



3.5m



5.0m





Figure 9.11: SNR of dark target versus mean photon density for 2 search radii and at 4 target depths.

9.11.1 Summary of Factors that Influence SNR

Overall, it appears that target reflectance (and/or optical depth since both influence the magnitude of the direct reflected radiance), water IOPs, map search radius, and mean photon density all interact to produce a given SNR. Admittedly, some of the aforementioned variables are correlated, and it is further acknowledged that additional variables (not yet considered) may in fact contribute to the observed SNR. However, several broad conclusions can be drawn:

1. The SNR increases as the search radius increases. The rate of change is not *explicitly* related to water clarity since bay water shows the greatest rate of increase (*e.g.* greater than the rate of change for the coastal water scenario, although both the bay and coastal water scenarios have nearly equivalent total extinction coefficients). Although the turbid water scenario does show the least rate of increase, it is unclear exactly which IOP variable this rate is most correlated with, although it appears to be inversely related to the optical depth, and for a fixed optical depth, inversely related to the single scattering albedo.
2. The SNR increases as the mean photon density increases, however, tripling the mean photon density but keeping a small search radius provides very little gain in SNR. Again, the highest rate of increase is associated with the bay water scenario.
3. The SNR increases as the average cosine decreases (becomes less forward peaked).
4. For dark objects, (and in general, for bright objects), the highest mean and lowest variance are associated with higher backscatter fractions.
5. The highest SNR is associated with the most shallow geometric depth (1m) for highly reflective targets, but the greatest geometric depth (5m) for highly absorptive target.
6. The highest SNRs are associated with the bay water IOPs when coupled with greater backscatter returns, but the *characteristics* of the phase function are also believed to play a role.
7. When the sensor reaching radiance is largely a function of backscatter, a target at a fixed (but high) geometrical depth can generate SNRs that are increasingly impervious to the water coefficients (note the convergence of SNR for different water types within Figure 9.11).
8. If a water type is coupled with the Petzold phase function, then the SNR appears to be dominated by the characteristics of the SPF, and less so by water coefficients. The validity of this should be examined in light of the perceived artifacts associated with this phase function.

9. Unless the absorptive and reflective targets have equal SNR, the SNR for the absorptive target is always lower than that for the reflective target. However, if the SNR of the bright and dark target are nearly equal, then the water coefficients (*e.g.* turbidity or optical depth) are great enough that image transfer is lost. In this case, it is unlikely that the simulation will reveal much of the underlying structure reflected from the benthic plane, and therefore the user should consider if the simulation actually warrants the 3D benefits associated with photon mapping.

Despite the academically interesting trends that have been loosely identified, the user is still left wondering whether an optimal SNR exists. From Figures 9.25 through 9.48, it seems that an SNR of approximately 40 to 50 provides reasonable and acceptable results for bright targets, while SNRs of 70 to 100+ are superior. The bigger tradeoff is in the spatial correlation of the noise. Smaller search radii (lower SNR) generate imagery that looks noisy, but larger search radii, although associated with a higher SNR, generate imagery that has a much more *mottled* appearance. The mottling is associated with the fact that the solution is expected to converge to the correct result, but convolved with the kernel function. For larger search radii, the kernel function is a larger circular search area, which at times, may be visually apparent. Subjectively, the less mottled appearance is actually more pleasing to the author, but it is acknowledged that others may disagree.

In addition to the noted trends, and perhaps of equal or greater utility, the results of this section can be used to help drive the user-settings of future simulations. A user can find the IOP configuration that most closely matches his or her simulation, and based on the qualitative and quantitative results that have been provided here, determine first-order estimates for mean photon density and search radius based on their simulation goals. Alternatively, this section provides proof of concept for a *virtual* step function, which if integrated within DIRSIG, can facilitate the automatic selection of these user-set variables.

9.12 Results: PSFs

The composite ESF was obtained from the slanted step function using the method described by Schott (2007). Regardless of the search radius employed, the composite was somewhat noisy, and certainly too noisy to directly convert to a PSF. The results were therefore smoothed using a box car average, and the smoothed ESF was then fitted using a sum of three Fermi functions as described by Tzannes and Mooney (1995). Figure 9.12 depicts the averaged composite ESFs and PSFs for Group #1 (coastal) water IOPs with HG ($g = 0.92$), modeled using a mean volume photon density of $9000 \frac{\text{photons}}{\text{m}^3}$ and search radius of 0.25 (solid line), 1.0 (dotted line), 2.0 (dashed line) and 3.0 (dash-dotted line). The top row depicts results associated with a target at 1.0m from the air-water interface. The remaining curves depict results from targets at 2.5, 3.5 and 5.0m from the air-water interface. It is evident from the ESFs that increasing the search radius is associated with a degree of edge bias, but not a direct loss in the actual resolution of the edge. This is confirmed when examining the PSFs, since larger search radii generate higher PSF tails (adjacent to the central peak) and yet the FWHM is nearly constant (especially after noise considerations).

The nearly-constant FWHM indicates that the resolution of the actual edge is the same, regardless of the search radius. This is likely a product of three things. First, the sensor reaching radiance is dominated by direct radiance, especially at shallow optical depths. Bias in the direct radiance is a function of the surface photon map. Although the search radius for the surface photon map has been set equal to that of the volume (*e.g.* either 0.25, 1.0 2.0 or 3.0m), this Radiometry Solver uses an adaptive sampling algorithm which decreases edge bias. In essence, the algorithm looks for rapid fluctuations in radiance (*e.g.* edges), and if found, the algorithm will average along the edge, rather than across it. This helps maintain high edge resolution, regardless of how the user configures the nominal surface map search radius. The second contributing factor is the sensor GSD (0.25m). If the edge does exhibit a loss in resolution with increasing search radius, it is below the limit of resolution of the sensor. Finally, small angle scattering events have been recast as direct radiance, which is expected to decrease the central peak of the PSF.

Figures 9.13 and 9.14 show corresponding ESFs and PSFs for Group #3 (bay) and #4b (turbid) water IOPs, respectively. The same pattern persists; the actual edge is nearly equally well resolved at all search radii, although biased on either side, generating higher tails for the PSF. Note that the highest search radius is 12 times the sensor's GSD, which results in a search *area* with a diameter of 6.0m (for a scene that is 62.5 x 50.0m in total size.)

Depth ↓

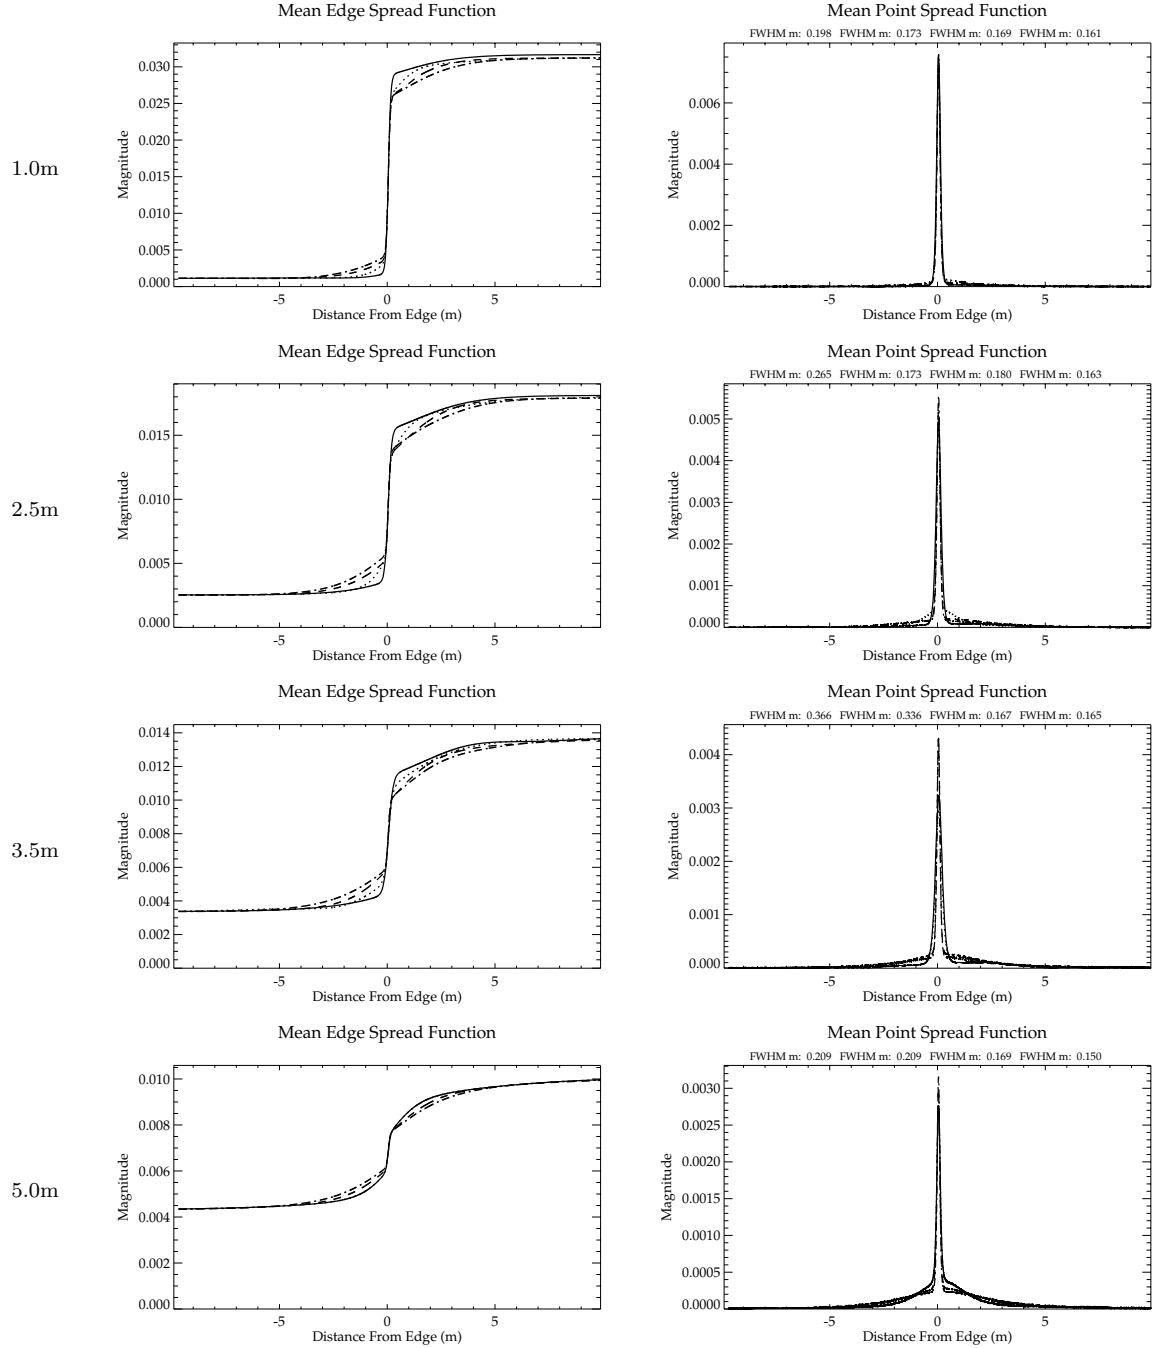


Figure 9.12: ESFs (left) and PSFs (right) for targets at 1.0, 2.5, 3.5 and 5.0m from the air-water interface (top to bottom) simulated using (surface and volume) search radii of 0.25 (solid line), 1.0 (dotted line), 2.0 (dashed line) and 3.0m (dash-dotted line). The FWHM is shown on the plot for search radii of 0.25, 1.0, 2.0 and 3.0m (from left to right) for Group #1 (coastal) IOPs, HG ($g = 0.9231$)

Depth ↓

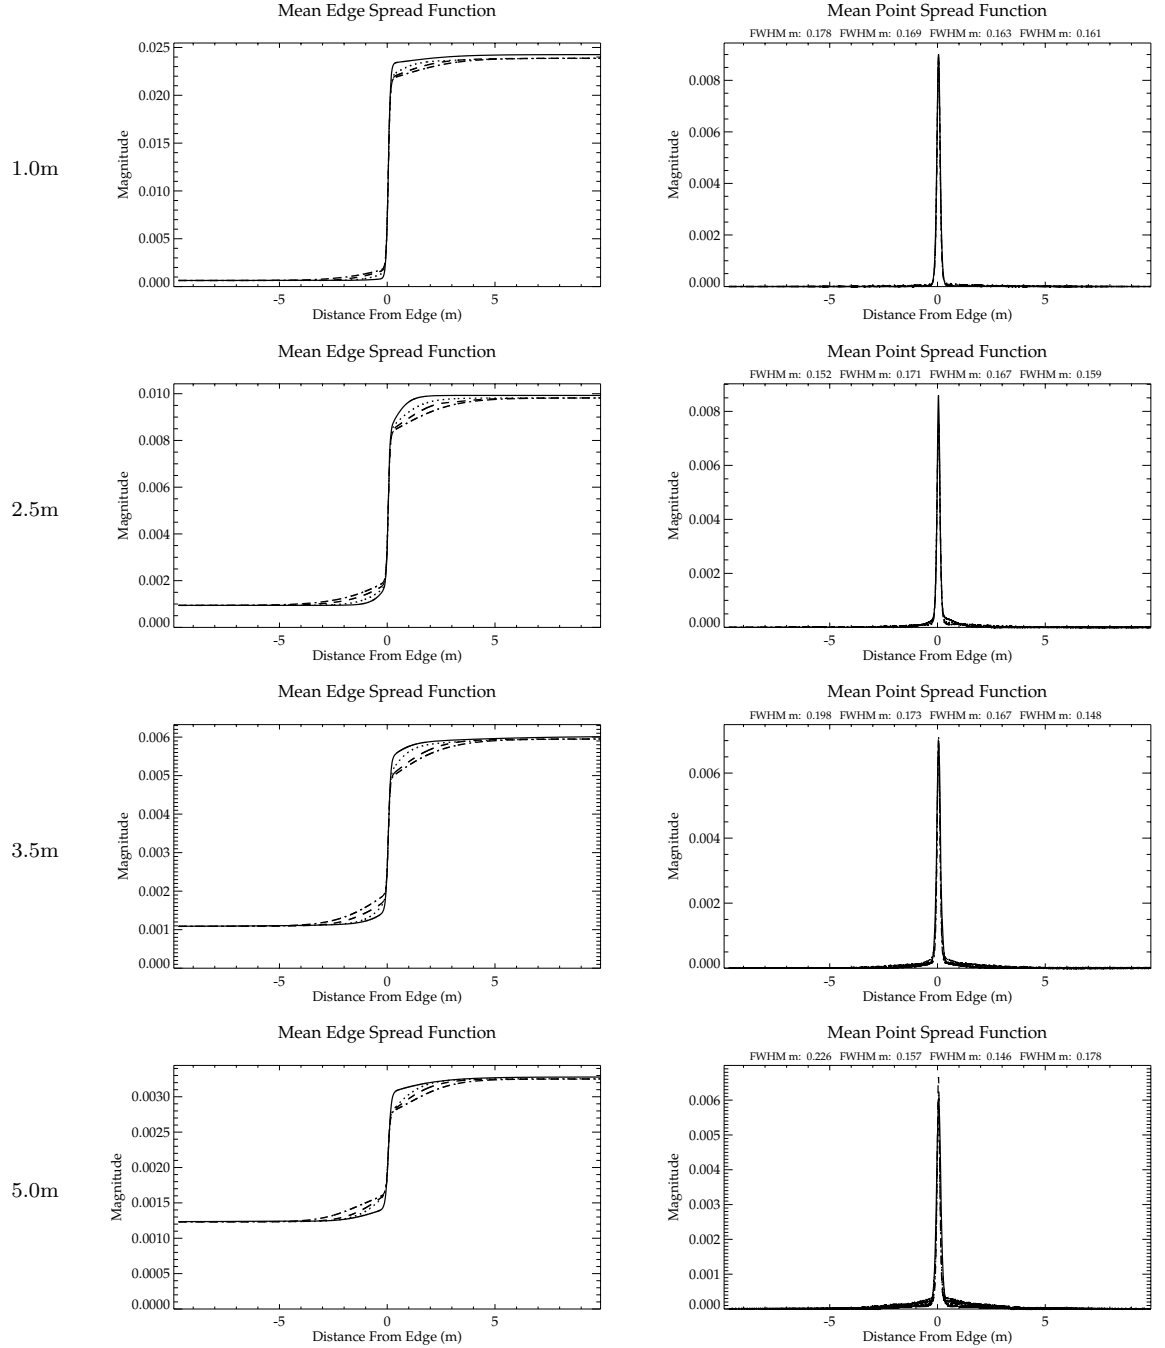


Figure 9.13: ESFs (left) and PSFs (right) for targets at 1.0, 2.5, 3.5 and 5.0m from the air-water interface (top to bottom) simulated using (surface and volume) search radii of 0.25 (solid line), 1.0 (dotted line), 2.0 (dashed line) and 3.0m (dash-dotted line). The FWHM is shown on the plot for search radii of 0.25, 1.0, 2.0 and 3.0m (from left to right) for Group #3 (bay) IOPs, HG ($g = 0.92$). 232

Depth ↓

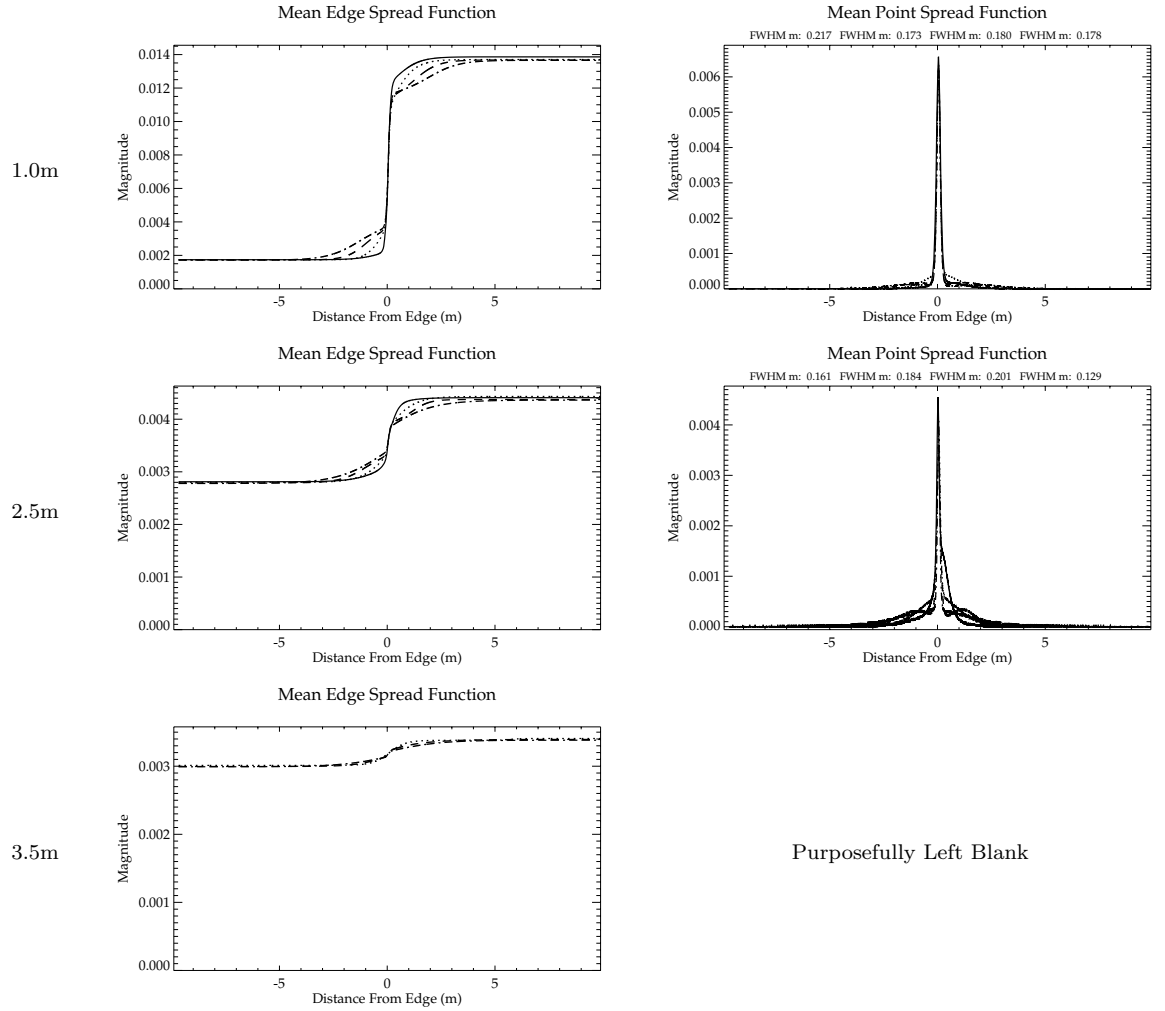


Figure 9.14: ESFs (left) and PSFs (right) for targets at 1.0, 2.5, 3.5 and 5.0m from the air-water interface (top to bottom) simulated using (surface and volume) search radii of 0.25 (solid line), 1.0 (dotted line), 2.0 (dashed line) and 3.0m (dash-dotted line). The FWHM is shown on the plot for search radii of 0.25, 1.0, 2.0 and 3.0m (from left to right) for Group #4b (turbid) IOPs, HG ($g = 0.92$).

9.12.1 Summary of PSF Results

The FWHM of the PSFs shown in Figures 9.12, 9.13 and 9.14 implies that the central peak of the water scattering PSF is extremely narrow over several optical depths. Although there is slight variation in the FWHM with search radius, it is asserted that none of the reported variability is above noise and/or methodological uncertainty in the ESF approximation.

The consistent central peak within the DIRSIG simulated results are believed to be a function of one simulation property, and two model properties. First, the GSD is 0.25m which means the simulated results cannot capture variability that is less than 25cm. Second, the adaptive sampling method employed within the SurfacePM search radius appears to be well-suited toward resolving the actual edge (for moderate IOPs), and keeping any changes in the central peak of the PSF below the GSD. Finally, the effective scattering coefficient (see Section 6.4.7) serves to reallocate small angle scattering events (within 0.01 steradians of the propagating ray) to direct radiance. Although this is an extremely small solid angle, it certainly decreases the magnitude of the scattered radiance, and reduces the number of single scattering events that are stored in the photon map, both of which will keep the central peak of the PSF narrow (without changing the absolute computed radiance).

This is not to suggest that increasing search radius is without consequence. Rather, the consequence of increasing SR first manifests itself in the tails of the PSF distribution, rather than the central peak. However, for a GSD of 0.25m, increasing the search radius to 3.0m (and therefore searching in an area of nearly 29 m^2) still confines much of the change in the PSF to the tails. Plus, the change is most likely the result of residual bias in the surface photon mapped solution, rather than the actual volume (*e.g.* the adaptive boundary compensation in the surface solution cannot divide the user-set nominal search area into an infinite number of subareas in order to test for rapid changes in the distribution, therefore some bias will persist, and this should be correlated with the size of the nominal search radius). Of course, a simple metric like FWHM does not discern the variation in the tails, but perhaps of greater importance is that if a user can accept the consequences of increasing SR, then a great deal of SNR can be gained while still providing reasonable resolution of the actual edge. *In fact, since it is likely that much of the bias observed in the ESF is actually residual bias that persists due to less than perfect boundary compensation in the surface photon mapped solution (rather than the actual volume solution), the volume and surface search radii can be varied in an arbitrary manner (instead of each set nominally equal to each other), to obtain various combinations of desired SNR and bias-reduction.*

The next question is how much the air-water interface will interfere with the PSF resolution, and if refractive events at the air-water interface will introduce greater variability than the search radius. To test this, the variability in the ESF associated with a wind roughened air-water interface was examined. Using the Beaufort Scale of wind speed ver-

sus probable wave height, two facetized surfaces were generated, as shown in Table 9.5, making use of IDL code supplied by Gartley (2010) (with slight modifications). Every effort was made to generate a reasonable mesh surface with sufficient facets, an appropriate RMS wave height, and a reasonable geometry file size and generation run time.

Table 9.5: Beaufort scale and attributes of facetized wave surfaces. The first two columns describe the wind and associated waves based on the Beaufort Scale. The last three columns report the predicted and actual RMS wave height, wave slope, and correlation length for the simulated surface.

Wind Speed $\frac{m}{s}$ (mph)	Wave Description	RMS Wave Height (m) (predicted, actual)	RMS Wave Slope (radians) (predicted, actual)	Correlation length (m)
6.0 (13.4)	Small Waves	0.576, 0.572	0.0337, 0.0335	24.2
10.0 (22.4)	Moderate Waves	1.600, 1.643	0.0542, 0.0557	41.7

The composite ESFs generated using the facetized wave surface are shown in Figure 9.15. They were generated using the same source and extinction attributes associated with Group #3 water IOPs (HG, $g = 0.92$), but void of any photon mapped solution, and using GRS on the edge. The results have been overlaid on the ESF curves associated with the Group #3 water IOPs, using search radii of 1.0 and 3.0m (and for a target-surface depth of 2.5m). Since the impact of the refractive interface on edge bias is proportional to the depth of the target, the magnitude of the effects shown in Figure 9.15 will be greater for targets at greater depth, and less for targets at shallower depths. However, the methodology associated with generating the composite ESF decreases the noise in the edge, and in so doing, minimizes the localized bias as well (*e.g.* refractive-type bias), so the results presented in Figure 9.15 are definitely indicative of a “best-case” scenario.

Figure 9.15 suggests that the capillary wave induced variations are just below the sensor GSD. This is confirmed by inspection of ESFs and PSFs associated from a field collect conducted on October 12th 2009, and shown in Figure 9.16. The wind speed was approximately 10 mph, and several images were collected while the test target was fixed at 10 inches below the air-water interface. At this shallow depth the actual ESF will be largely free of water-induced PSF effects. In addition, it will be largely free of refractive-depth induced variability. This means that targets viewed at this depth will have the least refractive induced edge variability for a constant wind speed and direction, and that this variability is expected to increase as the target’s depth is likewise increased. Since the field collect configuration allowed for very high resolution (GSDs on the order of centimeters), Figure 9.16 shows that the interface, (at this shallow depth) can contribute a PSF that drops to nearly zero beyond a central width of approximately 20cm. Of course, this is below the resolution of most remote sensing applications, and is in agreement with the




Figure 9.15: ESFs (left) and PSFs (right) as a function of capillary waves (for 6 (solid line) and 10 (dotted line) $\frac{m}{s}$ wind speeds), as compared to the results induced by 1m and 3m search radii for Group #3 water IOPs (HG, $g = 0.92$), and at 2.5m between surface and target. The FWHM is shown on the plot for a wind speed of 6 and 10 $\frac{m}{s}$, and search radii of 1.0 and 3.0m (from left to right).

results shown in Figure 9.15, although targets at deeper depths will surely show more wave-induced degradation.

Although the impact of refractive-induced ESF bias has only been briefly touched upon here, it is reasonable to assert that the bias contribution of this interface will be relatively small when the following is applicable (barring any consideration of breaking waves):

- The presence of low to moderate wind speeds within which the resolution of subsurface target edges is a reasonable pursuit.
- When targets of interest are positioned at relatively shallow depths.
- When analytical ESFs are generated using composite transects of an actual edge, since individual small wavelets can generate pronounced (but localized) edge bias (*Translation: objects with sufficiently large size such that localized edge bias does not disrupt imaging goals.*)
- When reasonable remote sensing GSDs are expected (*e.g.* on the order of approximately 0.25m or greater).

Although the competition and interplay of search radius and refractive interface bias has not been exhaustively evaluated here, the user should keep in mind that the latter is expected to impact the entire PSF, including the central lobe, rather than being restricted to the tails of the PSF as was found with increasing search radii in the synthetic data.




Figure 9.16: Top: ESFs (left) and PSFs (right) as a function of random capillary waves associated with a wind speed of 10 mph and for a target at 10 inches below the air-water interface. Since this Figure is meant to illustrate random variability in the PSF associated with capillary waves, a one-to-one correlation between FWHM and line type is not provided. Bottom: Sample imagery from field collect.

9.13 Spectral Considerations

9.13.1 Results: Spectral SNR

A spectral image quality investigation was carried out using IOPs obtained from a field collect dated October 12th 2009. The original IOPs are shown below in Table 9.6. These IOPs represent a very turbid water with maximum visibility of about 1m. The IOPs were therefore reduced to 0.25 of their original value to render a less turbid environment, but still consistent with the spectral character of the field samples.

Figures 9.49 through 9.52 in Section 9.16 depict the ESFs, sampled at spectral points of 0.4, 0.5 and 0.6 μm , using the IOPs shown in Table 9.6, and the HG ($g = 0.92$) phase function. Each image compares the SNR, per spectral band, for a multi-band simulation (labeled “3 Band Simulation”), with that of 3-single band simulations that were generated sequentially, and then combined (labeled “3 Band Combination”). For the 3 band sim-

Table 9.6: IOPs based on field collect, and utilizing specific coefficients from generalized Lake Ontario data (Raqueno 1999). All IOPs were reduced to 0.25 of their original values since the original coefficients represent extremely turbid water with maximum visibility of about 1m.

Coefficient $\frac{1}{m}$	0.4 microns	0.5 microns	0.6 microns
Chlorophyll σ_a	0.18623	0.146969	0.0401584
Chlorophyll σ_s	0.518596	0.565279	0.541272
TSS σ_a	1.147829	0.396356	0.152813
TSS σ_s	7.9333	6.34667	5.288889
CDOM σ_a	0.7222	0.1472	0.0184

ulation, the bands have a combined mean volume photon density of $3375 \frac{\text{photons}}{m^3}$, and a combined mean surface photon density of $4500 \frac{\text{photons}}{m^2}$. Conversely, the 3-single band simulations were generated such that each spectral band has the same reported mean photon density. Figure 9.17 visually illustrates the simulation combinations.

Figure 9.49 compares the 3 band simulation to the 3-single band simulations (combined after simulation) with mean photon densities of $1125 \frac{\text{photons}}{m^3}$ and $1500 \frac{\text{photons}}{m^2}$, and using search radii of 0.25, 1.0, 2.0 and 3.0m (and target depths of 1.0, 2.5 and 3.5m). This comparison attempts to discern how variation in the spectral coefficients impact the perceived results. Based on visual inspection, the 3 band simulation has a slightly better SNR over the single band combinations. Directly related to SNR, the single band simulations appear more “green” with increasing depth, while the multi-band simulation has a more “blue” appearance. This seems to indicate a variation in the mean signal (numerator of SNR computation) per band with simulation configuration. Figure 9.50 compares the 3 band simulation to 3-single band simulations (combined after simulation) with mean photon densities of $1125 \frac{\text{photons}}{m^3}$ and $4500 \frac{\text{photons}}{m^2}$. This comparison attempts to discern how the mean photon density of the surface map influences the results. However, little discernible difference is visually apparent. This is not to suggest that the photon density of the surface map is negligible, but instead shows that once a reasonable map density has been obtained for a given simulation configuration, there is little to be gained by adding more photons to the surface map. Figure 9.51 compares the 3 band simulation to 3-single band simulations (combined after simulation) with mean photon densities of $3375 \frac{\text{photons}}{m^3}$ and $1500 \frac{\text{photons}}{m^2}$. This comparison attempts to discern how the volume map influences the results. Visually, the SNR appears to be higher than that for the $1125 \frac{\text{photons}}{m^3}$ and $1500 \frac{\text{photons}}{m^2}$ combination shown in Figure 9.49.

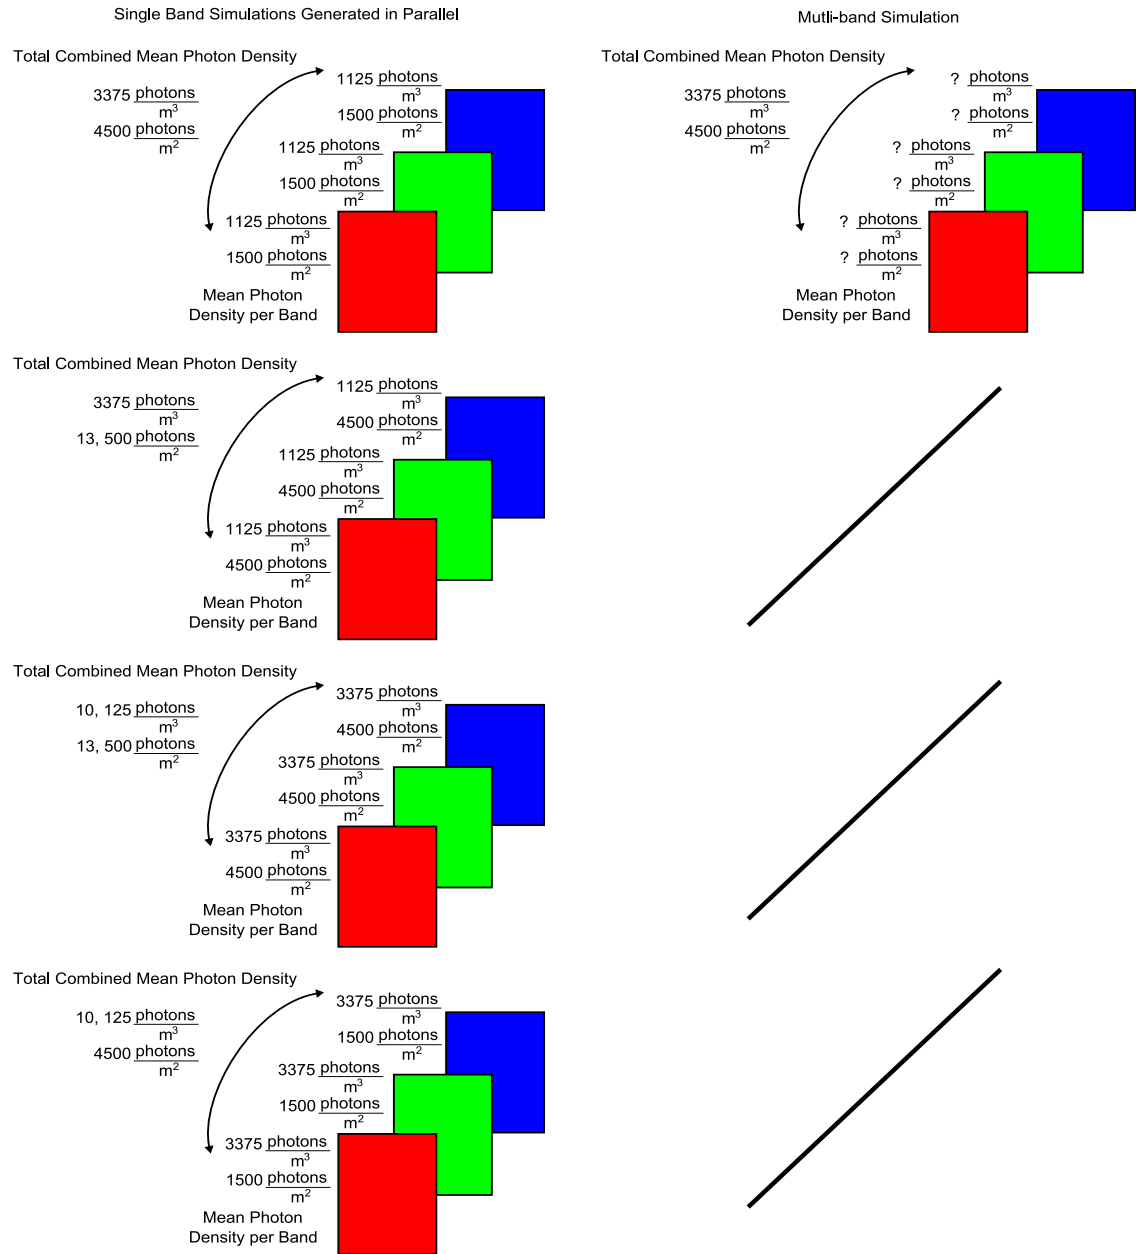


Figure 9.17: Schematic to explain the spectral simulations. The right column shows the series of single band combinations, including the volume and surface photon density per band, and the photon density of the combined 3-band image. The left column depicts the 3-band simulation that was generated with fixed volume and surface total photon densities, allowing the water IOPs to dictate the actual density per spectral band.

Finally, Figure 9.52 compares the 3 band simulation to 3-single band simulations (combined after simulation) with mean photon densities of $3375 \frac{\text{photons}}{\text{m}^3}$, and $4500 \frac{\text{photons}}{\text{m}^2}$. This means that each spectral band has a mean photon density that is equal to the total mean photon density in the 3-band simulation (or that the 3-single band simulations, when combined, have 3X the number of photons as the 3 band simulation). Visually, this simulation configuration generates an SNR that is nearly equivalent to the 3 band simulation, and which is in agreement with the cumulative results shown in Figures 9.49, 9.50 and 9.51.

The computed SNR for each band in Figures 9.49 through 9.52 is plotted in Figure 9.18. The x-axis shows the spectral points of interest at 0.4, 0.5 and $0.6\mu\text{m}$. The y-axis plots the SNR per spectral band, as a function of the search radius (white = 0.25, light gray = 1.0, dark gray = 2.0 and black = 3.0m), and for each mean photon density that was examined. The “order” of the plots is shown by labels above the first ($0.4\mu\text{m}$) spectral point (*e.g.* $[1125 \frac{\text{photons}}{\text{m}^3} + 1500 \frac{\text{photons}}{\text{m}^2}]$, $[1125 \frac{\text{photons}}{\text{m}^3} + 4500 \frac{\text{photons}}{\text{m}^2}]$, $[3375 \frac{\text{photons}}{\text{m}^3} + 4500 \frac{\text{photons}}{\text{m}^2}]$, $[3375 \frac{\text{photons}}{\text{m}^3} + 1125 \frac{\text{photons}}{\text{m}^2}]$, and the last (designated with an asterisk) as the 3 band simulation with $[3375 \frac{\text{photons}}{\text{m}^3} + 4500 \frac{\text{photons}}{\text{m}^2}]$). The left column shows the results for the bright target, while the right column shows the results for the dark target (at depths of 1.0, 2.5 and 3.5m from top to bottom).

The mean and variance of the signal for the bright and dark targets (whose SNR is shown using black bars in Figure 9.18 *e.g.* search radius = 3.0m) can be found in Figures 9.19 and 9.20. For both the bright and dark targets, the single band simulations have nearly equivalent mean radiance signals. In comparison to the single band simulations, the multi-band simulation has a slightly lower mean signal for the blue and green spectral bands, but a slightly higher signal for the red spectral band. This helps explain the above observations concerning the slight “color” variation between the simulation configurations (*e.g.* the single band simulation (for all mean photon densities) is more “green” than the multi-band counterpart).

For the dark targets, the multi-band simulation always has a slightly higher variance. The lower mean signal and higher variance for the blue and red bands in this simulation configuration explains the lower SNR for the dark targets in Figure 9.18 (see the cluster of bars denoted with an asterisk). Although the red band has a higher mean signal, the competition of the higher mean signal and variance still leads to a lower SNR for this spectral channel.

As for the bright target, the variance of the signal for the multi-band simulation is sometimes greater, and sometimes less than its single band counterpart. This leads to an SNR that is sometimes greater, and sometimes less than the equivalent single-band simulations as shown in Figure 9.18.

Depth ↓

Bright Target

Dark Target

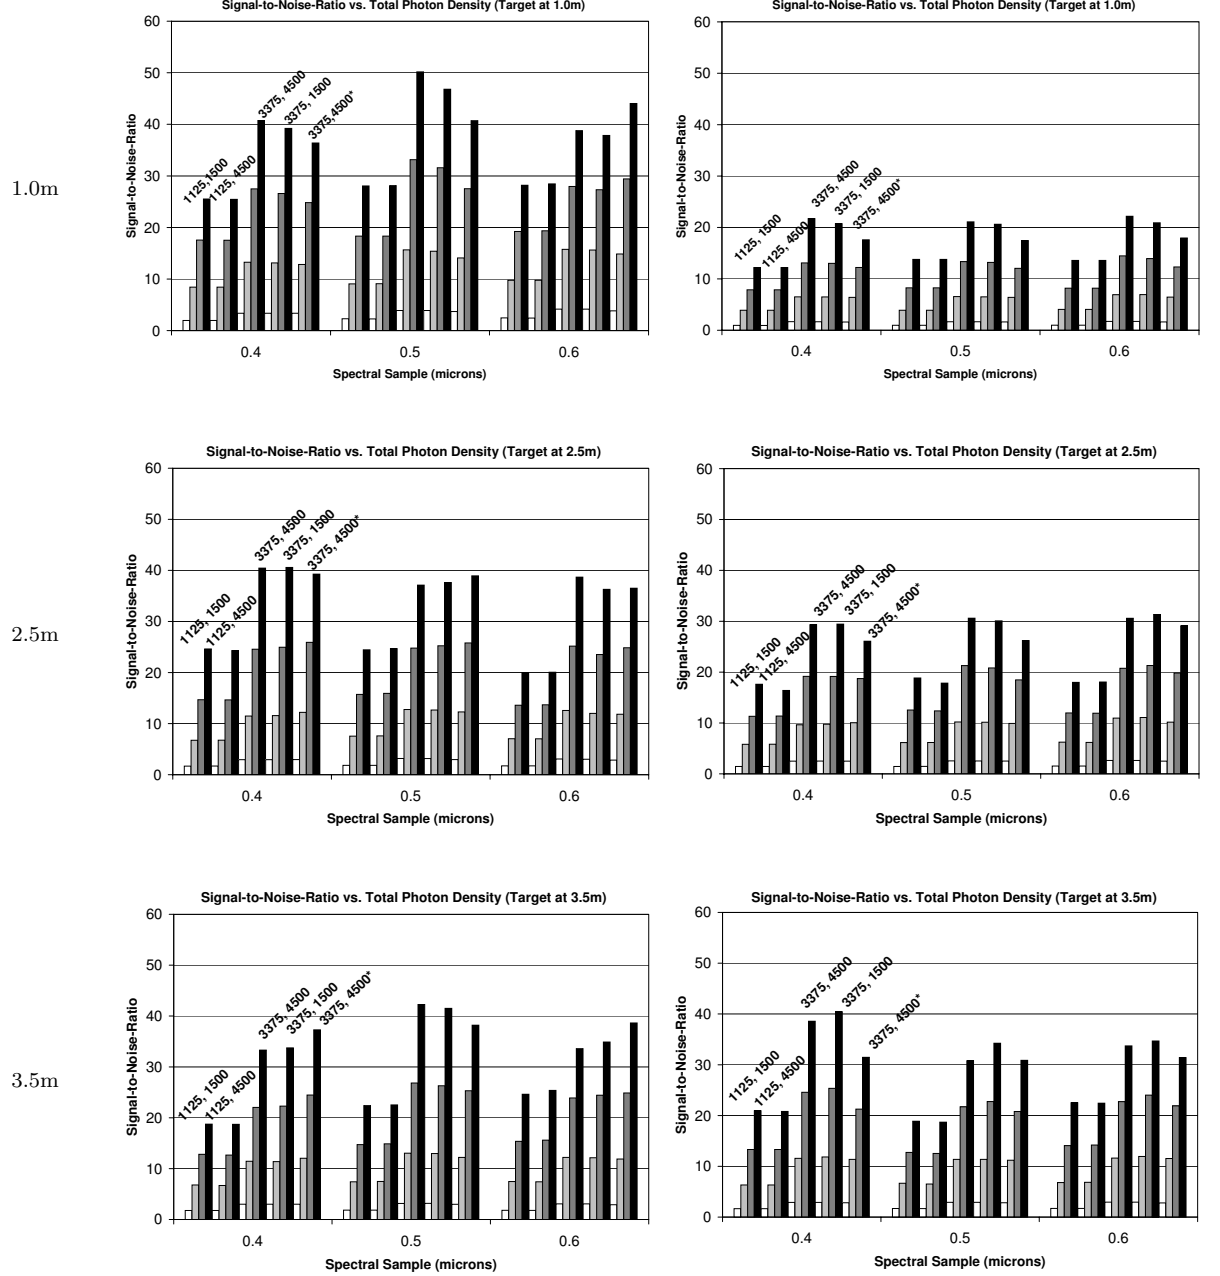


Figure 9.18: SNR at 3 target depths (from surface). The bar color corresponds to search radius as follows: white (0.25), light gray (1.0), dark gray (2.0) and black (3.0m). The plot labels report mean volume and surface photon density, respectively, in $\frac{\text{photons}}{\text{m}^3}$ and $\frac{\text{photons}}{\text{m}^2}$ per spectral band, and the total density for the 3-band simulation (designated by an asterisk).

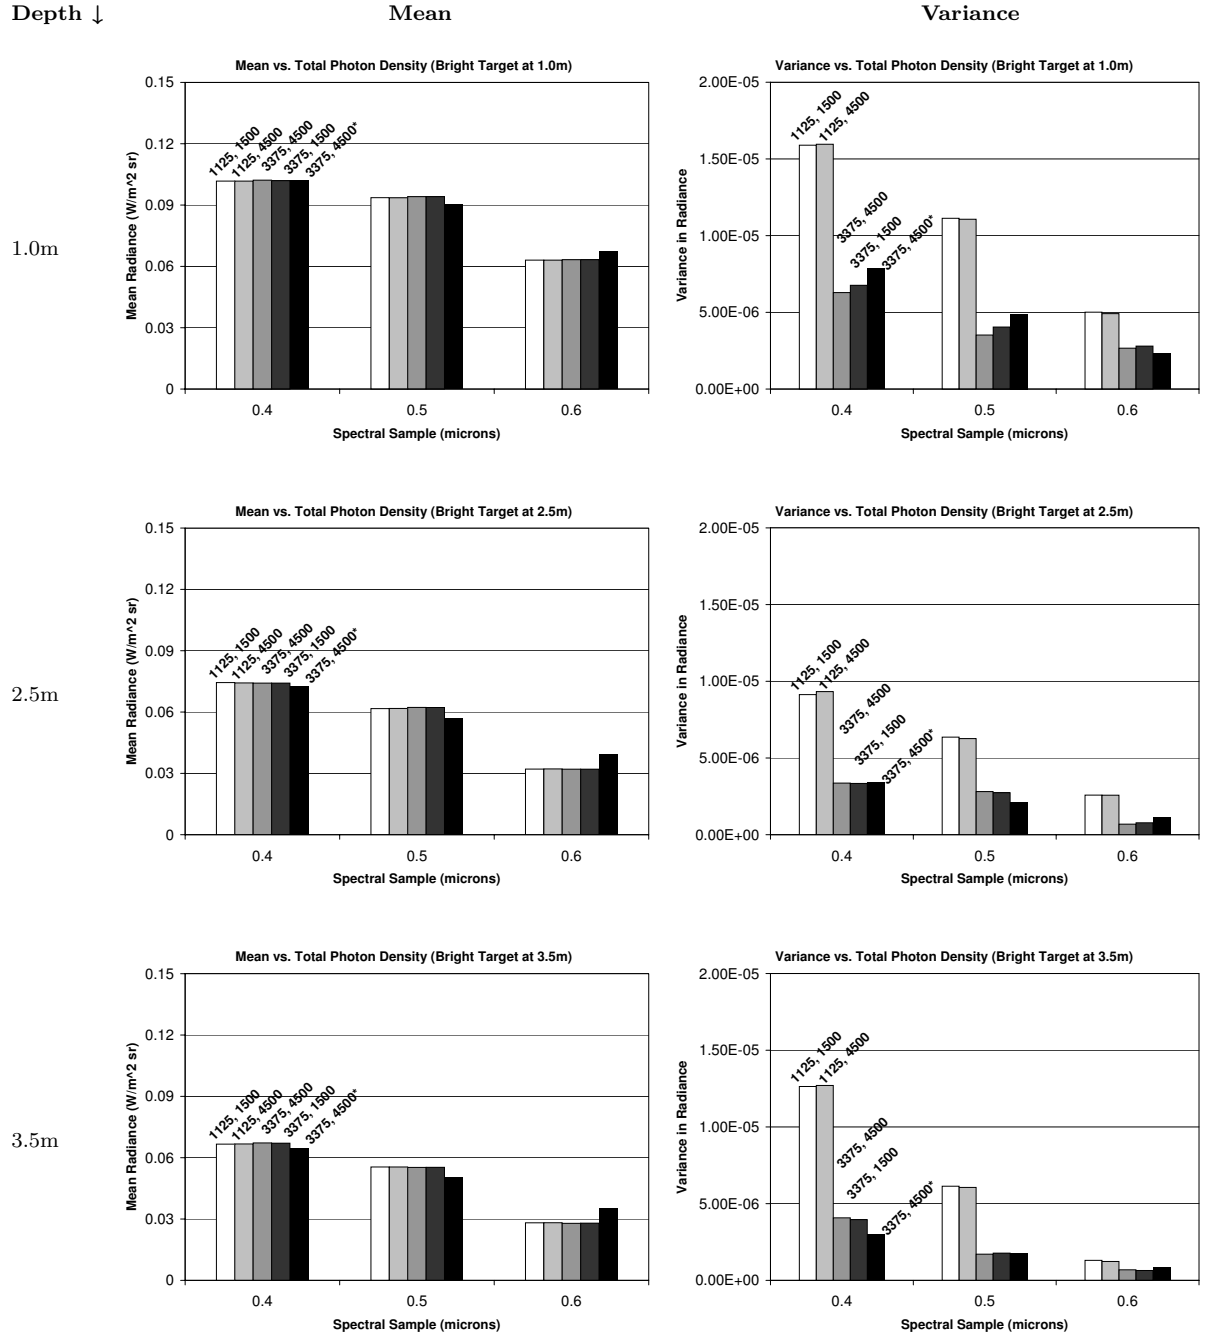


Figure 9.19: Mean and variance of bright targets at 3 depths (from surface) and for a fixed search radius of 3.0m. The bar color corresponds to the total mean photon density as labeled above spectral point = $0.4\mu m$ (mean volume and surface photon density, respectively, in $\frac{photons}{m^3}$ and $\frac{photons}{m^2}$ per spectral band, and the total density for the 3-band simulation [designated by an asterisk]).

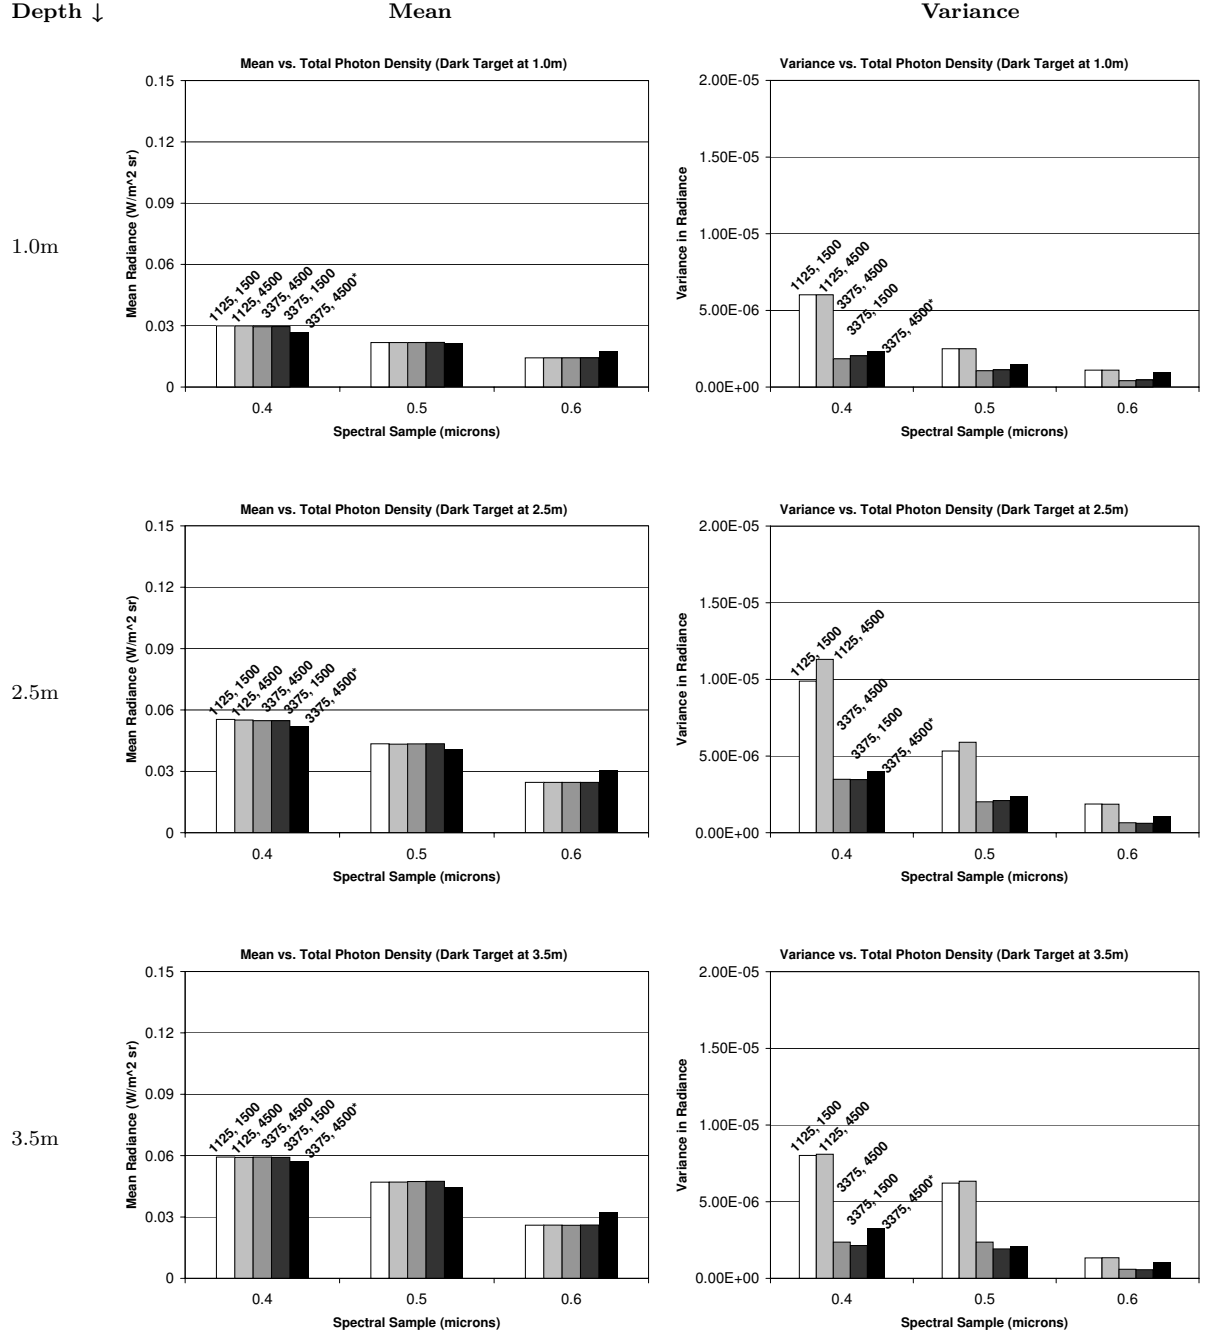


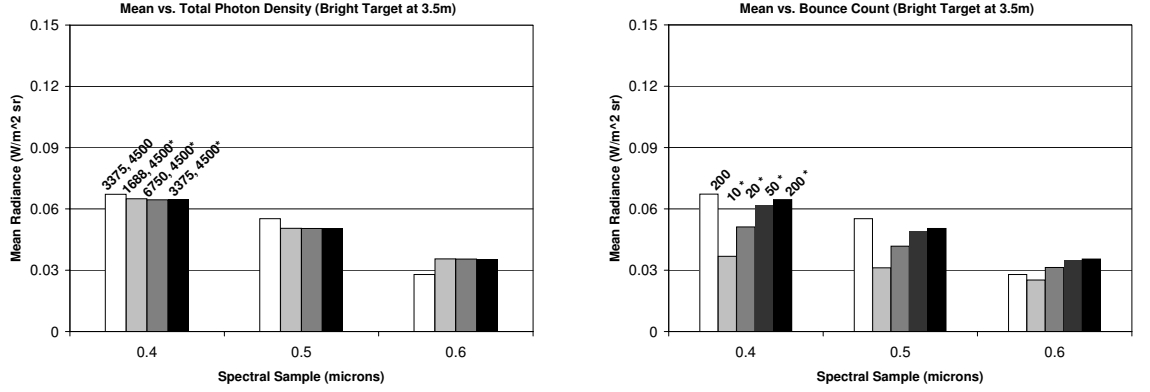
Figure 9.20: Mean and variance of dark targets at 3 depths (from surface) and for a fixed search radius of 3.0m. The bar color corresponds to the total mean photon density as labeled above spectral point = $0.4\mu\text{m}$ (mean volume and surface photon density, respectively, in $\frac{\text{photons}}{\text{m}^3}$ and $\frac{\text{photons}}{\text{m}^2}$ per spectral band, and the total density for the 3-band simulation [differentiated by an asterisks]).

A few general conclusions can be drawn after inspection of the images from Figures 9.49 through 9.52 (in Section 9.16), and summarized in terms of target mean, variance and SNR in Figures 9.18, 9.19 and 9.20. First, for single spectral band simulations, the SNR is more strongly dominated by the mean photon density in the volume as opposed to that on the surface (assuming a reasonable surface photon density to begin with). This is based on the negligible change in SNR that is associated with tripling the photon density on the surface, while keeping the volume density constant. Second, for bright targets at depths of 1.0 and 3.0m the SNR of the $0.5\mu\text{m}$ spectral band is slightly greater than its 0.4 and $0.6\mu\text{m}$ neighbors. For the dark target at greater depths, the $0.4\mu\text{m}$ band has a slightly higher SNR than its 0.5 and $0.6\mu\text{m}$ neighbors. The exact band with the highest SNR will likely vary with water IOPs, but always be shifted toward the shorter wavelengths. Finally, a larger search radius combined with a smaller mean photon density can generate SNR results on par with a higher mean photon density and smaller search radius.

With regard to single versus multi-band simulations, several additional trends are apparent. First, it is clear that DIRSIG is able to reduce the variance in the in-scattered radiance computation by utilizing all samples (across all bands) when computing the per band in-scattered radiance in a multi-band simulation. To illustrate, consider a simple thought experiment with 100 available samples confined to a small volume around a propagating ray of interest. Now assume that the spectral tags associated with those 100 samples result in 10 *red photons*, 40 *green photons* and 50 *blue photons*. One computational option is to use the spectral tag to cull the available photons that can be used when computing the in-scattered radiance per band (*e.g.* 10 samples would be available to compute the in-scattered radiance for the red band, 40 for the green band, and 50 for the blue band). Clearly this would result in greater noise for the red band computation (with fewer samples), as opposed to the neighboring blue and green bands. However, the lack of dispersion within the scattering phase function means that all 100 photons can be repeatedly used *for each band*, provided the spectral distribution (10:40:50 for R:G:B) is used to weight the radiance per band in the final computation. The net effect is variance-reduction, and an SNR that is nearly equal to using 3X the number of photons per spectral band (when those bands are generated sequentially). However, the mean radiance per spectral band is slightly different than the comparable single band simulation, and this difference can vary between a few percent (at shorter wavelengths), to as high as 20% (at longer wavelengths with lower mean signals), as shown in Figures 9.19 and 9.20. *Note that the mean signal is only slightly influenced by the spectral search radius in both the single and multi-band simulations. This means that the trend shown in Figures 9.19 and 9.20 (for a search radius of 3.0m) is still indicative of the other (smaller) search radii that were examined.* The reason for this variation cannot be fully explained. At first, it was postulated to be a photon starvation issue. However, varying the mean photon density in the 3-band simulation does not vary the mean signal. This is shown in Figure 9.21 (a.) which plots the mean photon density (SR = 3.0m) for the reflective target, and using 3 different mean

photon densities for the 3-band simulation. The second postulate was suggested by the DIRSIG development team, and related to the bounce count used in the simulation. All of the spectral simulations shown in this Chapter are based on a bounce count of 200. This means that any given photon can undergo 200 events, before the significance of the photon is considered negligible (e.g. the photon history is truncated after 200 events). As a safety precaution, this value should always be set fairly high. However, Figure 9.21 (b.) shows that increasing bounce count, although associated with a more accurate mean signal, will not ultimately cause the the results to converge. Clearly the mean signal increases (as it should) as the bounce count is increased. However, it begins to generate diminishing returns and starts to plateau (e.g. compare the mean signal using 50 versus 200 bounces). This suggests that a bounce count of approximately 200 is *reasonable* for the coefficients used in this simulation (although in hindsight, a larger value should have been utilized). However, it also demonstrates that changing the bounce count will not cause the multi-band and single-band mean signals to converge. Without additional research, the exact explanation for the continued variation in mean signal cannot be further explained, except to say that it is clearly a function of the volume (and not the surface), since the trend persists even for completely absorptive targets. These conclusions have been shared with the DIRSIG development team, and since the mean photon density and bounce count have been eliminated as contributors to the discrepancy, future research might consider the actual variance-reduction technique. In terms of the actual physics, the simulations shown here utilize a spectrally flat source, so there is equal probability of generating a photon within each spectral bandpass. This means that the water IOPs are driving the spectral variation in the final result. Since the spectral single scattering albedo determines the probability of an absorption versus a scattering event (per extinction event), this step may be driving the divergence in the mean signal, and therefore a new starting point for continued troubleshooting efforts.

The second observation is the lack of any general trend in SNR for the bright target. Instead, the variance per simulation configuration for the highly reflective target bounces around per target depth and per spectral point. On the other hand, a *slight* trend can be found concerning SNR and dark targets. For these targets, generating spectral bands sequentially and with a *greater* combined mean photon density *can* provide a modest gain in SNR over multi-band simulations. This gain is most strongly associated with shorter wavelengths, increasing depths and larger search radii (see $0.4\mu m$, right column, SNR comparison in Figure 9.18). Comparing the SNR of the single band simulation $[3375 \frac{\text{photons}}{m^3} + 4500 \frac{\text{photons}}{m^2}]$ with the multi-band simulation (designated with an asterisk), the SNR is nearly equivalent when using search radii of 0.25 and 1.0m, but has the great disparity for a search radius of 3.0m, and at depth = 3.5m. Why is this the case? With increasing search radius, the disparity in the variance between single and multi-band simulations decreases. Stated another way, the *difference in the variance* between a sequential and multi-band simulation disappears with increasing search radius. On the other hand, the



(a.) Mean vs. Total Photon Density.

(b.) Mean vs. Bounce Count.

Figure 9.21: Mean radiance for SR 3.0m, at 3.5m depth, for highly reflective targets. (a.) Plot of mean radiance associated with 3-single band simulations (white bar), versus multi-band combination (all other bars, using asterisks designation), at three different mean volume photon densities. (b.) Plot of mean signal versus bounce count. The asterisks are used to designate the multi-band simulations with four different bounce counts of 10, 20, 50 and 200. The white bar is the mean signal from the single-band simulations, using a bounce count of 200. The results indicate that the difference in the mean signal per simulation configuration is not dependent on the total number of photon histories that have been traced, and that increasing the bounce count, although associated with a more accurate mean signal, will not ultimately cause the the results to converge.

difference in the mean is nearly constant between single and multi-band simulations, and across all search radii, also noting that the mean signal for the single-band simulations is always slightly higher (by a few percent) than that for the multi-band simulation. This means that with increasing search radius, the SNR for the single-band simulations can increase at a greater rate than for the multi-band simulation.

Finally, based on all the combined observations, it is apparent that an increase in SNR as a function of parallel computing can only be obtained when the total mean photon density per spectral band is much *greater* than the mean photon density for the equivalent multi-band simulation. The modest difference in SNR per spectral band (and for a fixed search radius) between single and multi-band simulations does not warrant independent per-band-optimization except when memory and simulation efficiency requirements are of concern. The only remaining discrepancy is to ensure that the variance-reduction technique is accurately describing the spectral mean signal, since variations of up to 20% were observed.

9.13.2 Results: Spectral PSF

Spectral PSFs were created by averaging transects of the slanted-ESF and then slightly smoothing the result (without attempting to fit the result using Fermi functions). The spectral PSFs show the same trend as that seen with the panchromatic PSFs; the tails of the distribution broaden well before the central peak is significantly degraded. This is illustrated in Figure 9.22 for SR of 1.0, 2.0 and 3.0m and for a target depth of 1.0m. Figure 9.23 is a repeat of Figure 9.22 but for a target depth of 2.5m. At this greater depth, the SR dependency is lost, as evidenced by the ESFs (and any remaining PSF variation is noise-induced).

Figure 9.24 overlays the ESF and PSF for the red, green and blue spectral bands (dotted line, dashed line, dot-dashed line, respectively) at target depths of 1.0 (top row) and 2.5m (middle row) and using a fixed 3.0m search radius. The bottom row depicts the ESFs from targets at 1.0m, but normalized by the maximum value of the ESF (left), and the total or sum of the ESF (right). Although the red spectral band makes the smallest contribution to sensor reaching radiance, and the green band, the most, after normalization, variations in the ESF per spectral band are very minor.

9.13.3 Summary of Spectral Considerations

Several broad conclusions can be drawn from the spectral results shown here:

1. Spectral SNR varies most strongly with the mean photon density of the volume, as opposed to the mean photon density on the surface of the ESF. This is likely true provided a sufficient number of photons occupy the surface, and for the scenarios investigated here, $1500 \frac{\text{photons}}{\text{m}^2}$ was sufficient, and very little additional gain in SNR was incurred when this was increased to $4500 \frac{\text{photons}}{\text{m}^2}$.
2. Generating a 3-band or 3-single band simulations does not buy a considerable increase in SNR. The user should therefore consider memory requirements and simulation efficiency as equally important when deciding how best to simulate multi- or hyperspectral data.
3. The spectral PSFs follow the same trend as the panchromatic PSFs in that both are robust, showing only moderate bias when associated with reasonable increases in search radius. A band-specific solution for SR is unwarranted, at least given the results to-date.

Depth = 1.0m

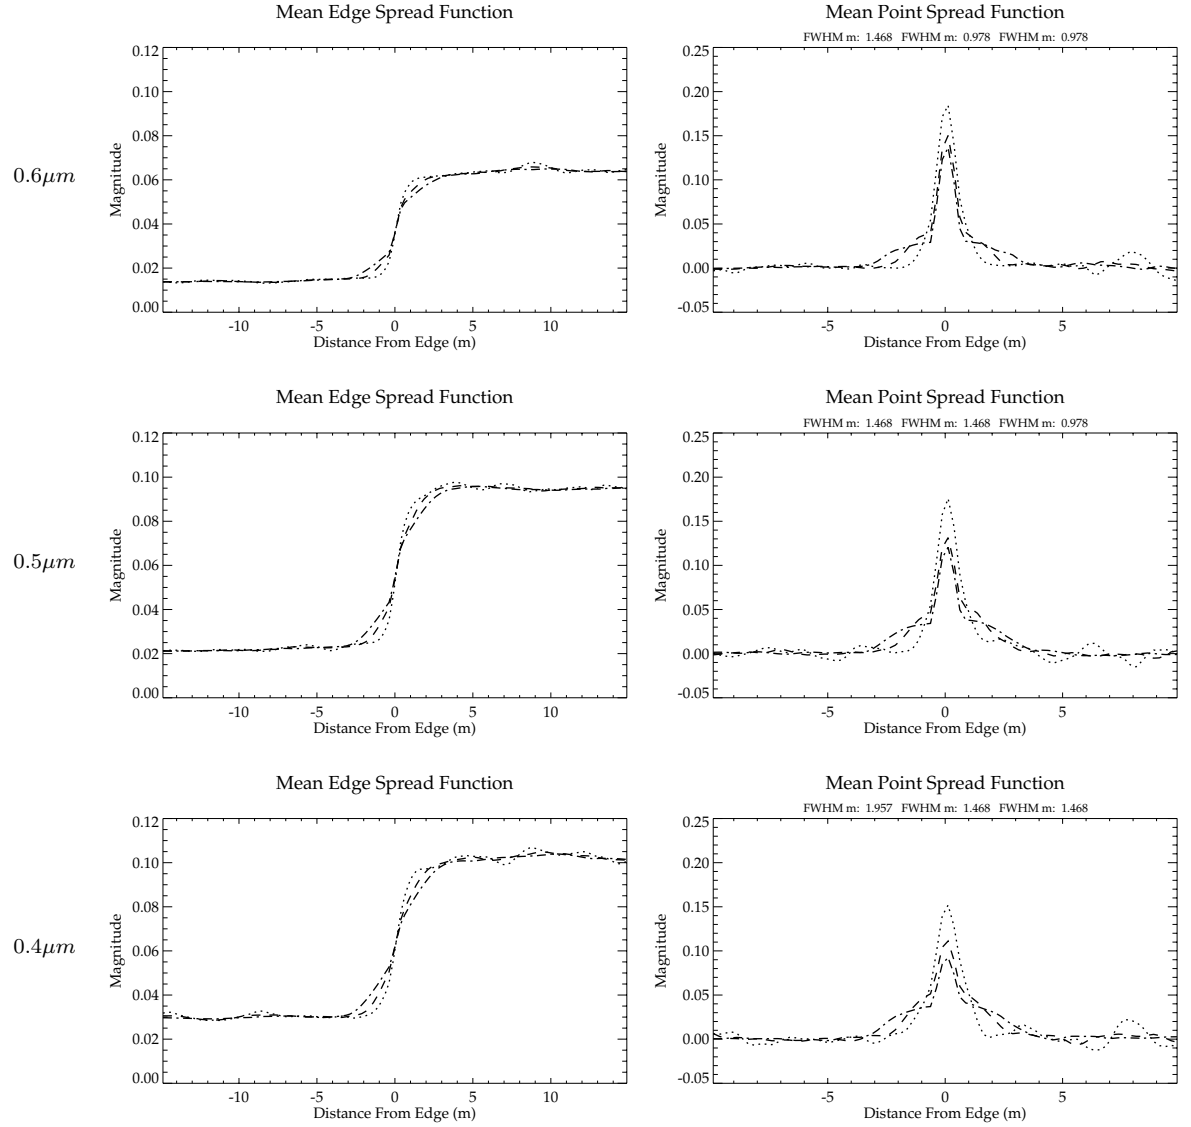


Figure 9.22: Spectral ESFs and PSFs associated with SR of 1.0, 2.0 and 3.0m (dotted line, dashed line, dot-dashed line, respectively) for red, blue and green spectral bands (top to bottom) with a target depth of 1.0m.

Depth = 2.5m

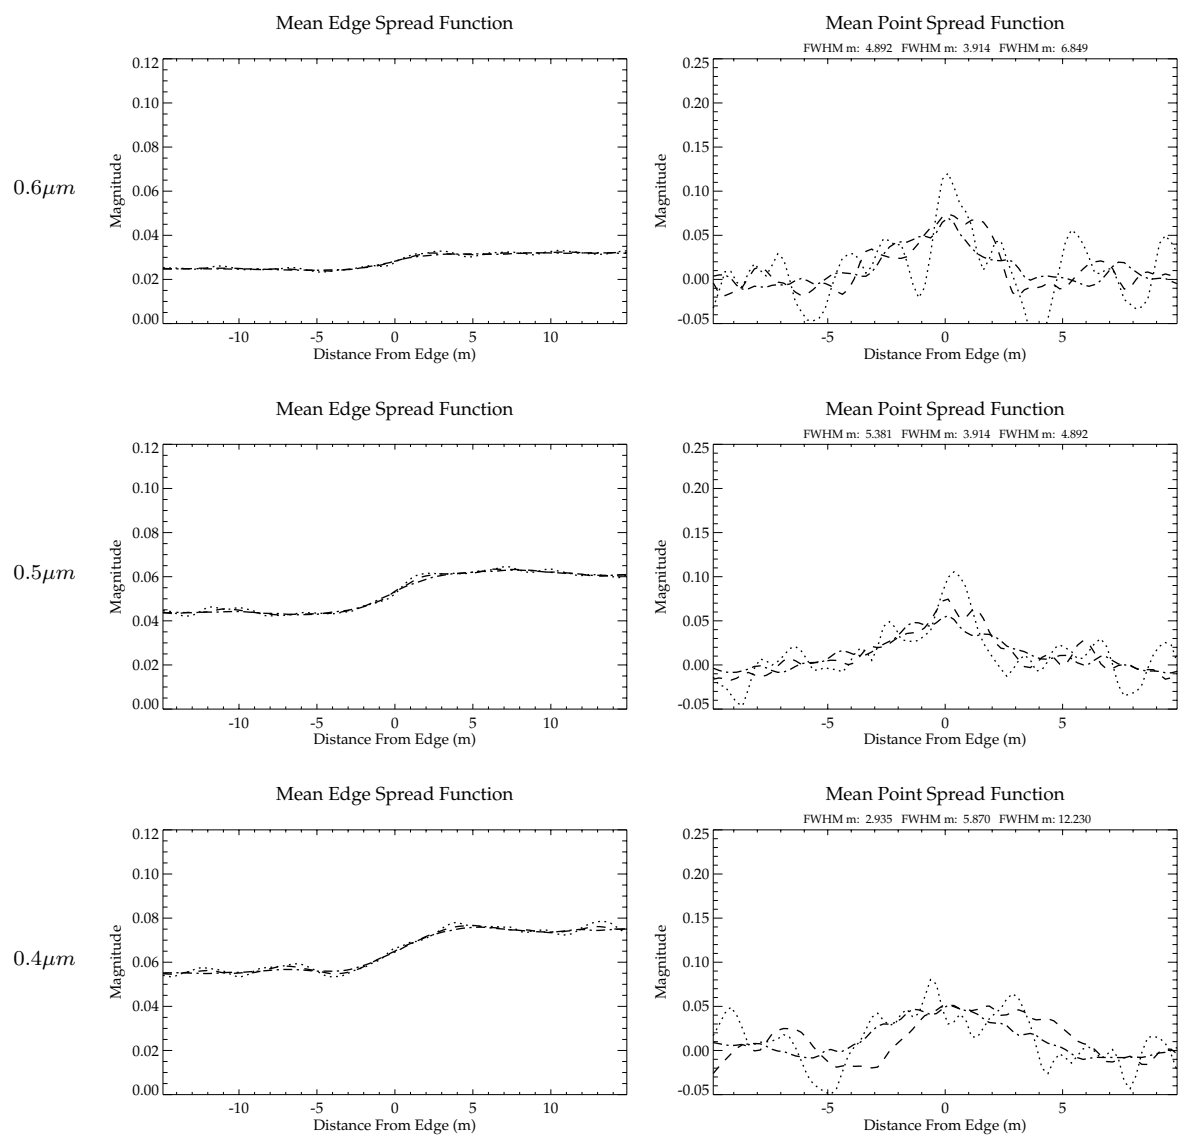


Figure 9.23: Spectral ESFs and PSFs associated with SR of 1.0, 2.0 and 3.0m (dotted line, dashed line, dot-dashed line, respectively) for red, blue and green spectral bands (top to bottom) with a target depth of 2.5m.

Depth ↓

Search Radius = 3.0m

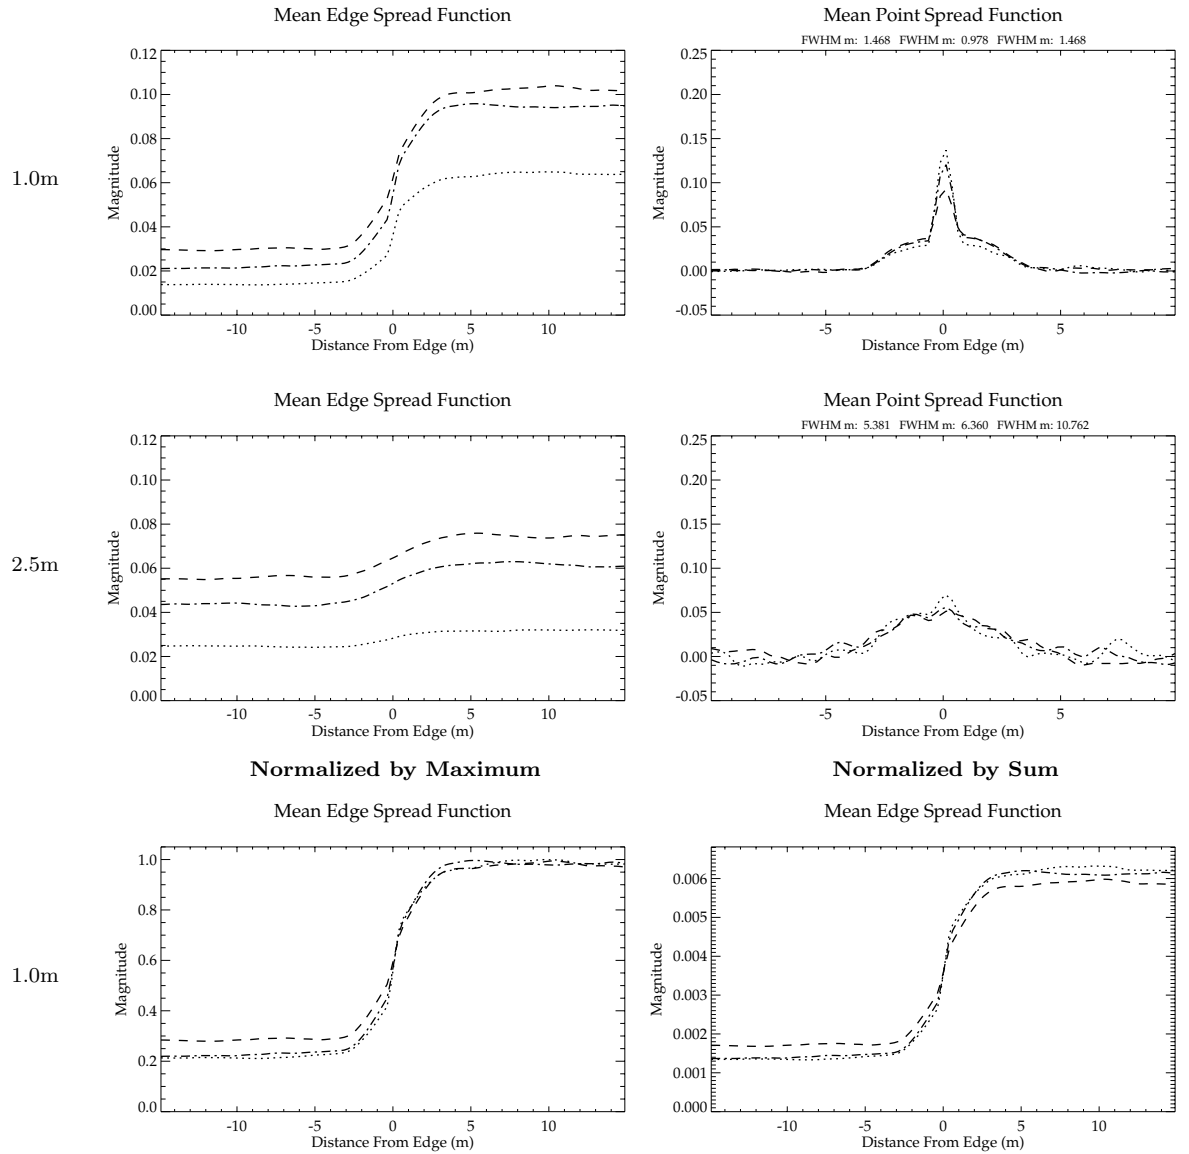


Figure 9.24: Spectral ESFs and PSFs associated with SR of 3.0m, for red, green and blue spectral bands (dotted line, dashed line, dot-dashed line, respectively), for target depths of 1.0 (top row) and 2.5m (middle row). The bottom row depicts the ESFs from targets at 1.0m, but normalized by the maximum value of the ESF (left), and the total or sum of the ESF (right).

9.14 Image Quality Overview

This chapter examined the SNR, variance and bias associated with ESF targets at variable depths and as viewed through water volumes with variable IOPs. The results indicate that an SNR of at least 50 is desirable for bright targets, but that the user must make a subjective assessment of quality when it comes to the SNR of dark targets, trading a “speckled” appearance for a “mottled” appearance, or vice versa. However, when the SNR converges for both dark and bright targets, the target’s geometrical and/or optical depth is such that very little image transfer will be obtained. For these situations, a 3D photon mapped solution may not be warranted and/or very high search radii can be used since bias will not be relevant (excluding convolution concerns).

Spectrally, very little SNR is gained by generating high photon density single band simulations, versus a multiband simulation. This was actually very surprising (and contrary to the author’s intuition), and yet nicely highlights an effective variance-reduction tool. It also suggests that the user should consider memory and efficiency requirements as equally important to SNR when deciding how best to generate multi- or hyperspectral photon mapped data.

With regard to PSF bias, recent changes to the computational mechanisms within the SurfacePM and MediumPM Radiometry Solvers (boundary compensation and the effective scattering coefficient, respectively) seems to generate some degree of “edge sharpening” despite a slowly mounting bias adjacent to actual edges. This allows the user to specify a large search radius, with the consequent gain in SNR, without incurring a large bias penalty (examined here for search radii that are 12 times the GSD). It is therefore argued that SNR should be the dominant image quality metric that drives the Radiometry Solver configuration.

9.15 Collection of Panchromatic Images

The following caption is applicable to all Figures shown in this section:

- **Caption**

The approximate mean photon density is listed in top left corner ($\frac{\text{photons}}{m^3}$). Each row represents a fixed target depth from the surface (*e.g.* 1.0, 2.5, 3.5 and 5.0m) as labeled (top to bottom). Each column depicts a fixed search radius (increasing across a row from left to right) with values of 0.25, 0.5, 1.0, 1.5, 2.0, 2.5 and 3.0m (or SR = GIFOV, SR = 2xGIFOV, etc.). The left-most column reports the water type, IOPs and phase function (P = Petzold, or g = average cosine of one-parameter Henyey-Greenstein phase function).

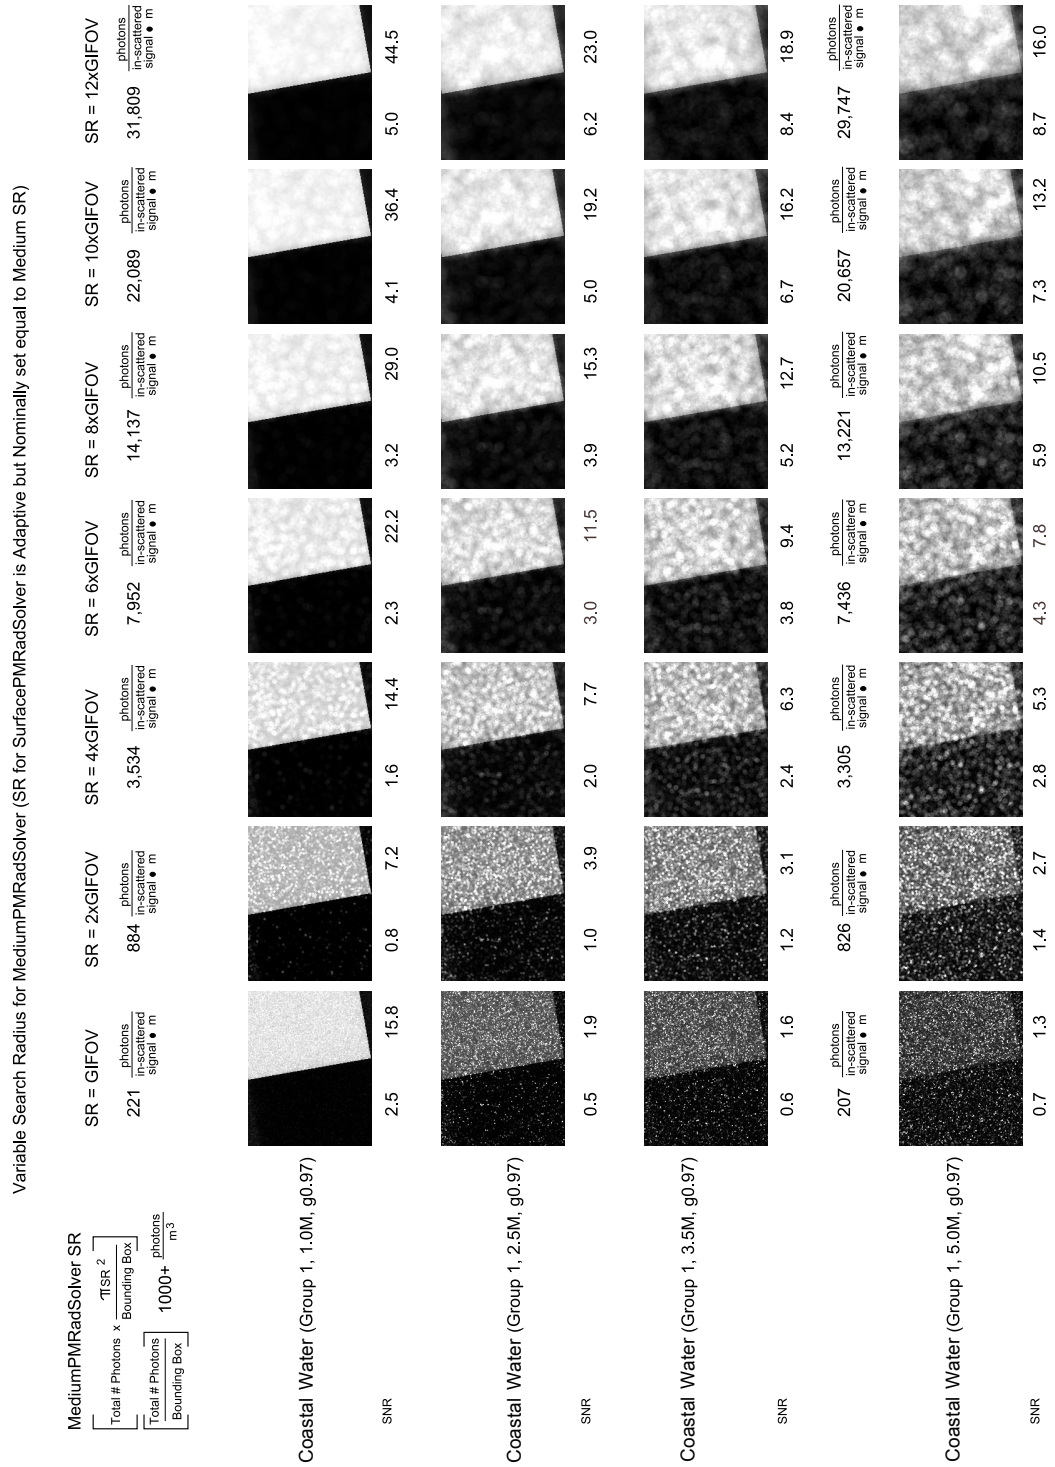


Figure 9.25: See Section 9.15 for Figure caption.

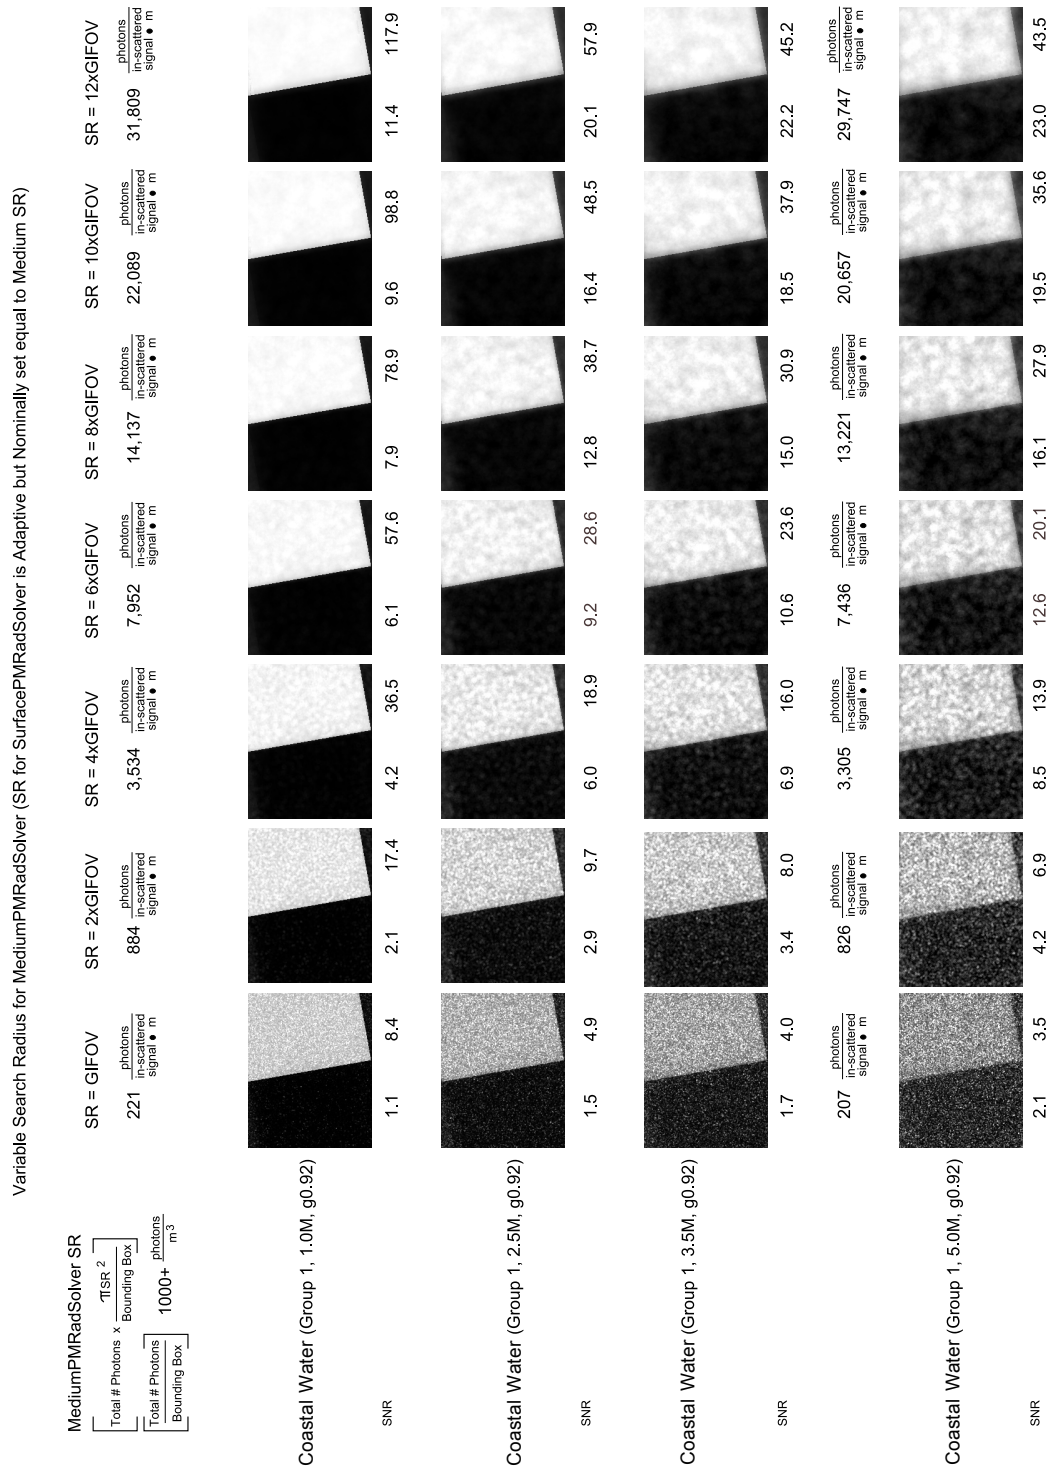


Figure 9.26: See Section 9.15 for Figure caption.

MediumPMRadSolver SR		SR = GIFOV		SR = 2xGIFOV		SR = 4xGIFOV		SR = 6xGIFOV		SR = 8xGIFOV		SR = 10xGIFOV		SR = 12xGIFOV	
$\left\lceil \frac{\text{Total \# Photons} \times \frac{\pi \text{SR}^2}{\text{Bounding Box}}}{1000+} \right\rceil \frac{\text{photons}}{\text{m}^3}$		photons in-scattered signal ● m	photons in-scattered signal ● m	photons in-scattered signal ● m	photons in-scattered signal ● m	photons in-scattered signal ● m	photons in-scattered signal ● m	photons in-scattered signal ● m	photons in-scattered signal ● m	photons in-scattered signal ● m	photons in-scattered signal ● m	photons in-scattered signal ● m	photons in-scattered signal ● m	photons in-scattered signal ● m	photons in-scattered signal ● m
Coastal Water (Group 1, 1.0M, g0.85)															
SNR		2.3	14.2	4.6	31.1	9.2	62.5	14.0	91.6	18.7	119.8	23.4	149.1	27.8	174.6
Coastal Water (Group 1, 2.5M, g0.85)															
SNR		3.0	9.1	5.9	18.0	11.6	35.4	17.0	53.3	22.4	70.7	27.5	87.6	32.6	103.3
Coastal Water (Group 1, 3.5M, g0.85)															
SNR		3.4	7.8	6.6	15.6	13.0	31.0	19.4	46.5	25.5	62.9	31.3	81.0	36.5	99.6
Coastal Water (Group 1, 5.0M, g0.85)															
SNR		4.0	6.5	8.1	12.8	16.4	25.8	25.0	38.4	33.5	50.2	41.6	60.4	49.6	70.1

Variable Search Radius for MediumPMRadSolver (SR for SurfacePMRadSolver is Adaptive but Nominally set equal to Medium SR)

MediumPMRadSolver SR

$$\left[\frac{\text{Total \# Photons} \times \frac{\pi \text{SR}^2}{\text{Bounding Box}}}{\frac{\text{Total \# Photons}}{\text{Bounding Box}}} \right] 1000 + \frac{\text{photons}}{\text{m}^3}$$


[illegible]

Figure 9.28: See Section 9.15 for Figure caption.

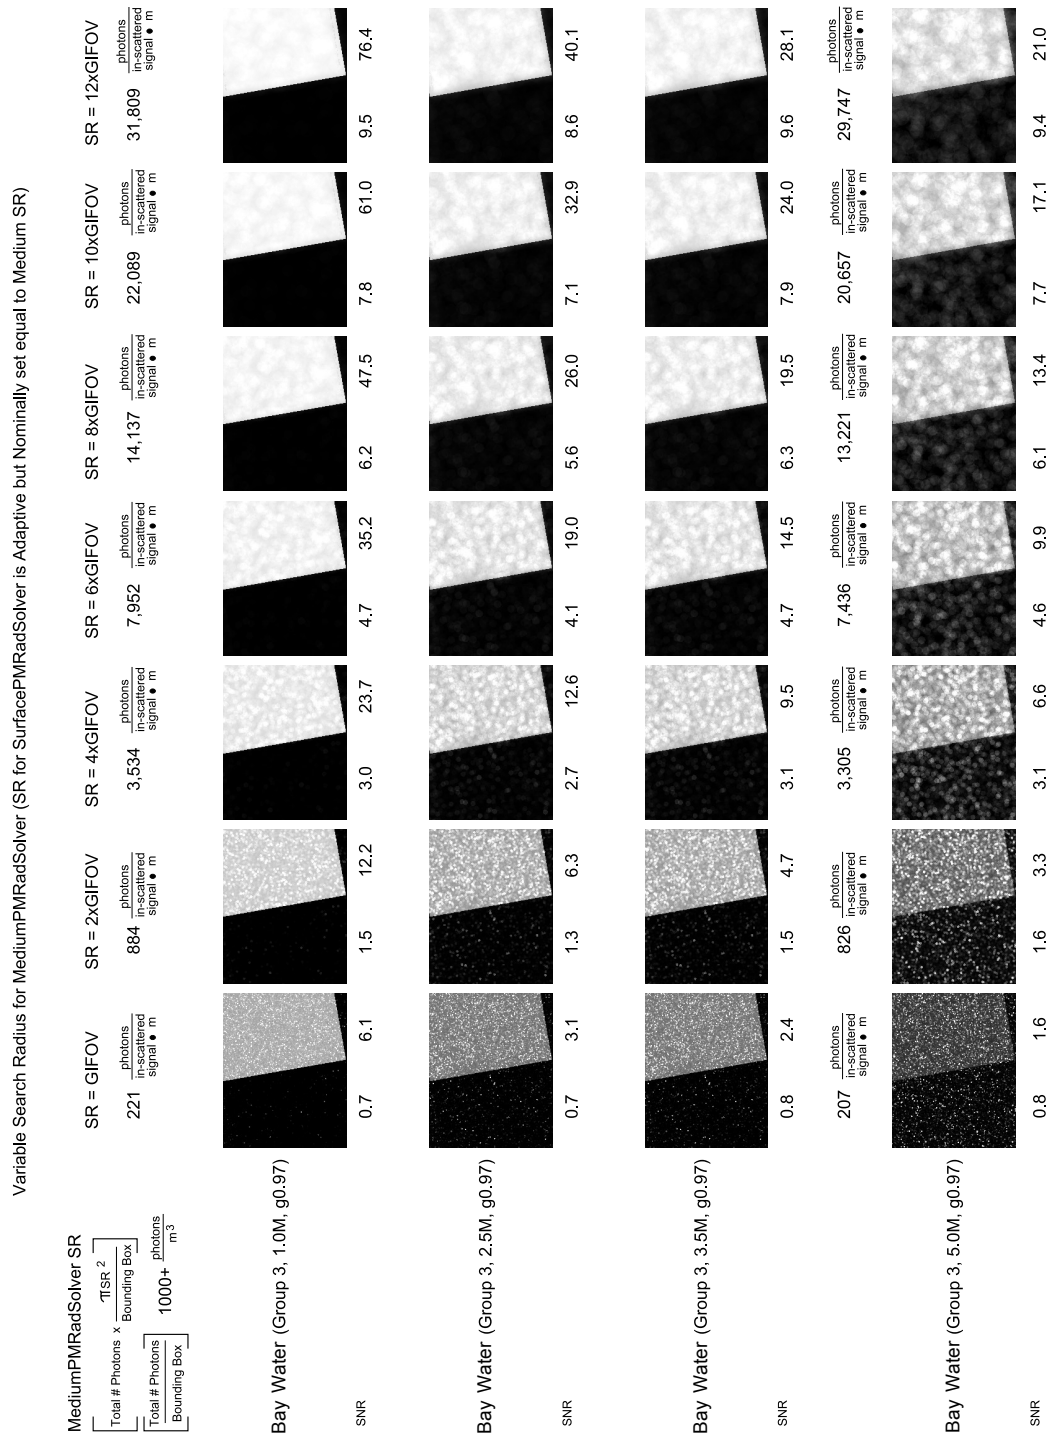


Figure 9.29: See Section 9.15 for Figure caption.

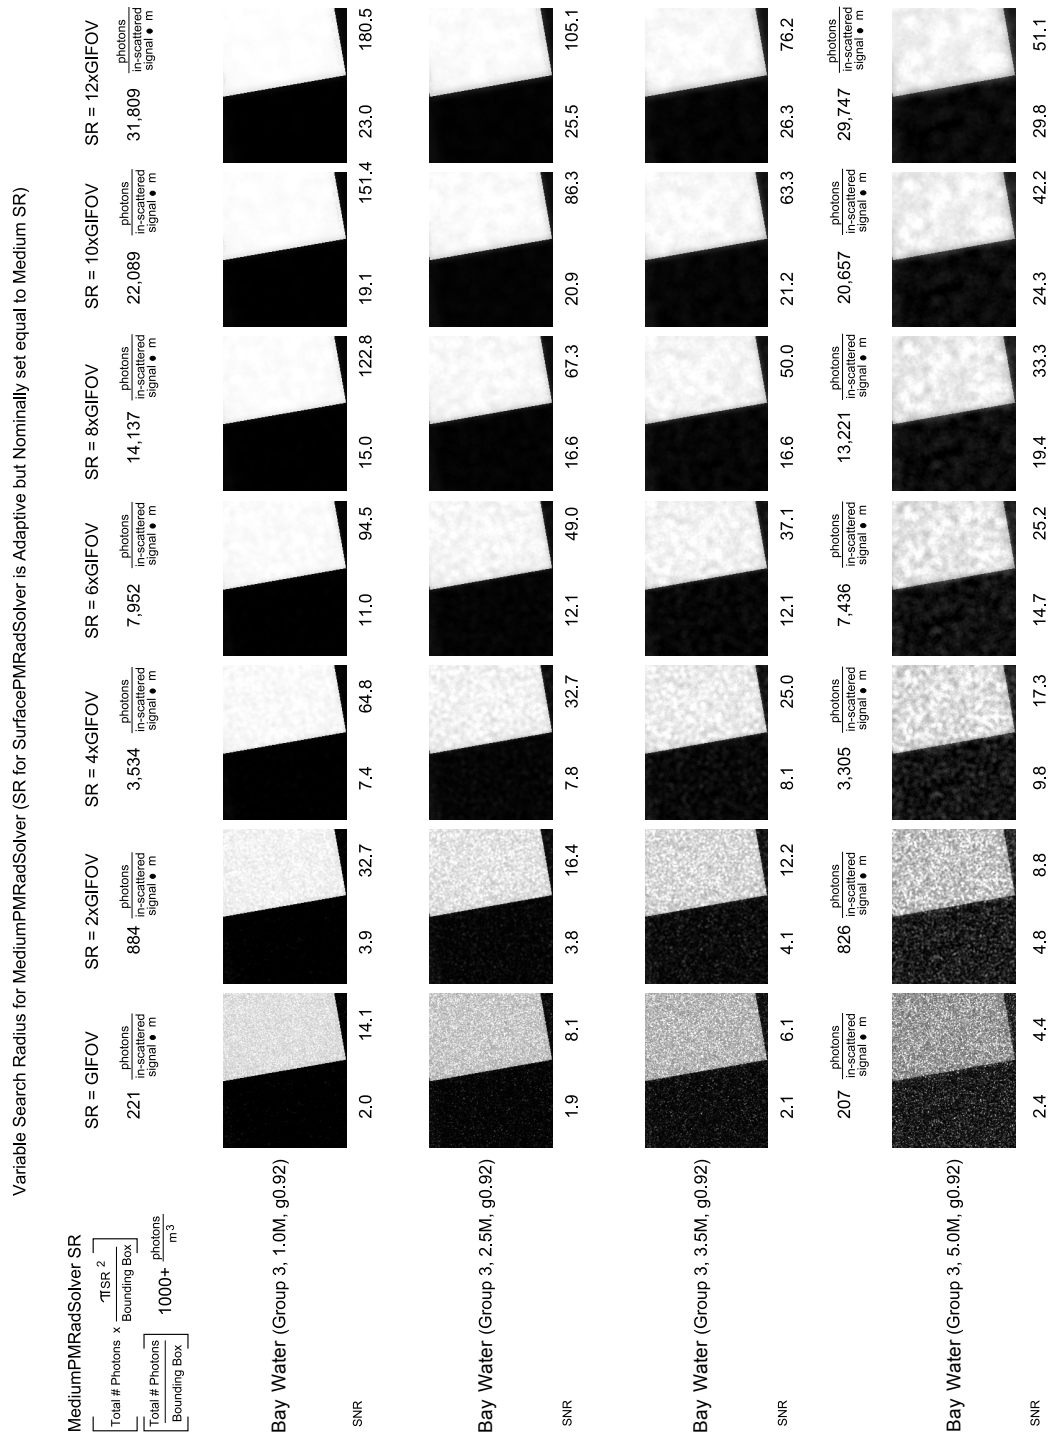


Figure 9.30: See Section 9.15 for Figure caption.

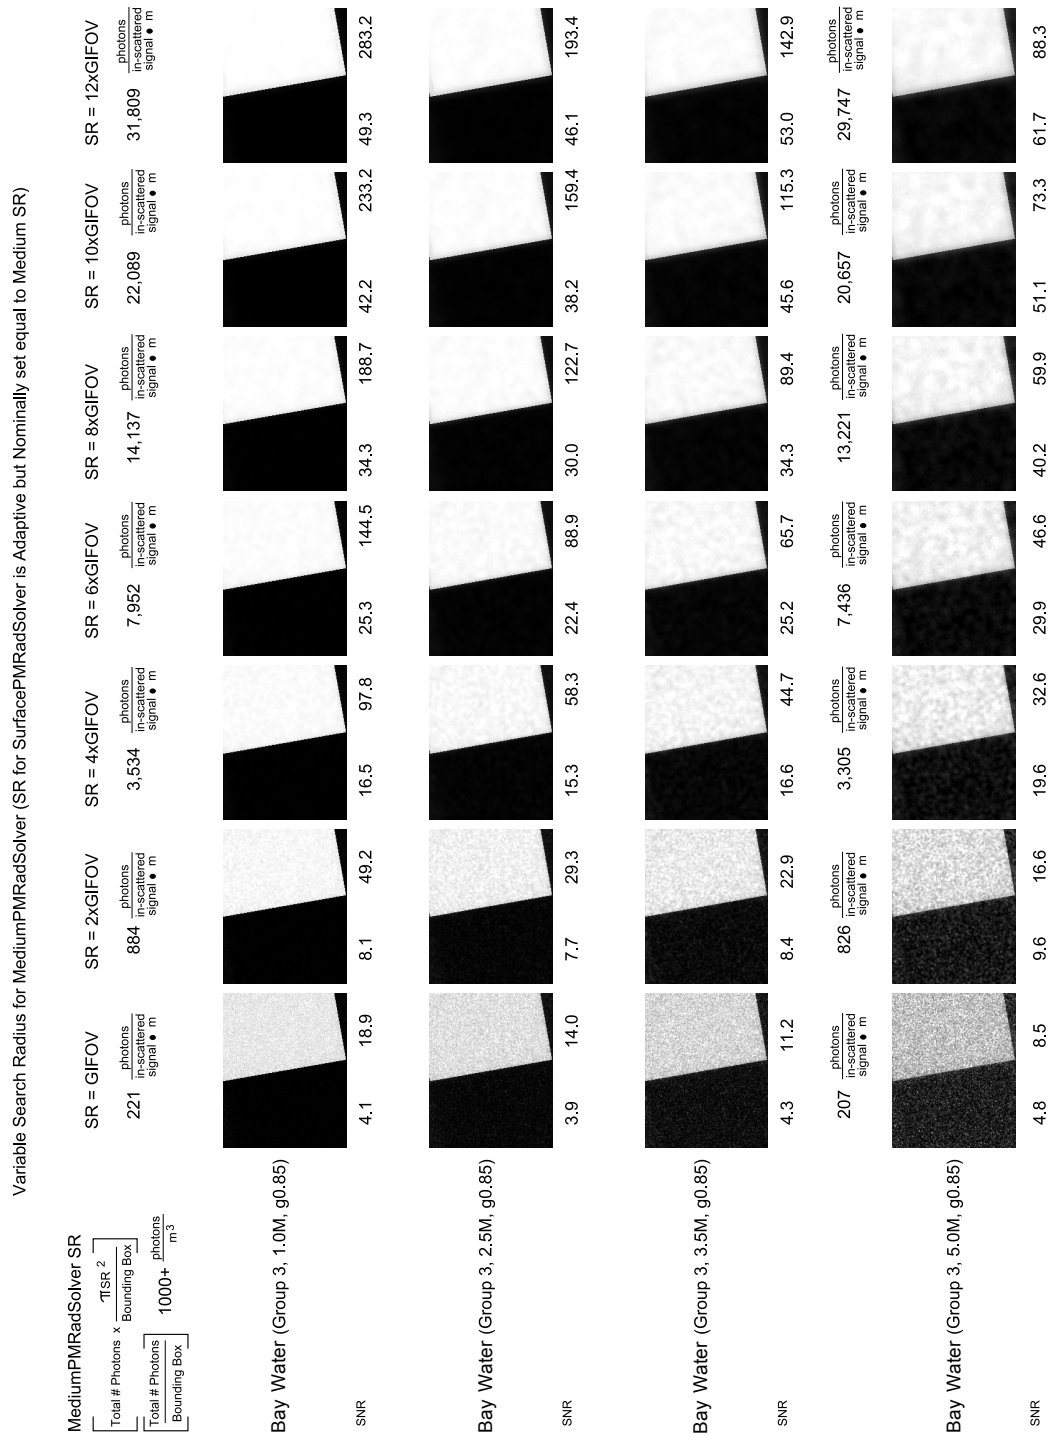


Figure 9.31: See Section 9.15 for Figure caption.

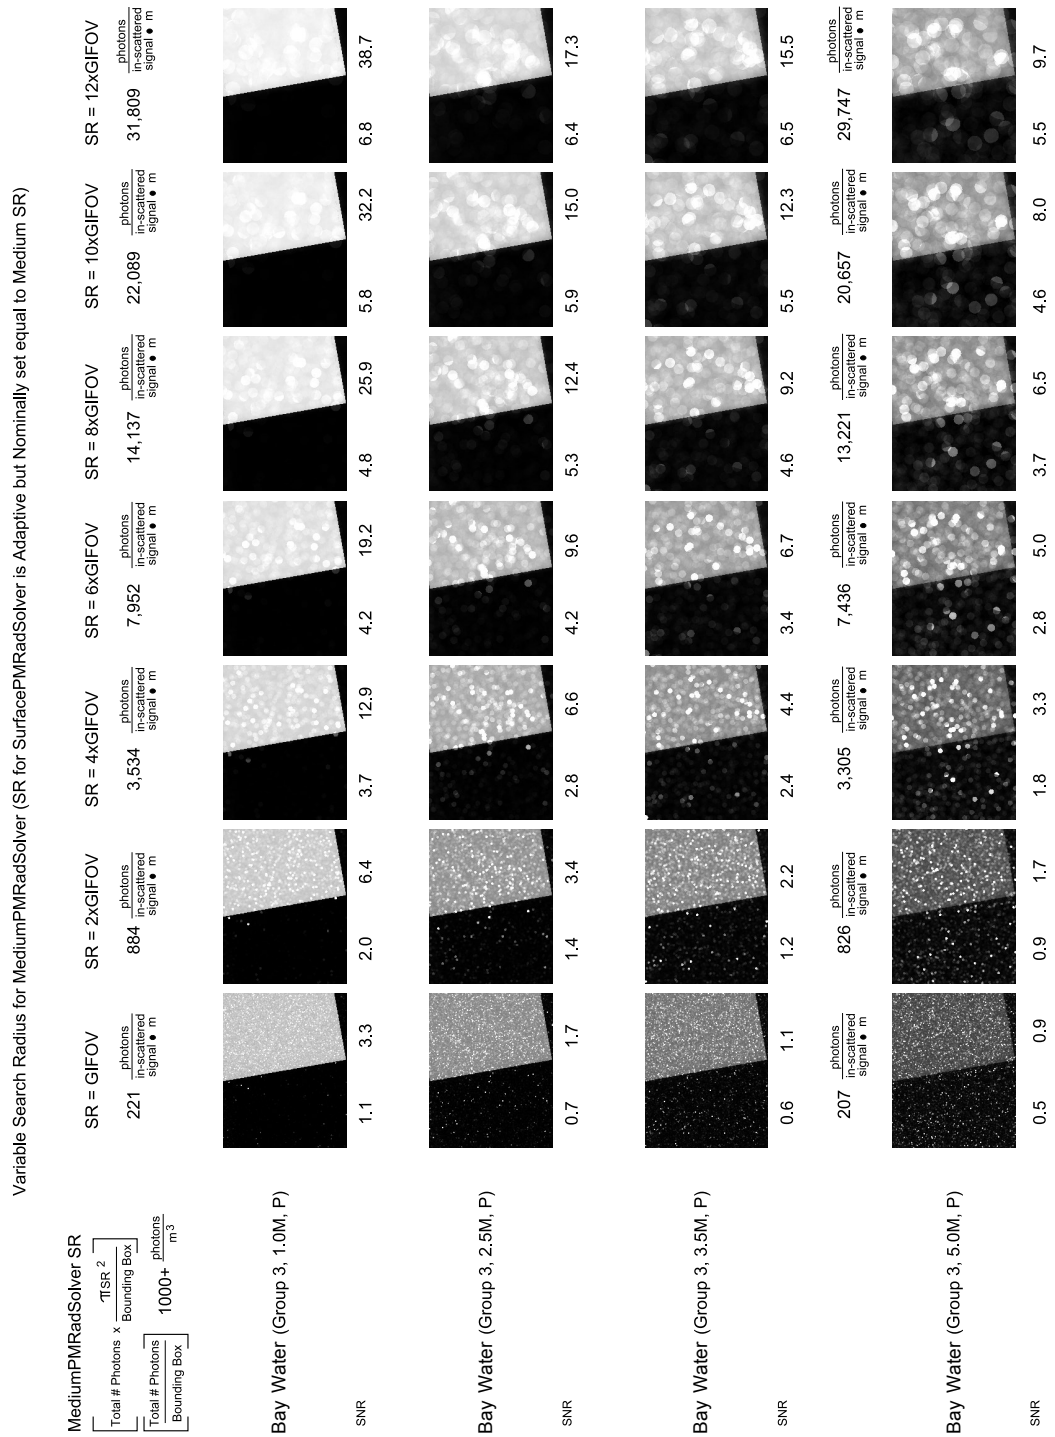


Figure 9.32: See Section 9.15 for Figure caption.

MediumPMRadSolver SR

SR = GIFOV	SR = 2xGIFOV	SR = 4xGIFOV	SR = 6xGIFOV	SR = 8xGIFOV	SR = 10xGIFOV	SR = 12xGIFOV
$\left[\frac{\text{Total \# Photons}}{\text{Bounding Box}} \times \frac{\pi \text{SR}^2}{\text{Bounding Box}} \right]$ $\left[\frac{\text{Total \# Photons}}{\text{Bounding Box}} \right] 1000+ \frac{\text{photons}}{\text{m}^3}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$ 221	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$ 884	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$ 3,534	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$ 7,952	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$ 14,137	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$ 31,809
0.4	0.8	3.1	1.6	6.0	12.4	20.0
0.6	1.2	1.6	2.5	3.4	5.1	9.8
0.7	1.4	1.6	2.8	3.1	5.7	11.1
207	826	3,305	7,436	13,221	20,657	29,747
$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$

Turbid Water (Group 4b, 1.0M, g0.97)

Turbid Water (Group 4b, 2.5M, g0.97)

Turbid Water (Group 4b, 3.5M, g0.97)

Turbid Water (Group 4b, 5.0M, g0.97)

SNR

[illegible]

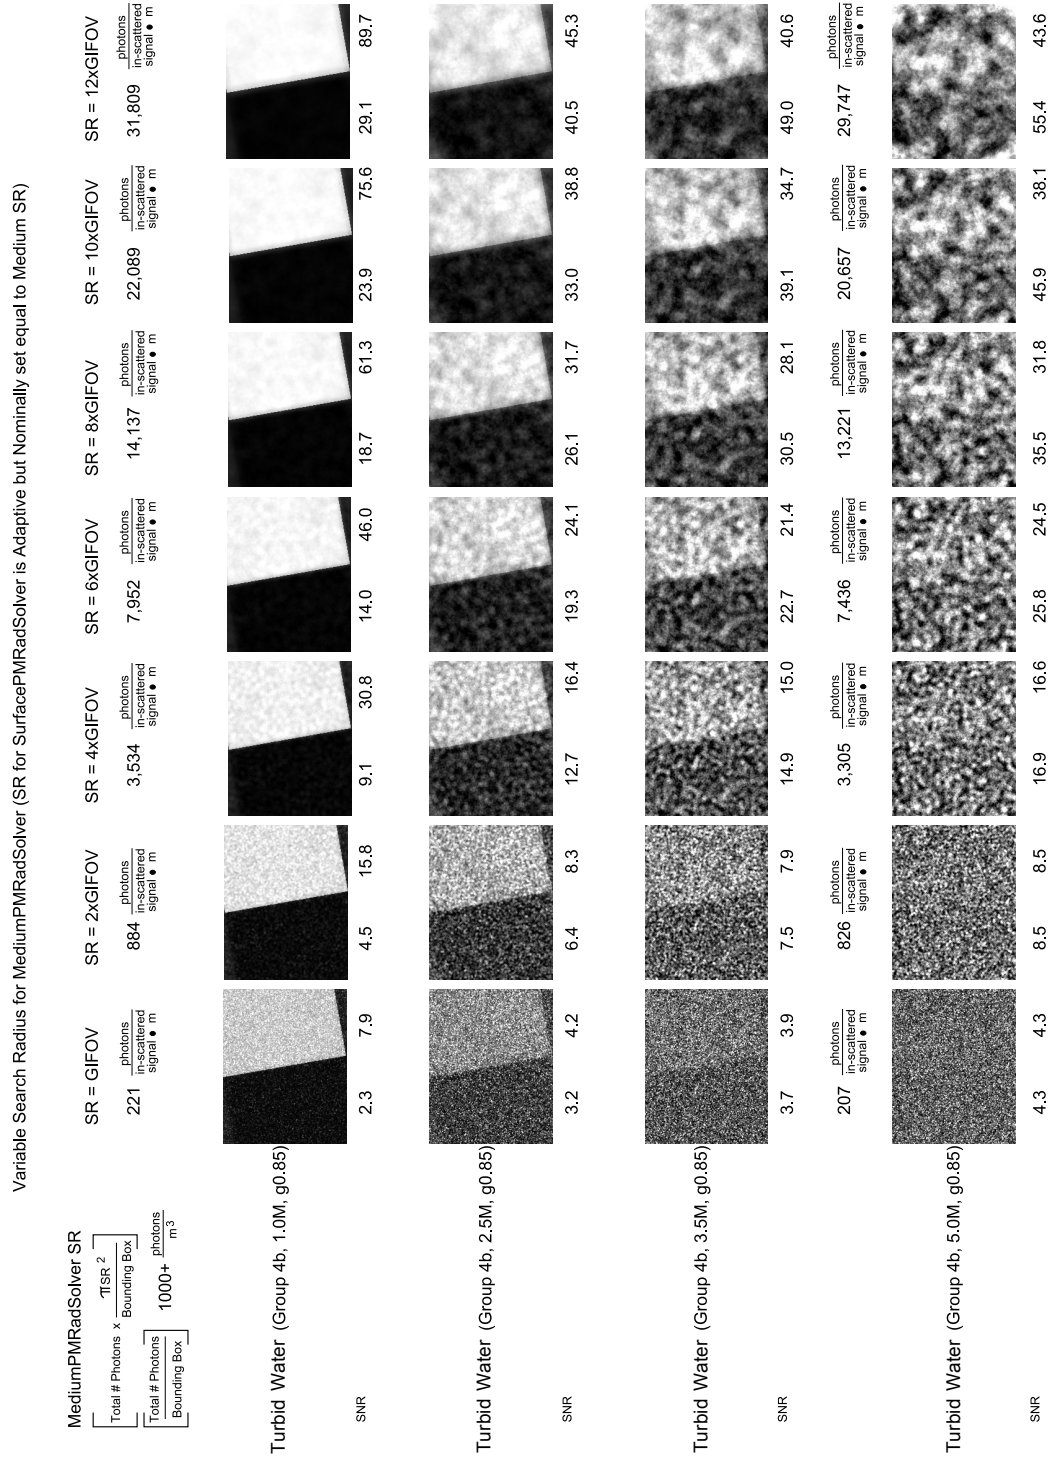


Figure 9.35: See Section 9.15 for Figure caption.

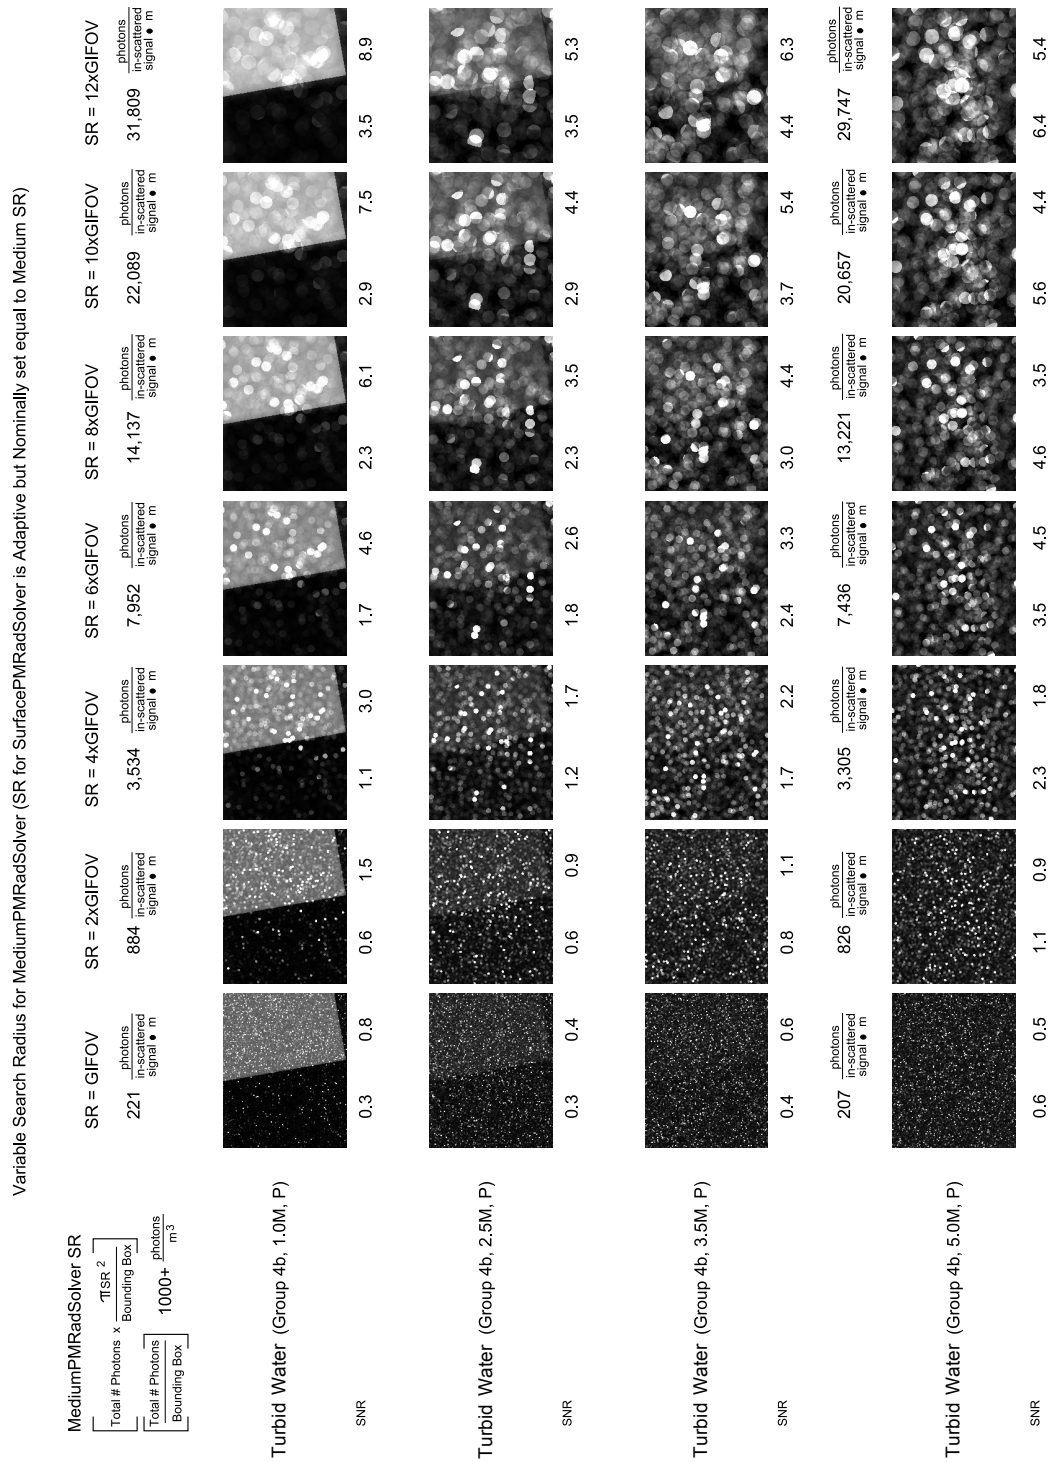


Figure 9.36: See Section 9.15 for Figure caption.

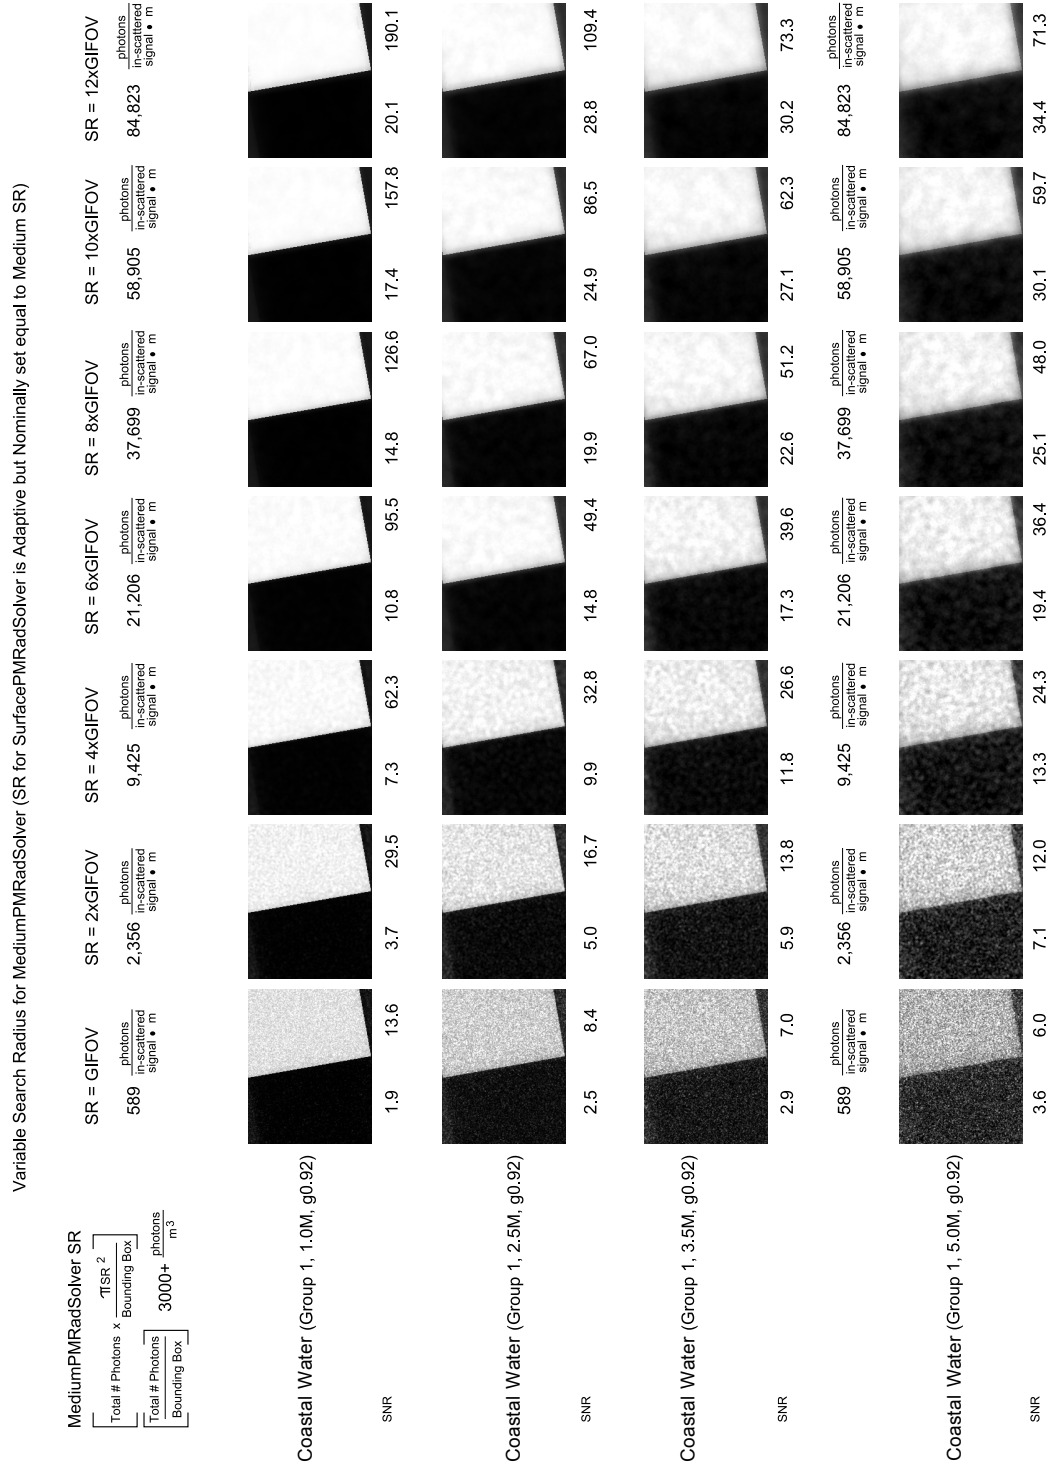


Figure 9.37: See Section 9.15 for Figure caption.

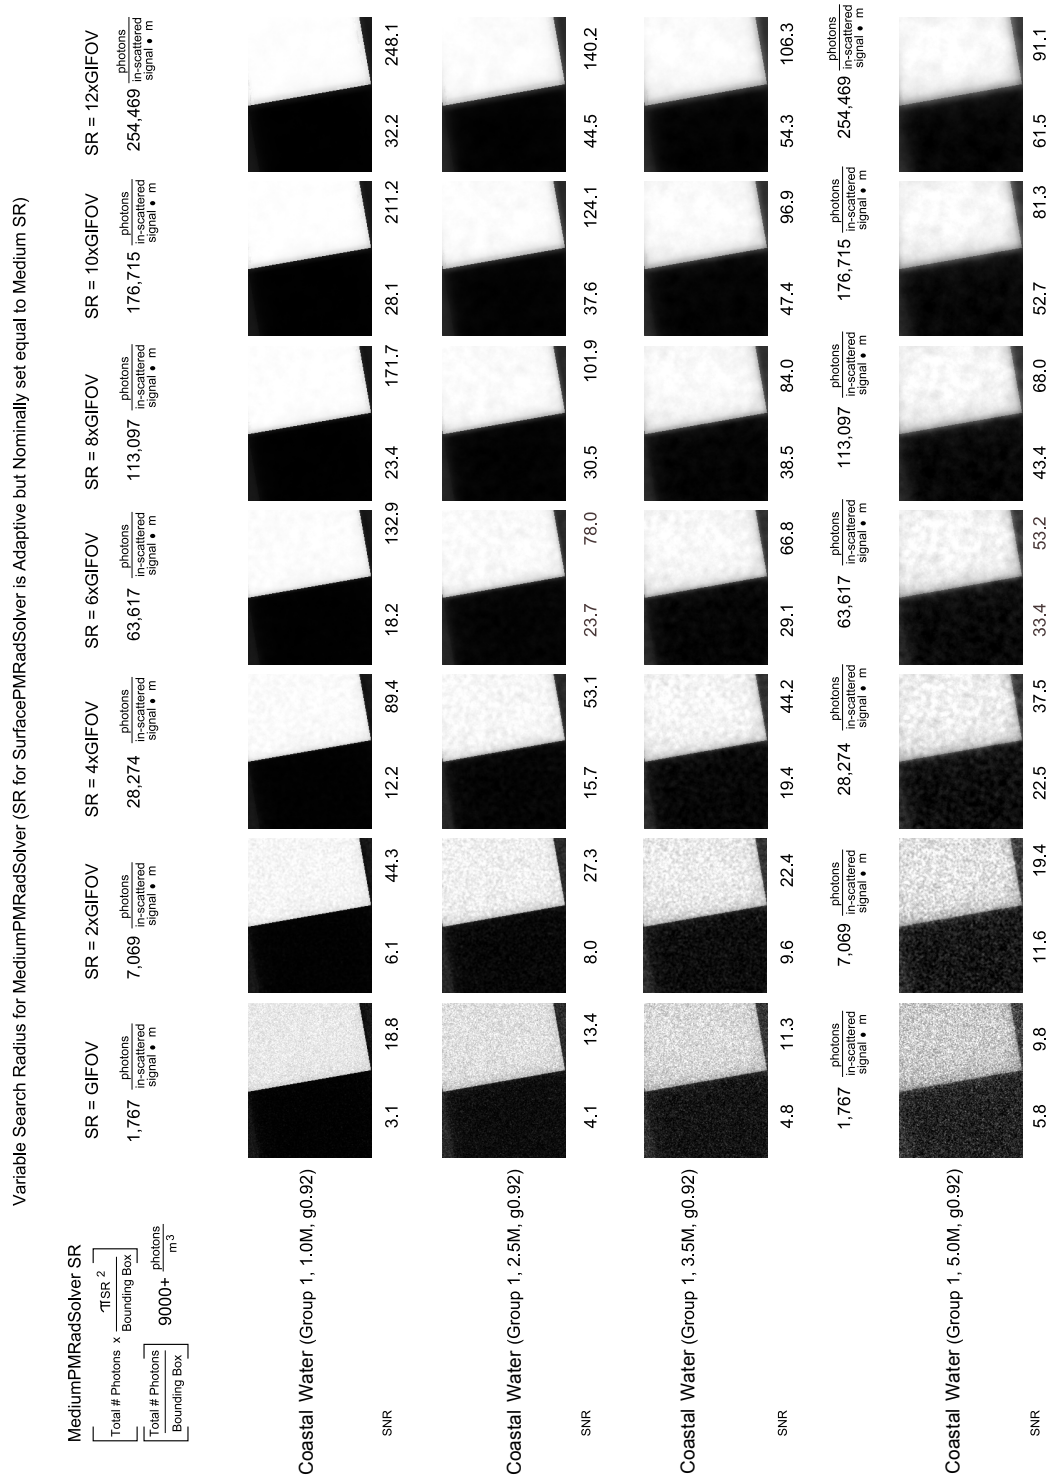


Figure 9.38: See Section 9.15 for Figure caption.

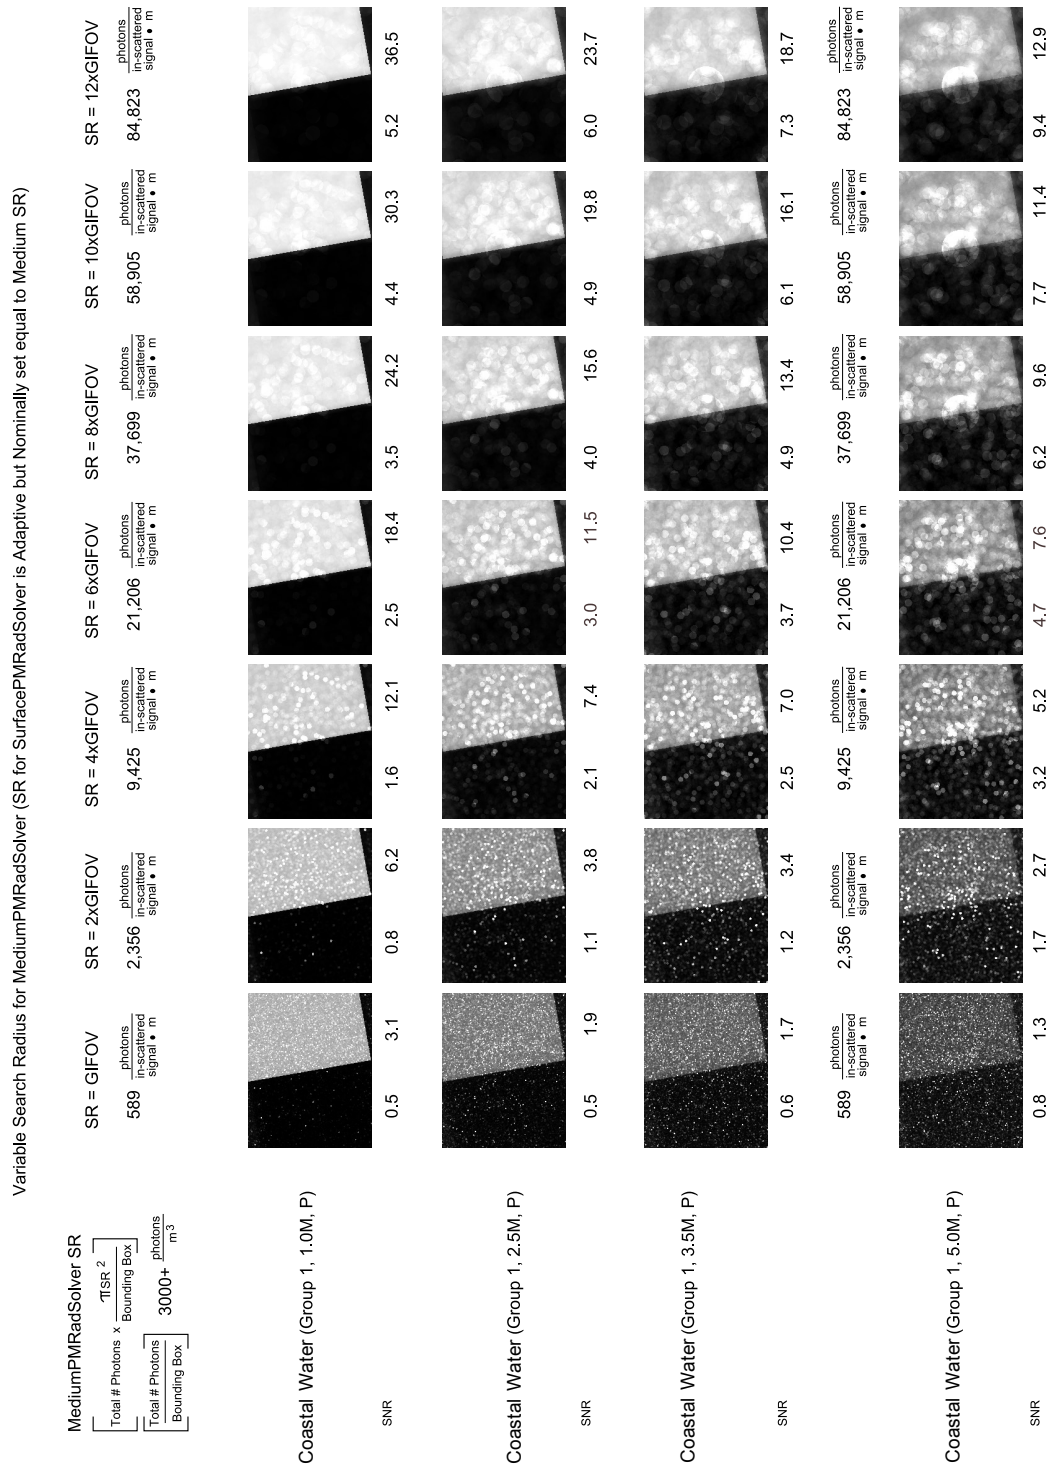


Figure 9.39: See Section 9.15 for Figure caption.

MediumPMRadSolver SR		SR = GIFOV		SR = 2xGIFOV		SR = 4xGIFOV		SR = 6xGIFOV		SR = 8xGIFOV		SR = 10xGIFOV		SR = 12xGIFOV		
$\left\lceil \frac{\text{Total \# Photons}}{\text{Bounding Box}} \times \frac{\pi \text{SR}^2}{\text{Bounding Box}} \right\rceil$	$\left\lceil \frac{\text{Total \# Photons}}{\text{Bounding Box}} \right\rceil$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	0.8	5.2	1.4	10.2	2.9	20.7	4.2	31.6	5.5	40.6	6.8	49.1	7.9	57.5
Coastal Water (Group 1, 1.0M, P)																
SNR			0.8	5.2	1.4	10.2	2.9	20.7	4.2	31.6	5.5	40.6	6.8	49.1	7.9	57.5
Coastal Water (Group 1, 2.5M, P)																
SNR			0.8	3.0	1.6	5.8	3.3	11.3	5.0	16.6	6.5	21.4	8.8	23.6	8.8	26.4
Coastal Water (Group 1, 3.5M, P)																
SNR			1.0	2.6	1.9	5.0	3.6	9.6	4.9	13.7	6.1	17.1	7.0	20.3	7.8	22.8
Coastal Water (Group 1, 5.0M, P)																
SNR			1.3	2.0	2.6	4.2	4.9	8.1	6.9	11.5	8.5	14.3	9.9	16.5	10.8	18.3

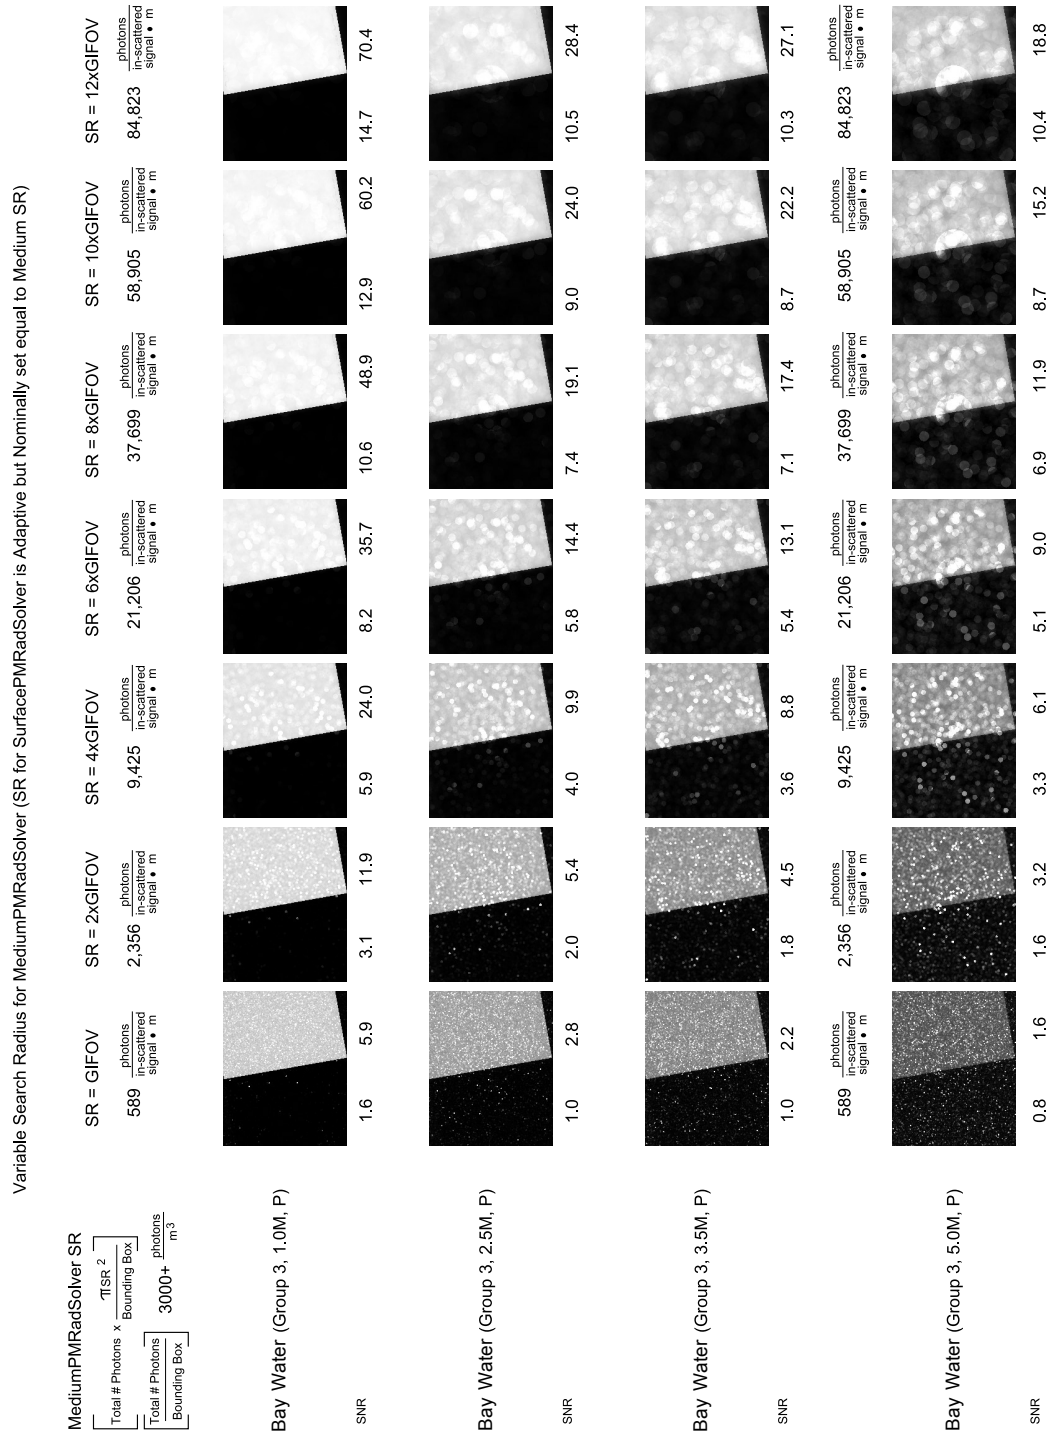


Figure 9.41: See Section 9.15 for Figure caption.

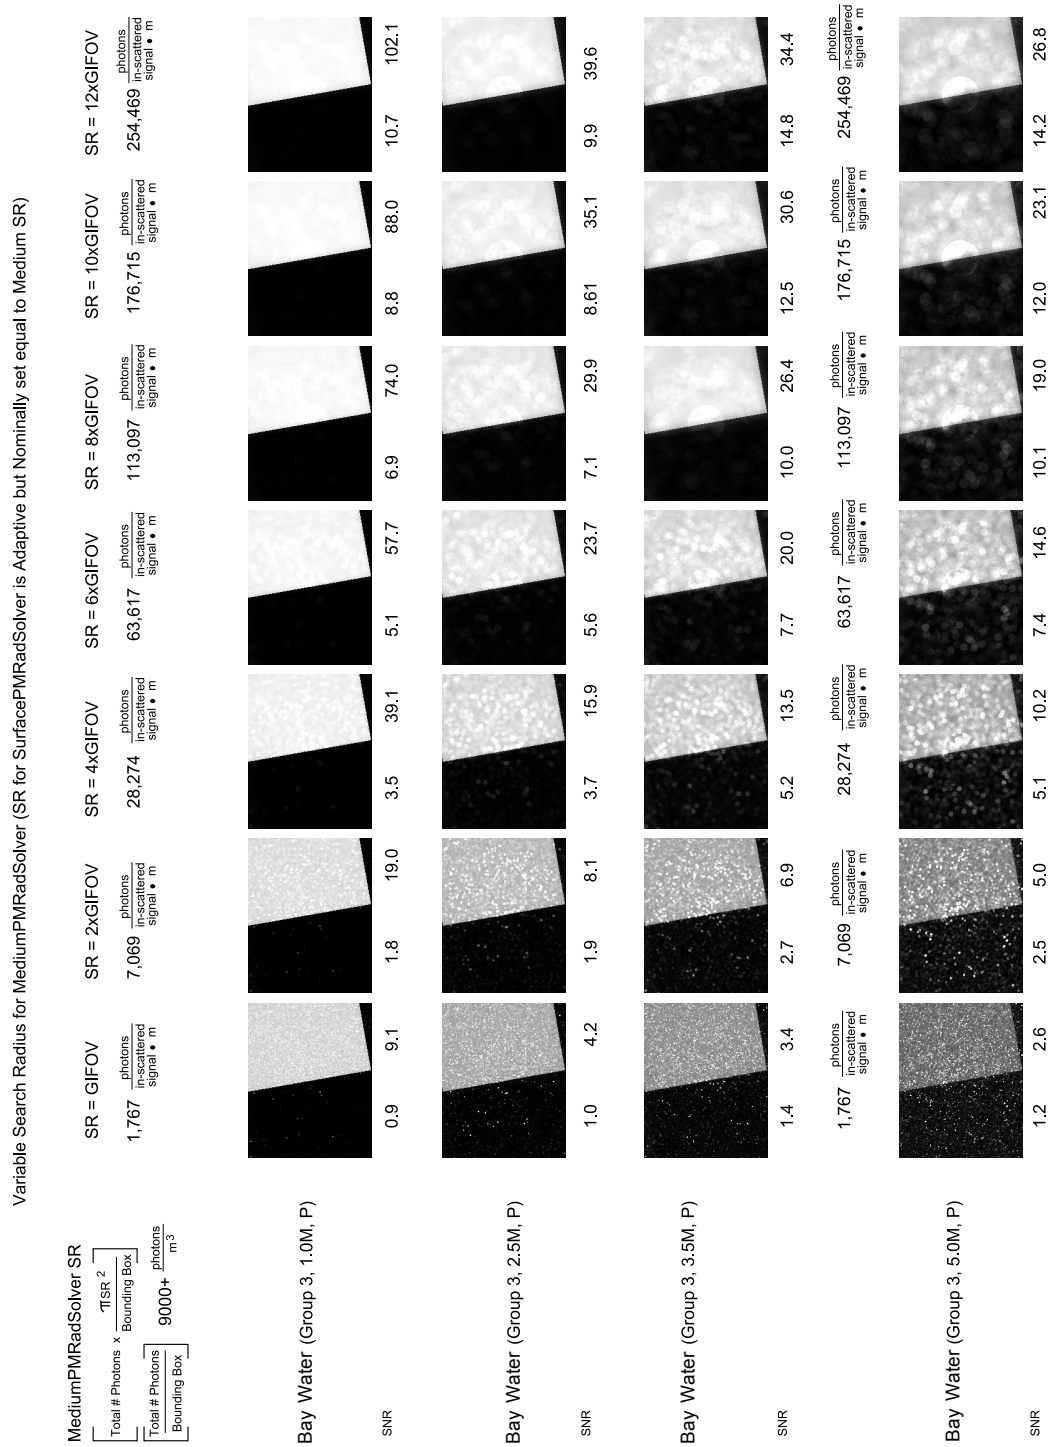


Figure 9.42: See Section 9.15 for Figure caption.

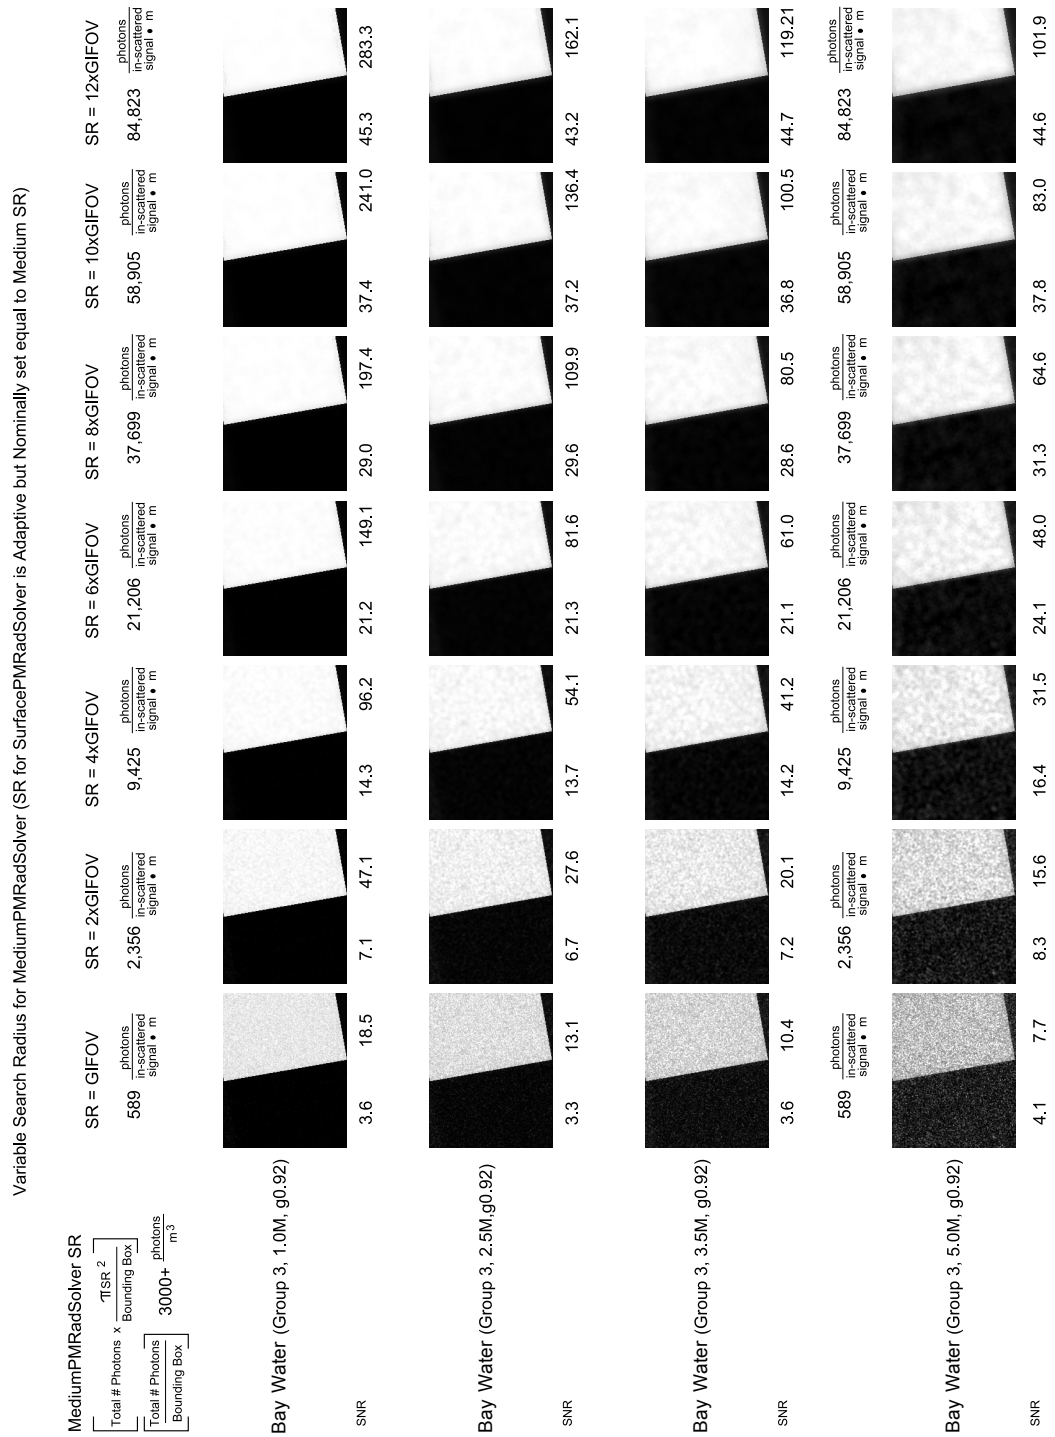


Figure 9.43: See Section 9.15 for Figure caption.

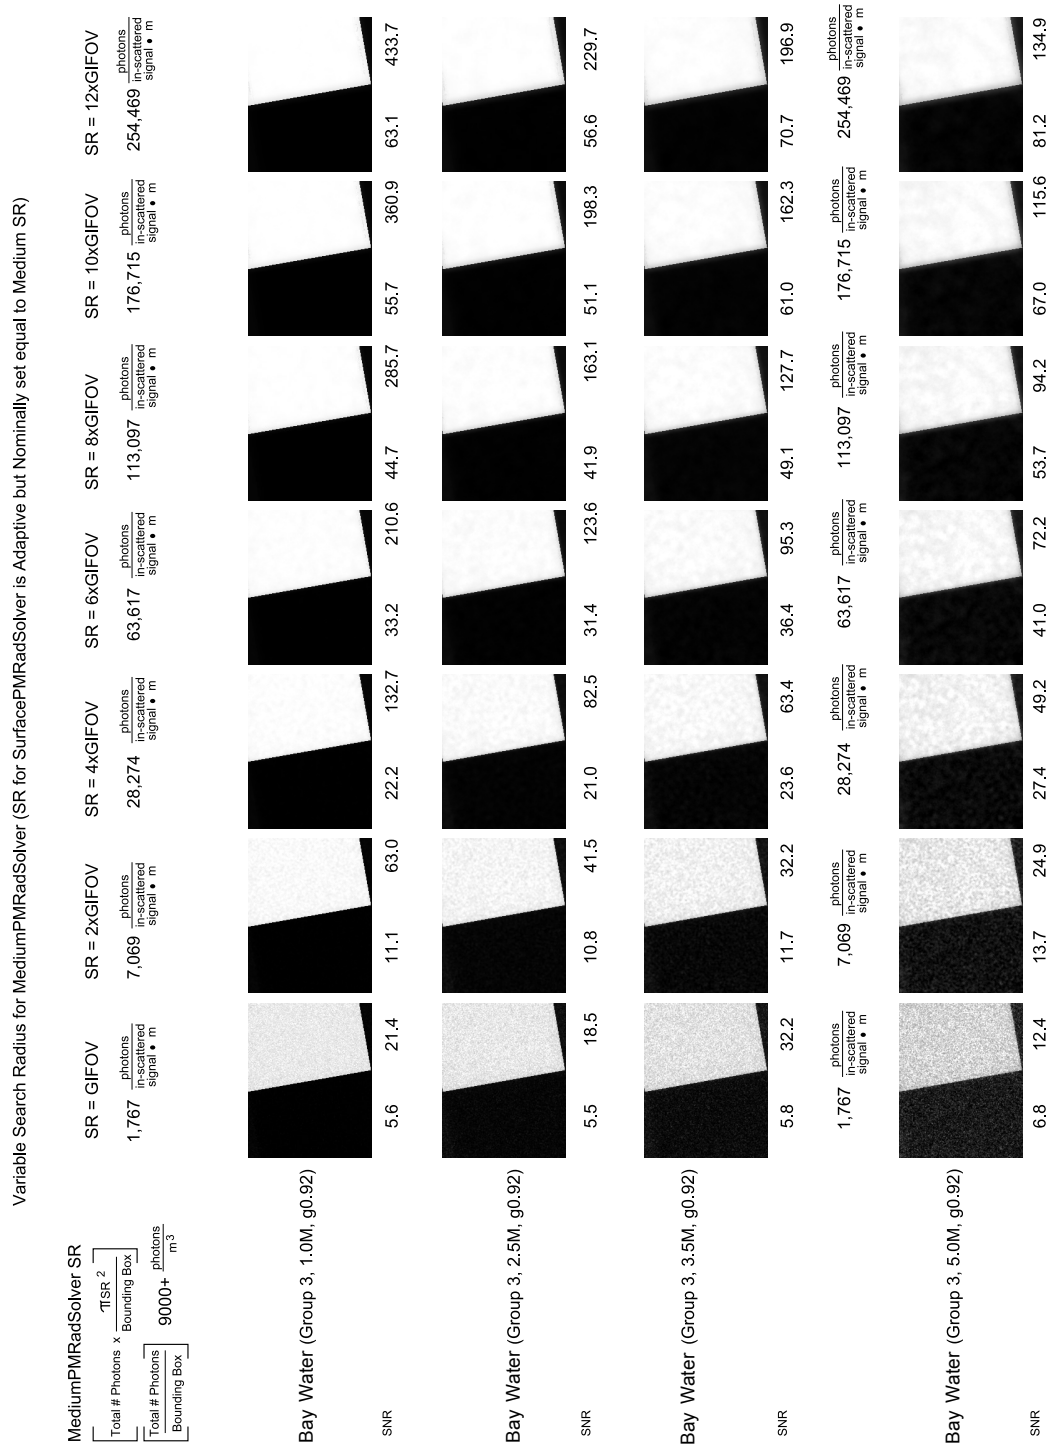


Figure 9.44: See Section 9.15 for Figure caption.

MediumPMRadSolver SR

$$\left[\begin{array}{l} \text{Total \# Photons} \times \frac{\pi \text{SR}^2}{\text{Bounding Box}} \\ \text{Total \# Photons} \left[\frac{\text{Bounding Box}}{3000+} \right] \end{array} \right] \frac{\text{photons}}{\text{m}^3}$$

SR = GIFOV	SR = 2xGIFOV	SR = 4xGIFOV	SR = 6xGIFOV	SR = 8xGIFOV	SR = 10xGIFOV	SR = 12xGIFOV
589	2,356	9,425	21,206	37,699	58,905	84,823
$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$
0.4	0.9	1.7	2.5	3.4	4.4	5.5
1.4	2.8	5.4	8.0	10.5	13.2	15.9
0.6	1.3	2.6	3.9	5.1	6.3	7.7
0.7	1.5	3.1	4.7	6.4	8.1	9.8
0.8	1.7	3.4	4.8	6.0	7.2	8.4
0.8	1.5	2.9	4.3	5.7	7.0	8.1
589	2,356	9,425	21,206	37,699	58,905	84,823
$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$	$\frac{\text{photons}}{\text{in-scattered signal} \bullet \text{m}}$
0.8	1.7	3.4	5.1	6.8	8.5	10.3
0.9	1.7	3.4	5.1	6.4	7.8	9.3

Turbid Water (Group 4b, 1.0M, P)

Turbid Water (Group 4b, 2.5M, P)

Turbid Water (Group 4b, 3.5M, P)

Turbid Water (Group 4b, 5.0M, P)

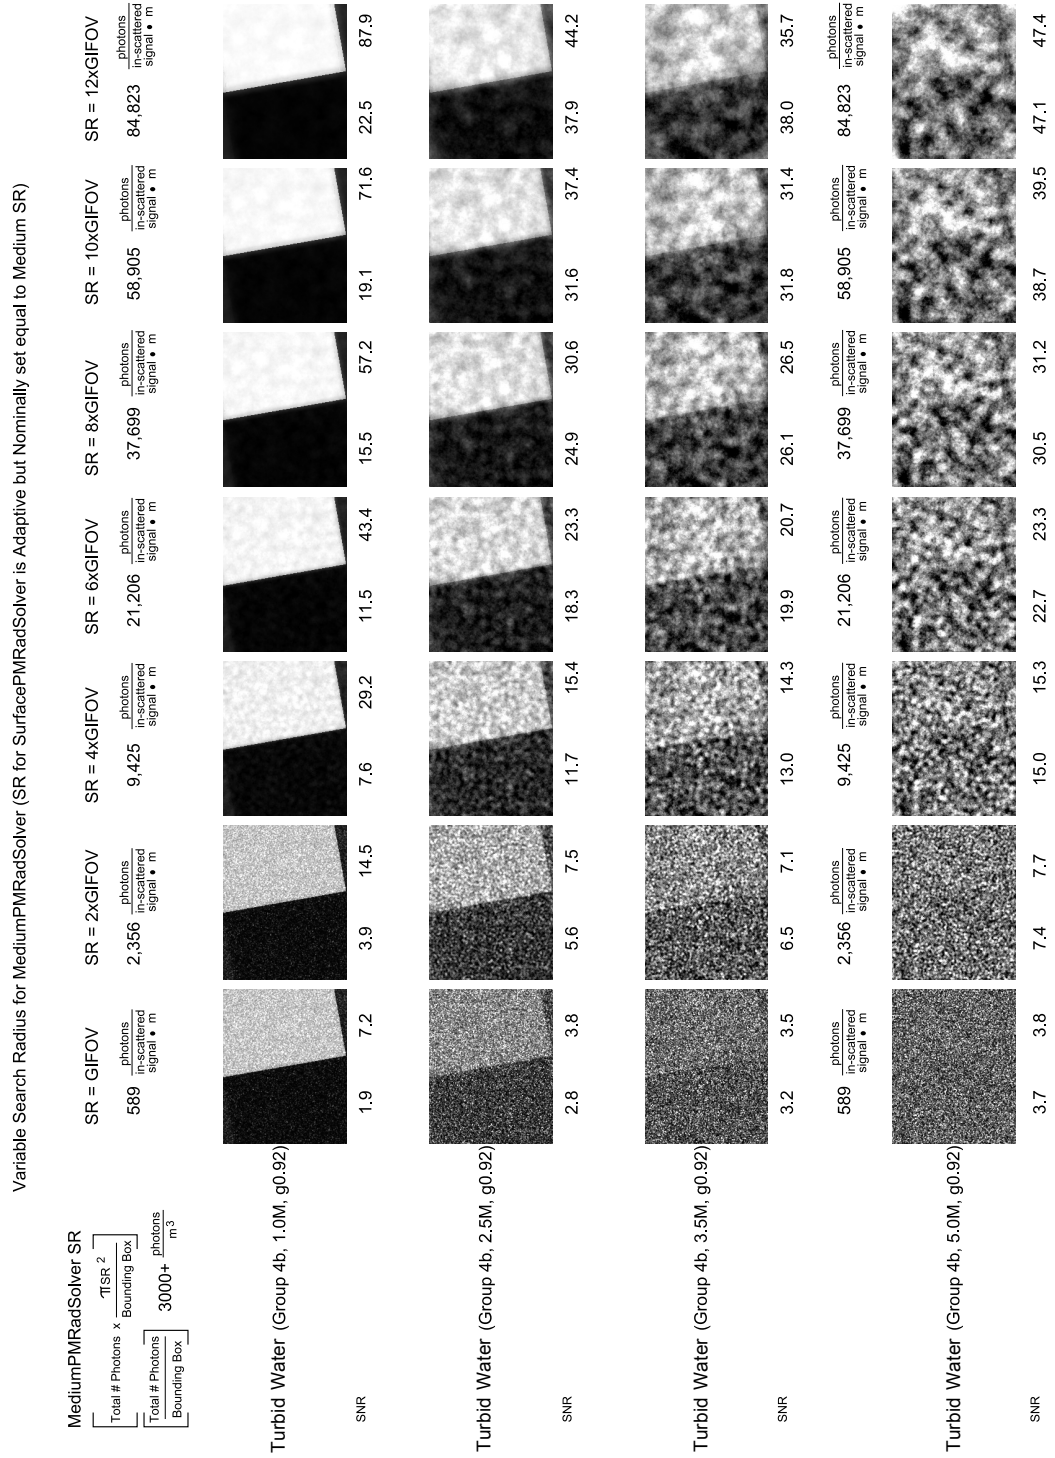


Figure 9.46: See Section 9.15 for Figure caption.

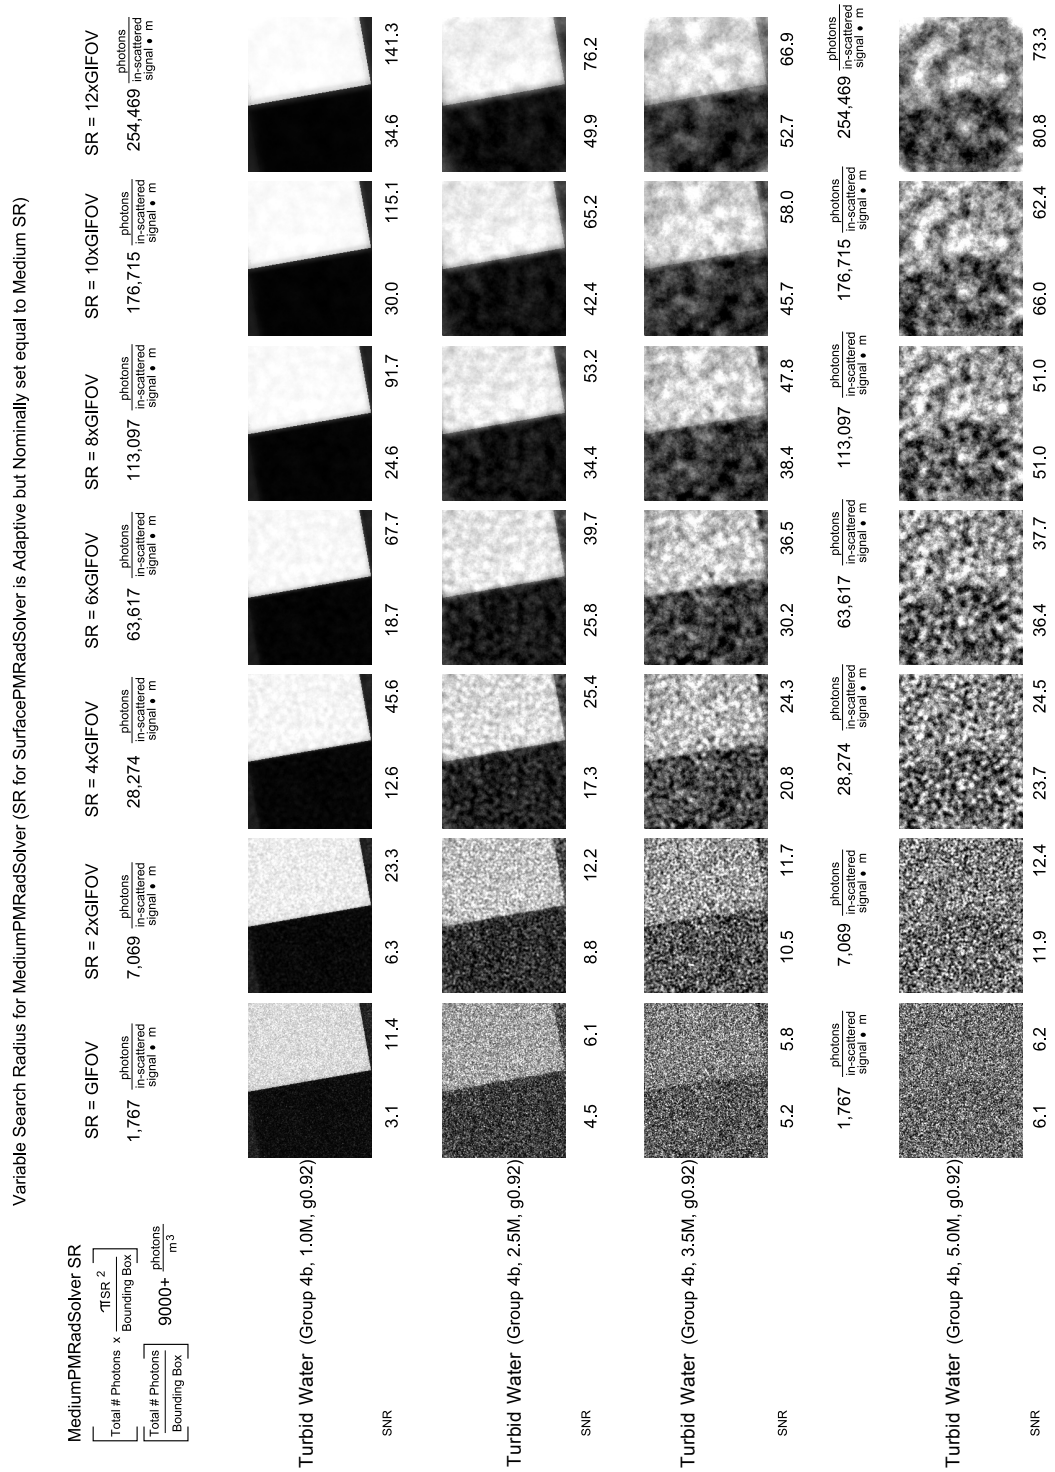


Figure 9.47: See Section 9.15 for Figure caption.

9.16 Collection of Spectral Images

The following caption is applicable to all Figures shown in this section:

- **Caption**

Each row represents a fixed target depth from the surface (*e.g.* 1.0, 2.5 and 3.5m) (top to bottom). Columns are paired. In each pair, the left-most column depicts the 3-band simulation (labeled “3 Band Simulation”), and the right-most column depicts the 3-single band simulations that have been combined (labeled “3 Band Combination”). There are a total of 4 pairs, with search radii from 0.25, 1.0, 2.0 and 3.0m (left to right). Mean photon density is shown at the top of the Figure for both the volume and surface. For the 3-band simulation, this is the *total mean photon density*. For the 3-single band simulations that have been combined, the reported mean density is on a *per band basis*. The left-most column reports the phase function (g = average cosine of one-parameter Henyey-Greenstein phase function).




Figure 9.49: See Section 9.16 for Figure caption.




Figure 9.50: See Section 9.16 for Figure caption.




Figure 9.51: See Section 9.16 for Figure caption.




Figure 9.52: See Section 9.16 for Figure caption.

Chapter 10

Visual Demonstrations

This chapter is divided into two main sections. The first Section (10.1) describes the systematic creation of DIRSIG imagery that depicts a coastal scene with constant underlying geometry and surface material properties. The goal is to qualitatively evaluate photon mapped imagery with variable source-sensor, wind speed and IOP configurations. The second Section (10.4) contains visual demonstrations of miscellaneous (but interesting) phenomenology, including shadows and caustics, which are phenomena most likely to be absent from the coastal scene by virtue of the low reflectance of in-water objects, volume turbidity, and in some cases, the sensor vantage point.

10.1 Coastal Scene

A visual demonstration of the phenomenology associated with a photon mapped coastal scene was created, taking inspiration from the image shown in Figure 10.1 (a.). This base image was classified into six separate materials, and the final classification image was used to map separate materials onto the base geometry. Next, each color channel of the original image was adjusted to remove any piece of 3D geometry that was destined to become a facetized surface in the final image, including the breaking waves near the shore line. The end result is shown in Figure 10.1 (b.), (c.) and (d.). The touched-up R, G and B channels were then manipulated to remove wind-induced surface wavelets, after which a statistical adjustment was carried out to map each class mean and variance to a desired value, thereby generating the final texture maps. Last but not least, in an effort to create a reasonable amount of scene clutter, shadow and character, several additional pieces of geometry were added to embellish the image. When appropriate, texture was added to each embellishment (*e.g.* wood grain, cement, etc.).

A total of ten different simulation configurations were prepared. Table 10.1 itemizes these simulation variations concerning IOPs, wind speed and sensor-source orientation. Simulation *A*, *B* and *C* represent three sensor-source configurations, characterized by a




Figure 10.1: Original imagery (a.) and touched-up red (b.), green (c.) and blue (d.) channels.

fixed solar orientation, but adjusting the sensor orientation to generate a simulation where the sensor captures backscattered/reflected sunlight (*A*), specular or forward scattered sunlight (assuming a horizontal planar surface) (*B*), and finally, with the sensor in a nadir-viewing position (*C*). Simulation *D* is a variant of *C* which positions the sun at a larger zenith angle (*e.g.* the source in simulation *D* is 30° greater than in simulation *C*). Simulations *E*, *F* and *G* are meant to accentuate the details of reflections from a wind-roughened interface, including a change in wind speed from 6.0 to $10.0 \frac{m}{s}$. Details regarding the facetized air-water interface are shown in Table 10.2. Simulation *H* places the sensor at the low elevation angle of 10° , meant to mimic the vantage point of someone standing aboard a floating vessel (or a “rowboat-line” view angle). Finally, simulations *I* and *J*, when combined with simulation *D*, depict three constituent concentration configurations (*Moderate*, *High* and *Low*, respectively). For each IOP configuration, the spectral character has been held constant, while the total concentration has been varied to generate the coefficients shown in Table 10.3. All scattering coefficients are modeled by the single parameter Henyey-Greenstein phase function with $g = 0.88$.

Table 10.1: Simulation configurations details.

Simulation	Description	Solar Configuration	Sensor Configuration	IOPs	Wind Speed
A	Backscatter	$\theta_z = 30, \phi = 315$	$\theta_e = 60, \phi = 315, d = 500\text{m}$	Low	$6.0 \frac{m}{s}$
B	Specular	$\theta_z = 30, \phi = 315$	$\theta_e = 60, \phi = 135, d = 500\text{m}$	Low	$6.0 \frac{m}{s}$
C	Nadir	$\theta_z = 30, \phi = 315$	$\theta_e = 90, \phi = 270, d = 500\text{m}$	Low	$6.0 \frac{m}{s}$
D	Nadir	$\theta_z = 60, \phi = 315$	$\theta_e = 90, \phi = 270, d = 500\text{m}$	Low	$6.0 \frac{m}{s}$
E	Surface Detail	$\theta_z = 30, \phi = 315$	$\theta_e = 60, \phi = 180, d = 500\text{m}$	Low	$6.0 \frac{m}{s}$
F	Surface Detail	$\theta_z = 60, \phi = 315$	$\theta_e = 30, \phi = 180, d = 500\text{m}$	Low	$6.0 \frac{m}{s}$
G	Wind Speed Variation	$\theta_z = 30, \phi = 315$	$\theta_e = 60, \phi = 180, d = 500\text{m}$	Low	$10.0 \frac{m}{s}$
H	Rowboat	$\theta_z = 48, \phi = 315$	$\theta_e = 10, \phi = 262, d = 90\text{m}$	Low	$6.0 \frac{m}{s}$
I	Concentration Variation	$\theta_z = 60, \phi = 315$	$\theta_e = 90, \phi = 270, d = 500\text{m}$	Mod	$6.0 \frac{m}{s}$
J	Concentration Variation	$\theta_z = 60, \phi = 315$	$\theta_e = 90, \phi = 270, d = 500\text{m}$	High	$6.0 \frac{m}{s}$

Concentration: Mod = Moderate; actual coefficients are shown in Table 10.3, z = zenith, e = elevation.

All simulations have a GSD of 0.2m (along the optical axis), except for the following variants:

D , I and J (GSD = 0.1m) and H (GSD = 0.018m).

Table 10.2: Beaufort scale and attributes of facetized wave surfaces. The first two columns describe the wind and associated waves based on the Beaufort Scale. The last three columns report the predicted and actual RMS wave height wave slope, and correlation length for the simulated surface.

Wind Speed $\frac{m}{s}$ (mph)	Wave Description	RMS Wave Height (m) (predicted, actual)	RMS Wave Slope (radians) (predicted, actual)	Correlation length (m)
6.0 (13.4)	Small Waves	0.576, 0.502	0.0337, 0.029	24.15
10.0 (22.4)	Moderate Waves	1.600, 1.24	0.0542, 0.042	41.7

Additional IDL code simulation details: GSD = 0.25m, 1200 x 1200 pixels, seed = 500, and modeled using 30,000 triangles per surface.

Each material has been assigned the Classic Radiometry Solver, except for the items listed in Table 10.4. The material attributes for each surface have been borrowed from former DIRSIG material files (*e.g.* vegetation, sand, cement, shingles, fiberglass, etc.). Every simulation traces 50 million surface photons, and 70 million volume photons. The bounding box for each map assumes the entire scene, which is approximately 90 x 60 meters. This generates average photon densities of approximately $5500 \frac{\text{photons}}{m^2}$ for the surface map, and $3000 \frac{\text{photons}}{m^3}$ in the volume map. Since the volume is only about half the field of view, this average density will be nearly double the reported value. In a similar manner, the surface map excludes the grass and trees in the top quadrant of the image, which will make the average photon density of this map about a quarter higher than the reported value.

Table 10.3: Coefficients used for visual simulations.

Description	Coefficient $\frac{1}{m}$	0.4 microns	0.5 microns	0.6 microns
Low	Chlorophyll σ_a	0.001862	0.001469	0.000401
	Chlorophyll σ_s	0.005185	0.005652	0.005412
	TSS σ_a	0.011478	0.003963	0.001528
	TSS σ_s	0.079333	0.063466	0.052888
	CDOM σ_a	0.007222	0.001472	0.000184
Mod	Chlorophyll σ_a	0.009311	0.007348	0.002007
	Chlorophyll σ_s	0.025929	0.028264	0.027063
	TSS σ_a	0.057391	0.019817	0.007640
	TSS σ_s	0.396665	0.317334	0.264444
	CDOM σ_a	0.036110	0.007360	0.000920
High	Chlorophyll σ_a	0.018623	0.014696	0.004015
	Chlorophyll σ_s	0.051859	0.056527	0.054127
	TSS σ_a	0.114783	0.039635	0.015281
	TSS σ_s	0.793330	0.634667	0.528889
	CDOM σ_a	0.072220	0.014720	0.001840

Mod = Moderate.

Table 10.4: The Classic Radiometry Solver has been assigned to the majority of scene materials, with the exception of those listed below.

Radiometry Solver	Item/Material
SurfacePM Radiometry Solver	Pier Step functions Sand Underwater vegetation
Generic Radiometry Solver	Aluminum & fiberglass on boat
Fresnel Radiometry Solver	Air-water interface
MediumPM Radiometry Solver	Water volume

10.2 Results

Figure 10.2 depicts two photon mapped user-configurations for simulation *A*. The images in the top row were simulated using a search radius of 1.0m (top left), and 0.4m (top right). The larger SR is associated with a more mottled appearance. Based on the author’s subjective assessment, the image at right was deemed more visually pleasing. Since the simulation does not include oversampling, the pixelation artifacts were softened by convolution with a *notional* sensor response. Two convolution kernels were tested as shown in Figure 10.2; the SINC and a simple Gaussian function. Several gamma adjustments were also considered, but $\gamma = 0.65$ and convolution with the Gaussian kernel (Figure 10.2 bottom row, FWHM = 1 pixel) was subjectively determined to be the most visually pleasing option, and will therefore be applied to all other images.

Aside from using simulation *A* to test and compare post-processing options, the sensor-source configuration in this simulation was meant to capture backscattered/reflected solar illumination. Positioning the sensor directly in front of the source typically generates a hotspot, and an image largely devoid of shadows. This is clearly evident in Figure 10.2, as well as the larger version of simulation *A* that can be found in section 10.3 Figure 10.9.

In comparison, simulation *B* uses a sensor-source configuration that captures forward scatter and specular glints. These details are captured by placing the sensor at a zenith angle that is equal to that of the sun, but rotated 180° in azimuth. The resulting imagery is expected to have notable shadows, halos in illumination associated with transmissive surfaces, and for wind-roughened water, wave glints. The glints are evident in Figure 10.3 (convolved with the Gaussian function shown in Figure 10.2, and with a gamma of 0.65). A larger version of simulation *B* can be found in section 10.3, Figure 10.10. Although the specular glints are present, they are somewhat artificial looking. This is believed to be a function of the facetized nature of the air-water interface (despite the fact that this surface does not appear pixelated in any preliminary simulation), as well as the absence of any oversampling. Of additional interest are the shadows cast by the small boat (upper left corner of image) and the cement pier (lower right corner of image). The shadow cast by the boat is symmetric. By virtue of the geometry of the shadow, it must be a shadow that has been cast by the boat onto the air-water interface. In contrast, the shadow cast by the pier is asymmetric, becoming elongated at deeper water depths. Barring any drastic fluctuations in the shadow as a function of the waves on the air-water interface, this shadow appears to be part surface (air-water and ground plane), and part volumetric.


Simulation *C* is the nadir viewing angle for the source configuration depicted in simulation *A* and *B* (*e.g.* sensor zenith angle $\theta_z = 30$ and azimuth angle $\phi = 315$). Simulation *D* is a repeat of simulation *C*, except for a larger source zenith angle which is expected to cast much longer shadows. This is shown in Figure 10.4. Larger versions of simulation *C* and *D* are available for inspection in section 10.3, Figure 10.11 and 10.12, respectively. These images also clearly show the change in sensor reaching radiance associated with the submerged test targets. The step targets are positioned so that each is approximately 0.75m deeper than its neighbor. The more shallow the target, the more “white” the left edge of the step function, and the sharper the expected edge. For low constituent concentrations (and sensor resolution), the edges of all step targets, regardless of depth, appear equally well resolved. This is in contrast to the appearance of the step function when the volume is associated with higher constituent concentrations as shown in Figure 10.8, which depicts the results of scenario *I* and *J*.




SR = 1.0m




SR = 0.4m




SINC PSF



Gaussian PSF



SR = 0.4M with SINC & gamma = 0.65




SR = 0.4m with Gauss & gamma = 0.65

Figure 10.2: Top row: Backscatter sensor configuration (Scenario A) with SR = 1.0m and SR = 0.4m. The less mottled image at right was deemed more visually pleasing. Middle Row: SINC and Gaussian PSFs meant to mimic a notional sensor response function (X and Y axes are in pixels). Bottom row: The SR = 0.4m image (gamma = 0.65) convolved with the SINC and Gaussian PSF. Both PSFs have 1 pixel wide FWHMs. The kernel values near the central peak for the Gauss are 0.2909, 1.0000, 0.2909, and for the SINC are -3.8981e-017, 1.0000, -3.8981e-017.



(A) Backscatter



(B) Forward Scatter/Specular

Figure 10.3: Specular sensor configuration (Scenario *B*) (shown along side the backscatter simulation configuration *A*), with $SR = 0.4m$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2. The glints appear somewhat artificial, but are believed to be a function of the 30,000 triangular facets used to construct the air-water interface, and the absence of any oversampling during the second pass.




(C) Nadir, solar $\theta_z = 30^\circ$




(D) Nadir, solar $\theta_z = 60^\circ$

Figure 10.4: Nadir viewing sensor with source zenith angle of $\theta_z = 30^\circ$ (*C*) and $\theta_z = 60^\circ$ (*D*). Other variables are as follows: $SR = 0.4m$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2.

Simulation *E* and *F* position the sensor at $\theta = 90^\circ$ and $\phi = 135^\circ$ from the source, but with a source zenith angle of $\theta_z = 30^\circ$ and $\theta_z = 60^\circ$, respectively. The higher solar zenith angle in simulation *F* will generate longer shadows, as shown in Figure 10.5. To maintain a constant sensor-source angle of $\theta = 90^\circ$, the sensor elevation angle has been lowered to 30° in simulation *F*. This observation angle captures incident illumination at considerable grazing angles, associated with a higher Fresnel reflectance, thereby generating imagery that accentuates the details associated with the air-water interface, showing how difficult it is to “see through” the interface when the sensor is at a low grazing angle. A larger version of both simulation *E* and *F* can be viewed in Section 10.3, Figure 10.13 and 10.14, respectively.



(E) $\Delta_\theta = 90^\circ$ $\Delta_\phi = 135^\circ$, solar $\theta_z = 30^\circ$




(F) $\Delta_\theta = 90^\circ$ $\Delta_\phi = 135^\circ$, $\theta_z = 60^\circ$


Figure 10.5: Simulation *E* and *F* with $SR = 0.4m$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2.

Simulation *G* is a repeat of simulation *E*, but for a higher wind speed of $10.0 \frac{m}{s}$. The higher wind speed is expected to generate a surface with larger amplitude waves, as shown in Figure 10.6 (and section 10.3 Figure 10.15).

Simulation *H* positions the sensor from the vantage point of someone standing aboard a floating vessel (or a “rowboat-line” view angle). This configuration will show mirror-like reflections from the Fresnel interface, as well as a side-view through the water volume. This view will be the most reasonable attempt to capture the volumetric shadow cast by the deepest step function. The photon-mapped and non-photon mapped counterpart for



(E) Wind speed = $6.0 \frac{m}{s}$




(G) Wind speed = $10.0 \frac{m}{s}$


Figure 10.6: Simulation E (wind speed = $6.0 \frac{m}{s}$) and G (wind speed = $10.0 \frac{m}{s}$) with $SR = 0.4m$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2.

this simulation scenario are shown in Figure 10.7. A larger image of simulation H can be viewed in section 10.3, Figure 10.16. The surface reflections are clearly prominent, despite the lack of any volumetric shadows. This is not a deficit of the DIRSIG model, but instead a function of the simulation configuration. Proof of volumetric shadows were formerly shown in Chapter 8 via irradiance profiles, and a separate visual demonstration is included in section 10.4.

Scenarios I and J are variations in constituent concentrations that can be compared against scenario D . All three simulations have a constant source-sensor orientation, but parameter concentrations that vary between moderate, high, and low, respectively. The moderate and high turbidity scenarios are shown in Figure 10.8, and larger versions in section 10.3, Figures 10.17 and 10.18, respectively. Obviously, the higher the turbidity, the greater the extinction. This means the sensor will be less able to “see” through the water column and capture details associated with the benthic ground plane. This will also be apparent from inspection of the submerged step targets. The actual edge between the bright and dark portions of the target will be blurred, as will the edges between the target and underlying benthic materials.



(H) Non-photon mapped



(H) Low Elevation Angle $\theta_e = 10^\circ$

Figure 10.7: Photon mapped and non-photon mapped versions of simulation H , with $SR = 0.4m$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2. Note the prominent surface reflections.



(I) Moderate Concentration.



(J) High Concentration.

Figure 10.8: Simulations I and J with moderate and high turbidity, respectively, $SR = 0.4m$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2. Compare with simulation D in Figure 10.4 which has a lower turbidity than both I and J .

10.3 Image Summary

This section presents larger versions of all previously shown photon mapped images of the coastal scene.




Figure 10.9: Backscatter sensor configuration (Scenario *A*), with $SR = 0.4m$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2.



Figure 10.10: Specular sensor configuration (Scenario *B*), with $SR = 0.4m$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2.




Figure 10.11: Nadir viewing angle (Scenario *C*), with $SR = 0.4m$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2.




Figure 10.12: Nadir viewing angle but with larger source zenith angle (simulation D) than that depicted in Figure 10.11. Other variables are as follows: $SR = 0.4m$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2. Compare with simulation I (moderate turbidity, Figure 10.17) and J (high turbidity, Figure 10.18).




Figure 10.13: Simulation E which places the sensor $\theta = 90^\circ$ from the source, but with an azimuthal difference of $\phi = 135^\circ$. Other variables are as follows: $SR = 0.4\text{m}$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2.




Figure 10.14: Simulation F which places the sensor $\theta = 90^\circ$ from the source, but with an azimuthal difference of $\phi = 135^\circ$, and a larger solar zenith angle than that depicted in Figure 10.13. Other variables are as follows: $SR = 0.4\text{m}$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2.




Figure 10.15: Simulation G (same as simulation E shown in Figure 10.13 except for the larger wind speed of $10.0 \frac{m}{s}$). Other variables are as follows: $SR = 0.4m$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2.




Figure 10.16: Simulation H . Other variables are as follows: $SR = 0.4\text{m}$, $\text{gamma} = 0.65$, and convolved with the Gaussian function shown in Figure 10.2.




Figure 10.17: Simulation I (moderate turbidity). Compare with simulation C (Figure 10.11) and J (Figure 10.18) which have low and high turbidity, respectively. Other variables are as follows: $SR = 0.4m$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2.



Figure 10.18: Simulation J (high turbidity). Compare with simulation C (Figure 10.11) and I (Figure 10.17) which have low and moderate turbidity, respectively. Other variables are as follows: $SR = 0.4m$, $\gamma = 0.65$, and convolved with the Gaussian function shown in Figure 10.2.

10.4 Miscellaneous Visual Demonstrations

10.4.1 Surface & Volumetric Shadows

The preceding visual simulations do not fully illustrate two major phenomenon of interest; volumetric and surface shadows generated by floating obstructions. This is not a failure of the DIRSIG model, nor the simulation configuration, but instead is a product of the low reflectance of in-water objects, the volume IOPs, and the remotely sensed vantage point. Figures 10.19 and 10.20 illustrate both phenomenon, but from more contrived simulation conditions. Figure 10.19 was generated using a solar zenith angle of 45° with hemisphere irradiance of 5.0 units, with 90% attributed to the solar disc. The volume contained 5M photons (approximately $1400 \frac{\text{photons}}{\text{m}^3}$), and IOPs of $\sigma_a = \sigma_s = 0.2$ with a uniform scattering phase function. The numeric designations above each image indicate the units (in meters) between the bottom edge of the floating obstruction and the actual air-water interface. The surface shadow cast by the box (and onto the air-water interface) is apparent in the left most image, while the volumetric shadow is evident for all box positions. As the box is lowered through the depth of the water volume, the volumetric shadow becomes less pronounced, as expected. Figure 10.20 is from a slightly different vantage point (and different water IOPs), and depicts the shadow cast on the ground plane due to the presence of a floating obstruction overhead.




Figure 10.19: Interface and volumetric shadow cast by floating obstruction positioned at variable distances of 0.5, 1.5, 2.2, 3.0, 3.5 and 4.0m from the surface.

10.4.2 Caustics

A hand-held cube of epoxy generates interesting caustics that can be mimicked by DIRSIG. The numerous refractive events that generate the caustics are too complicated to explore in detail, but Figure 10.21 clearly illustrates DIRSIG's ability to model each refractive event and generate both the caustics highlighted by the blue arrows, as well as the caustics that are apparent in the interior of the cube.




Figure 10.20: Surface shadow cast onto ground plane by a floating obstruction overhead. Image has not been optimized to reduce spectral variance associated with the surface photon map.




Figure 10.21: DIRSIG simulated caustics (left), and those generated by an insulcast epoxy cube (right). The blue arrows denote corresponding caustics that are exterior to the cube. Additional corresponding caustics are also evident within the interior of the cube. Note that the image has not been optimized for noise reduction.

10.5 Chapter Summary

This chapter presents visual demonstrations that illustrate DIRSIG's ability to model interesting phenomena germane to coastal ocean-water simulations. It is meant as a capstone to all former chapters in which simulations were intentionally simplified in order to isolate and evaluate a limited number of interactions and contributions to sensor reaching radiance. However, not all phenomena of interest will always be apparent. For example, surface shadows cast onto the benthic ground plane by floating obstructions tend to be lost by virtue of the low reflectance of in-water objects, and scattering within the volume. These same surface shadows, but cast onto the air-water interface, tend to be very minor in horizontal extent due to the limited height of coastal objects, and desirable remote sensing conditions which usually include a high solar elevation angle in order to maximize the depth of penetration of sunlight within the volume of interest.

Caustics are another example of an interesting phenomenon, but one that is rare to observe from the remote sensing vantage point. The refractive focusing of light is nicely modeled by DIRSIG (for evidence, see Figure 10.21), but wave induced caustics are often absent from remotely sensed images owing to their high spatial and temporal frequency, sensor GSDs, volume turbidity, and the low reflectance of in-water objects. Therefore, contrived and sometimes simplistic simulation scenarios, such as those shown in the preceding chapters (and Section 10.4 of this chapter) are equally as valuable as realistic imagery, not only for illustrative purposes, but sometimes as the only means by which the phenomena of interest can be reasonably validated.

Chapter 11

Final Comments

It is proposed that the research presented in all previous chapters be taken as a reasonable validation of DIRSIG's ability to model both raytraced and photon mapped contributions to sensor reaching radiance in coastal environments. Although each investigation has been geared toward validation of in-water radiative transfer, several raytraced solutions are equally applicable to land observations, and several photon mapped solutions are equally applicable to other participating media, of which fog and clouds are obvious alternative examples.

Although the DIRSIG model can describe exceptionally complex physical interactions, the end product of a simulation is naturally limited by the state of the art in computational tools. As DIRSIG evolves, the computational tools become more sophisticated, as does the rendering. However, it is not reasonable to expect a linear improvement, especially when individual tools do not evolve at the same rate, and when small seemingly innocuous changes can sometimes generate unforeseen incompatibilities.

As a non-developer, it would be foolish to try to suggest how those charged with care of the code should best meet the mutual goals of code development and maintenance. Yet if asked to comment, it would seem that two obvious, although not mutually exclusive, schools of thought exist. The first is the repeated evaluation of a generalized suite of validations, which would be especially useful before major code changes are accepted. The second school of thought is less direct, but which seems more, if not equally, preferable. This involves generating sufficient interest in a modality such that it attracts a reasonable number of users with varied research goals. These users and their varied goals would naturally generate simulations that tax the code in unique ways, thereby allowing the research needs of the user-community to serve as a continual revalidation of the code as it evolves and changes over time. Hopefully the research presented here is a stepping stone for both modes of validation; providing a baseline reference against which future simulations can be expected to conform, and generating sufficient confidence in the existing tool to spark the interest of a user-community.

Bibliography

- Ashikhmin, M., S. Premoze, and P. Shirley (2000). A microfacet-based BRDF generator. *International Conference on Computer Graphics and Interactive Techniques, Proceedings of the 27th annual conference on Computer graphics and interactive techniques, SIGGRAPH: ACM Special Interest Group on Computer Graphics and Interactive Techniques*, pp. 65 – 74.
- Bannister, T. (1992). Model of the mean cosine of underwater radiance and estimation of underwater scalar irradiance. *Limnology and Oceanography Vol. 37, No. 4*, pp. 773–780.
- Berwald, J., D. Stramski, C. Mobley, and D. Kiefer (1995). Influences of absorption and scattering on vertical changes in the average cosine of the underwater light field. *Limnology and Oceanography Vol. 40, No. 8*, pp. 1347–1357.
- Carder, K., C. Liu, Z. Lee, D. English, J. Patten, F. Chen, and J. Ivey (2003). Illumination and turbidity effects on observing faceted bottom elements with uniform lambertian albedos. *Limnology and Oceanography Vol 48, No. 1, Part 2*, pp. 355 – 363.
- Cox, C. and W. Munk (1954). Measurement of the roughness of the sea surface from photographs of the sun’s glitter. *Journal of the Optical Soccity of America 44, no. 11*, 838 – 850.
- Cox, C. and W. Munk (1956). Slopes of the sea surface deduced from photographs of sun glitter. Technical report, University of California Press, Berkeley and Los Angeles, Bulletin of the Scripps Institution of Oceanography of the University of California.
- Doyle, J. and G. Zibordi (2002). Optical propagation within a three-dimensional shadowed atmosphere-ocean field: Application to large deployment structures. *Applied Optics Vol. 41, No. 21*, pp. 4283 –4306.
- Duntley, S. (1963). Light in the sea. *Journal of the Optical Society of America 53*, 214–233.
- Fleming, R., A. Torralba, and E. Adelson (2004). Specular reflections and the perception of shape. *Journal of Vision Vol. 4*, pp. 798–820.

- Gerstl, S., A. Zardecki, W. Unruh, D. Stupin, G. Stokes, and N. Elliott (1987). Off-axis multiple scattering of a laser beam in turbid media: Comparison of theory and experiment. *Applied Optics Vol. 26, No. 5*, pp. 779–785.
- Goodenough, A. (2007). *In-Water Radiative Transfer Modeling using Photon Mapping*. Ph. D. thesis, Rochester Institute of Technology.
- Goodenough, A. (2008). Personal communication.
- Goodenough, A. (2009). Personal communication.
- Gordon, H. (1985). Ship perturbation of irradiance measurements at sea 1: Monte carlo simulations. *Applied Optics Vol. 24, No. 23*, pp. 4172–4182.
- Gordon, H. and K. Ding (1992). Self-shading of in-water optical instruments. *Limnology and Oceanography Vol. 37, No. 3*, pp. 491–500.
- Hecht, E. (2002). *Optics*. Pearson, Addison Wesley.
- Hodara, H. (1973). Experimental results of small angle scattering. *AGARD Lecture Series, Optics of the Sea (Interface and In-water Transmission and Imaging) 61*, 3.4.1–3.4.17.
- Honey, R. (1979). Beam spread and point spread functions and their measurement in the ocean. *Ocean Optics VI. October 23-25, Monterey, CA Vol. 208*, pp. 242–248.
- Hou, W., D. Gray, A. Weidemann, and R. Arnone (2008). Comparison and validation of point spread models for imaging in natural waters. *Optics Express Vol. 16, No. 13*, pp. 9958–9965.
- Hou, W., D. Gray, A. Weidemann, G. Fournier, and J. Forand (2007b). Automated underwater image restoration and retrieval of related optical properties. *IEEE International, Geoscience and Remote Sensing Symposium, IGARSS 2007 N/A ISBN: 978-1-4244-1211-2*, pp. 1889–1892 [on-line PDF version appears to be missing a page].
- Hou, W., A. Weidemann, D. Gray, and G. Fournier (2007c). Imagery-derived modulation transfer function and its application for underwater imaging. *SPIE: Applications of Digital Image Processing XXX Vol. 6696*, pp. 669622–1–669622–8.
- Ishimaru, A. (1978). *Wave Propagation and Scattering in Random Media, Volume I and Volume II*. Academic Press, Inc.
- Jaffe, J. (1992). Validity and range of linear approximations in underwater imaging. *SPIE, Ocean Optics XI Vol. 1750*, pp. 388–396.
- Jaffe, J. (1995). Monte carlo modeling of underwater-image formation: Validity of the linear and small angle approximations. *Applied Optics Vol. 34, No. 24*, pp. 5413–5421.
- Jensen, H. W. (2001). *Realistic Image Synthesis Using Photon Mapping*. A. K. Peters.

- MacRobert, T. (1967). *Spherical Harmonics: An Elementary Treatise on Harmonic Functions with Applications*. Pergamon Press Ltd.
- Maffione, R. (1998). Theoretical developments of the optical properties of highly turbid waters and sea ice. *Limnology and Oceanography* Vol. 43, No. 1, 29–33.
- Marshak, A. and A. Davis (Eds.) (2005). *3D Radiative Transfer in Cloudy Atmospheres*. Springer.
- Martin, S. (2004). *An Introduction to Ocean Remote Sensing*. Cambridge University Press.
- McLean, J. and K. Voss (1991). Point spread function in ocean water: Comparison between theory and experiment. *Applied Optics* Vol. 30, No. 15, pp. 2027–2030.
- Mertens, L. and F. J. Replogle (1977). Use of point spread and beam spread functions for analysis of imaging systems in water. *Journal of the Optical Society of America* 67, No. 8, 1105–1117.
- Mobley, C. (1994). *Light and Water: Radiative Transfer in Natural Waters*. Academic Press.
- Mobley, C., B. Gentili, H. Gordon, Z. Jin, G. Kattawar, A. Morel, P. Reinersman, K. Stamnes, and R. Stavn (1993). Comparison of numerical models for computing underwater light fields. *Applied Optics* 32, no. 36, 7484 – 7504.
- Mobley, C., L. Sundman, and E. Boss (2002). Phase function effects on oceanic light fields. *Applied Optics* Vol. 41, No. 6, pp. 1035–1050.
- Piening, B. and N. McCormick (2003). Asymptotic optical depths in source-free ocean waters. *Applied Optics* Vol. 42, No. 27, pp. 5382–5387.
- Piskozub, J. (2004a). Effect of 3-D instrument casing shape on the self-shading of in-water upwelling irradiance. *Optics Express* Vol. 12, No. 14, pp. 3144–3148.
- Piskozub, J. (2004b). Effect of ship shadow on in-water irradiance measurements. *Oceanologia* Vol. 46, No. 1, pp. 103–112.
- Raqueno, R. (1999). <http://wiki.cis.rit.edu/bin/viewauth/dirs/hydrolightinputfiles>.
- Reinersman, P. and K. Carder (2004). Hybrid numerical method for solution of the radiative transfer equation in one, two or three dimensions. *Applied Optics* 43, no. 13, 2734–2743.
- Ross, V., D. Dion, and G. Potvin (2005). Detailed analytical approach to the gaussian surface bidirectional reflectance distribution function specular component applied to the sea surface. *Journal of the Optical Society of America, A*. Vol. 22, No. 11, pp. 2442 – 2453.

- Salinas, S. and S. Liew (2007). Light reflection from a rough liquid surface including wind-wave effects in a scattering atmosphere. *Journal of Quantitative Spectroscopy & Radiative Transfer* 105, pp. 414–424.
- Schenck, H. J. (1957). On the focusing of sunlight by ocean waves. *Journal of the Optical Society of America* 47, no. 7, 653 – 657.
- Schjoth, L., J. Frisvad, K. Erleben, and S. J. (2007, December 1-4). Photon differentials. pp. 179–186. Association for Computing Machinery, Inc. Graphite 2007 - Perth, Western Australia.
- Schott, J. R. (2007). *Remote Sensing: The Image Chain Approach, 2nd Edition*. Oxford University Press.
- Schott, J. R., S. D. Brown, R. V. Raqueño, H. N. Gross, and G. Robinson (1999). An advanced synthetic image generation model and its application to multi/hyperspectral algorithm development. *Canadian Journal of Remote Sensing* 25(2), 99–111.
- Shu, F. (1991). *The Physics of Astrophysics: Volume I, Radiation*. University Science Books.
- Silverman, B. (1998). *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, CRC.
- Spinrad, R. and E. Widder (1992). Ship shadow measurements obtained from a manned submersible. *SPIE, Ocean Optics XI Vol. 1750*, pp. 372–382.
- Suykens-De Laet, F. (2002). *On Robust Monte Carlo Algorithms for Multi-Pass Global Illumination*. Ph. D. thesis, Katholieke Universiteit Leuven.
- Tam, W. and A. Zardecki (1980). Off-axis propagation of a laser beam in low visibility weather conditions. *Applied Optics Vol. 19, No. 16*, pp. 2822–2827.
- Thomas, G. and K. Stamnes (1999). *Radiative Transfer in the Atmosphere and Ocean*. Cambridge University Press.
- Tzannes, A. and J. Mooney (1995). Measurement of the modulation transfer function of infrared cameras. *Optical Engineering Vol. 34, No. 6*, pp. 1808–1817.
- Voss, K. (1991a). Simple empirical model of oceanic point spread function. *Applied Optics Vol. 30, No. 18*, pp. 2647–2651.
- Voss, K. (1991b). Variation of the point spread function in the sargasso sea. *SPIE Underwater Imaging, Photography, and Visibility Vol. 1537*, 97–103.
- Walter, B., P. Hubbard, P. Shirley, and D. Greenberg (1997). Global illumination using local linear density estimation. *ACM Transactions on Graphics Vol. 16, No. 3*, pp. 217–259.
- Walter, B., S. Marschner, L. Hongson, and K. Torrance (2007). Microfacet models for refraction through rough surfaces. In *Eurographics Symposium on Rendering*.

- Weir, C., D. Siegel, A. Michaels, and D. Menzies (1994). In situ evaluation of a ship's shadow. *SPIE Ocean Optics XII Vol. 2258*, pp. 815–821.
- Wells, W. (1969). Loss of resolution in water as a result of multiple small angle scattering. *Journal of the Optical Society of America Vol. 59, No. 6*, 686–691.
- Wells, W. (1973). Theory of small angle scattering. *AGARD Lecture Series, Optics of the Sea (Interface and In-water Transmission and Imaging) 61*, 3.3.1–3.3.19.
- Zaneveld, R. (1989). An asymptotic closure theory for irradiance in the sea and its inversion to obtain the inherent optical properties. *Limnology and Oceanography Vol. 34, No. 8*, pp. 1442–1452.
- Zege, E., A. Ivanov, and I. Katsev (1991). *Image Transfer Through a Scattering Medium*. Springer-Verlag.
- Zege, E., I. Katsev, and I. Polonsky (1993). Multicomponent approach to light propagation in clouds and mist. *Applied Optics Vol. 32, No. 15*, 2803–2812.
- Zeisse, C. (1995). Radiance of the ocean horizon. *Journal of the Optical Society of America, A. Vol. 12, No. 9*, pp. 2022 – 2030.
- Zender, C. (2008). *Radiative Transfer in the Earth System*. Freely Available Community Test (FACT) Project.
- Zhai, P., G. Kattawar, and P. Yang (2008). Impulse response solution to the three-dimensional vector radiative transfer equation in atmosphere-ocean systems. II. the hybrid matrix operator - monte carlo method. *Applied Optics 47, No. 8*, 1063–1071.

Appendices

Appendix A

Two-Stream Solution & Related Variables

A.1 The Average Cosine of the Radiance Field

The full-range mean cosine of the radiance distribution is:

$$\bar{\mu}(z) = \frac{E_d(z) - E_u(z)}{E_{od}(z) + E_{ou}(z)} \quad (\text{A.1})$$

where z is depth, E_d and E_u are the downwelling and upwelling irradiance (all positive quantities) and E_{od} and E_{ou} are the downwelling and upwelling scalar irradiance (Bannister 1992), as shown below (Maffione 1998).

$$E_o(z) = 2\pi \int_0^\pi L(z, \theta) \sin \theta d\theta \quad (\text{A.2})$$

Since neither underwater irradiance nor scalar irradiance are available *a priori*, actual calculation of the specific intensity-weighted angular mean may not be feasible. If the specific intensity field is truly isotropic, then $\bar{\mu} = \frac{1}{2}$ at all depths and both hemispheres. If the specific intensity were nearly linear in μ (*e.g.* $L(\mu) \approx C\mu$ where C is a constant), then $\bar{\mu} = \frac{2}{3}$. Often the choice for $\bar{\mu}$ is determined by trial and error (Thomas and Stamnes 1999), although the following discussion will shed some light on how to thoughtfully interpret the average cosine variable.

Berwald et al. (1995) investigated how absorption, scattering and the nature of the phase function influence the average cosine of the in-water radiance field as a function of vertical depth through the water column. They postulate that the average cosine as a function of depth can be described according to:

$$\bar{\mu}(\tau) = \bar{\mu}_\infty + (\bar{\mu}_o - \bar{\mu}_\infty) \exp(-P_\tau \tau) \quad (\text{A.3})$$

where $\bar{\mu}_o$ is the average cosine just below the water surface, $\bar{\mu}_\infty$ is the asymptotic value of the average cosine, and P_τ is the rate of change of the average cosine. The value of $\bar{\mu}_o$ and $\bar{\mu}_\infty$ are a function of the single scattering albedo and the phase function for the medium under investigation. The assumption is that the rate is depth independent, although Berwald et al. (1995) reports that it may vary within a few optical depths of the air-water interface.

Berwald et al. (1995) indicates that at very low w_o values, absorption dominates and the average cosine is nearly 1 at all depths (assuming a point source at zenith) or it at least will approach a value of 1 at greater optical depths when the surface illumination conditions have less impact. This is because highly angular light, traveling longer path lengths, will be more rapidly removed by absorption than nearly vertical light. As such, the radiance field is more dominated by vertical forward propagating light that is able to penetrate to greater depths in the absence of scattering. Conversely, Berwald et al. (1995) reports that for water that can be described by a single scattering albedo greater than 0.1, the variable controlling the exponential decay is dominated by the scattering process (as compared to the absorption process). The scattering phase function will then influence how the scattered light is distributed. The forward peaked scattering phase function associated with water is expected to induce large average cosine values for the radiance field, at least at small optical depths, before becoming less peaked at greater optical depths. The general trend in average cosine of the radiance distribution as a function of optical depth and variations in the single scattering albedo is provided by Berwald et al. (1995) and illustrated in Figure A.1 (left), for the Petzold scattering phase function and a point sun at zenith (in the absence of skylight).

A similar examination of the change in the average cosine was conducted by Bannister (1992), but this time varying the sun's zenith angle [see Figure A.1 (middle), absorption coefficient of 0.2 m^{-1} and a $\sigma_s : \sigma_a$ ratio of 4]. The shape of the average cosine of the radiance field with depth for larger zenith angles shows “an initial downward course followed by an upward rise toward the asymptotic limit” (Bannister 1992). Bannister (1992) describes the shape as a *turning* of the “asymmetric distribution near the surface to a vertically symmetric distribution in the asymptotic region”. The average cosine therefore has a more complicated rate function. The variation in the initial subsurface $\bar{\mu}_o$ is a function of the zenith angle, as well as the ratio of the scattering to absorption coefficient, as shown in Figure A.1 (right).

Berwald et al. (1995) goes on to summarize how the rate of change of the average cosine of the radiance field varies with contrived scattering phase functions (*e.g.* equal scattering probability between zero and β_{max} where $\beta_{max} = 5, 10, 20, 45$ and 180 degrees) and a zenith sun. Essentially, these phase functions appear as rectangular, uniform probability functions with the limits of the function imposed by zero and the value of β_{max} . In general, highly forward peaked phase functions produce smaller magnitude rates with little variation with the value of the single scattering albedo. Conversely, larger angle scattering




Figure A.1: Left: Variation in average cosine with depth for a Petzold scattering phase function, a constant point sun at zenith (*e.g.* directly overhead of the observation point), the absence of any downwelling skylight, for four single scattering albedos. Figure adapted from Berwald et al. (1995). Middle: Variation in average cosine with depth for various zenith angles, and using Petzold’s scattering phase function for the San Diego Harbor. Figure adapted from Bannister (1992). Right: Variation in average cosine with the scattering to absorption coefficient ratio. Figure adapted from Bannister (1992). Note that the Figures illustrate trends, and are not meant to represent absolute magnitudes.

produces a larger magnitude rate (and therefore the radiance distribution more quickly reaches the asymptotic value), also showing a “distinct increase with w_o ” (Berwald et al. 1995). Overall, Berwald et al. (1995) concludes that scattering between approximately 20 and 45° plays a large role in the rate of change P_τ , at least for a zenith sun.

Of the variables in Eq. (A.3), all except $\bar{\mu}_\infty$ are a function of the illumination conditions, the scattering phase function and the single scattering albedo. Conversely, $\bar{\mu}_\infty$ can be determined from the asymptotic regime. The asymptotic light regime is reached when the radiance no longer has an angular dependence, and only varies with depth due to water IOPs (coefficients and the scattering phase function) (Piening and McCormick 2003). Intuitively, a comparison of diffuse surface versus monodirectional illumination shows that the apparent optical properties of a water volume will more quickly reach asymptotic values for the diffuse surface illumination case (Piening and McCormick 2003). Zaneveld (1989) even uses a variant of the diffuse attenuation coefficient to determine the asymptotic value of the average cosine of the radiance distribution as shown below:

$$\begin{aligned}
 \frac{k}{\sigma} &= 1 - 0.52w_o - 0.44w_o^2 \\
 \bar{\mu}_\infty &= \frac{\sigma_a}{k} \\
 \bar{\mu}_\infty &= \frac{1 - w_o}{1 - 0.52w_o - 0.44w_o^2}
 \end{aligned} \tag{A.4}$$

The point of this discussion is not to seek an analytical equation to estimate the average cosine under a given set of conditions, but rather to highlight the variables that influence the average cosine of the radiance field, so that when a value is settled upon when using the two-stream approximation, it can be thoughtfully interpreted.

A.2 Uniform Illumination, Isotropic Scattering & a Non-Reflecting Bottom Boundary

Development of the two-stream approximation left the reader with two expressions as shown in Eqs. (7.12) and (7.13). Neither represent a complete solution of the RTE. To illustrate the solution, consider the specific case of uniform illumination, isotropic scattering and a non-reflecting bottom boundary layer. For a scenario with uniform incident illumination, the first order scattering source term can be described as follows:

$$S^*(\tau) = \frac{w_o}{4\pi} \int_0^{2\pi} d\phi \int_0^{+1} L_n e^{-\frac{\tau}{\mu}} d\mu \quad (\text{A.5})$$

However, the diffuse half range specific intensity equations are more easily handled if the source function in Eq. (A.5) is set equal to zero [$S^* = 0$] and instead a boundary condition is applied. In the case of uniform downwelling incident radiation, the boundary condition is conveniently $L^-(\tau = 0) = L_n$. Furthermore, with the added simplification of a uniform scattering phase function (*e.g.* $p = 1$ therefore $b = 1/2$), the diffuse half range specific intensities are reduced to:

$$-\bar{\mu}^- \frac{dL_{dave}^-(\tau)}{d\tau} = L_{dave}^-(\tau) - \frac{w_o}{2} L_{dave}^+(\tau) - \frac{w_o}{2} L_{dave}^-(\tau) \quad (\text{A.6})$$

$$\bar{\mu}^+ \frac{dL_{dave}^+(\tau)}{d\tau} = L_{dave}^+(\tau) - \frac{w_o}{2} L_{dave}^+(\tau) - \frac{w_o}{2} L_{dave}^-(\tau) \quad (\text{A.7})$$

If the two-stream Eq. ((A.6) and (A.7)) are added and subtracted, the result will yield:

$$\bar{\mu} \frac{d}{d\tau} [L_{dave}^+(\tau) - L_{dave}^-(\tau)] = (1 - w_o) [L_{dave}^+(\tau) + L_{dave}^-(\tau)] \quad (\text{A.8})$$

$$\bar{\mu} \frac{d}{d\tau} [L_{dave}^+(\tau) + L_{dave}^-(\tau)] = [L_{dave}^+(\tau) - L_{dave}^-(\tau)] \quad (\text{A.9})$$

Differentiate of Eq. (A.9) will yield:

$$\bar{\mu} \frac{d^2}{d\tau^2} [L_{dave}^+(\tau) + L_{dave}^-(\tau)] = \frac{d[L_{dave}^+(\tau) - L_{dave}^-(\tau)]}{d\tau} \quad (\text{A.10})$$

The right side of Eq. (A.10) is equal to Eq. (A.8) so that the appropriate substitution will yield:

$$\frac{d^2}{d\tau^2}[L_{dave}^+(\tau) + L_{dave}^-(\tau)] = \frac{(1 - w_o)[L_{dave}^+(\tau) + L_{dave}^-(\tau)]}{\bar{\mu}^2} \quad (\text{A.11})$$

If the differentiation and substitution are repeated for Eq. (A.8), the following is obtained:

$$\frac{d^2}{d\tau^2}[L_{dave}^+(\tau) - L_{dave}^-(\tau)] = \frac{(1 - w_o)[L_{dave}^+(\tau) - L_{dave}^-(\tau)]}{\bar{\mu}^2} \quad (\text{A.12})$$

Eqs. (A.11) and (A.12) have the general form:

$$\frac{d^2 Y}{d\tau^2} = \Gamma^2 Y \quad (\text{A.13})$$

which has the general solution of a sum of positive and negative exponentials:

$$Y = X e^{\Gamma\tau} + Z e^{-\Gamma\tau} \quad (\text{A.14})$$

where X and Z are constants. Since the sum and difference of the half range specific intensities is a sum of exponentials, the individual half range specific intensities can also be expressed in the same manner, where A , B , C and D are constants, and where $\Gamma = \frac{\sqrt{(1-w_o)}}{\bar{\mu}}$:

$$L_{dave}^+(\tau) = A e^{\Gamma\tau} + B e^{-\Gamma\tau} \quad (\text{A.15})$$

$$L_{dave}^-(\tau) = C e^{\Gamma\tau} + D e^{-\Gamma\tau} \quad (\text{A.16})$$

The constants A , B , C and D are related as follows (Thomas and Stamnes 1999).

$$\frac{C}{A} = \frac{B}{D} = \frac{w_o}{2 - w_o + 2\bar{\mu}\Gamma} = \frac{1 - \bar{\mu}\Gamma}{1 + \bar{\mu}\Gamma} = \frac{1 - \sqrt{1 - w_o}}{1 + \sqrt{1 - w_o}} = \rho_\infty \quad (\text{A.17})$$

In this simplified example, the boundary conditions of $L^-(\tau = 0) = L_n = \text{constant}$ and $L^+(\tau = \tau^*) = 0$ will yield a simple analytical expression for the diffuse half range intensities. Using the ratio relationships and the boundary conditions the constants can be determined as shown below:

$$L_{dave}^+(\tau) = A e^{\Gamma\tau} + \rho_\infty D e^{-\Gamma\tau} \quad (\text{A.18})$$

$$L_{dave}^-(\tau) = \rho_\infty A e^{\Gamma\tau} + D e^{-\Gamma\tau} \quad (\text{A.19})$$

$$L_{dave}^-(\tau = 0) = \rho_\infty A + D = L_n \quad (\text{A.20})$$

$$L_{dave}^+(\tau = \tau^*) = Ae^{\Gamma\tau^*} + \rho_\infty De^{-\Gamma\tau^*} = 0 \quad (\text{A.21})$$

which yields:

$$A = \frac{-\rho_\infty L_n e^{-\Gamma\tau^*}}{e^{\Gamma\tau^*} - \rho_\infty^2 e^{-\Gamma\tau^*}} \quad (\text{A.22})$$

$$D = \frac{L_n e^{\Gamma\tau^*}}{e^{\Gamma\tau^*} - \rho_\infty^2 e^{-\Gamma\tau^*}} \quad (\text{A.23})$$

Therefore, the solutions to the diffuse *average* half range intensities are:

$$L_{dave}^+(\tau) = \frac{L_n \rho_\infty}{E} \left[e^{\Gamma(\tau^* - \tau)} - e^{-\Gamma(\tau^* - \tau)} \right] \quad (\text{A.24})$$

$$L_{dave}^-(\tau) = \frac{L_n}{E} \left[e^{\Gamma(\tau^* - \tau)} - \rho_\infty^2 e^{-\Gamma(\tau^* - \tau)} \right] \quad (\text{A.25})$$

where the denominator is:

$$E = e^{\Gamma\tau^*} - \rho_\infty^2 e^{-\Gamma\tau^*} \quad (\text{A.26})$$

The generalized source function stems from the full source function shown in Eq. (A.27) which includes thermal, first order and multiple scattering terms (using solid angle instead of θ and ϕ for brevity of expression).

$$\begin{aligned} S_v(\tau, \Omega) &= \underbrace{[1 - w_o(v)] B_v(T)}_{\text{thermal emission}} \\ &+ \underbrace{S^*(\tau, \Omega)}_{\text{first-order scattering}} \\ &+ \underbrace{\frac{w_o}{4\pi} \int_{4\pi} p(\Omega', \Omega) L_d(\Omega') d\Omega'}_{\text{multiple scattering}} \end{aligned} \quad (\text{A.27})$$

The first order scattering term S^* specifically describes the initial scattering event for the incident external source. For example, let L_s represent the solar radiation field that survives extinction and is equal to (Thomas and Stamnes 1999):

$$\begin{aligned}
L_s(\tau, \Omega) &= E_s e^{-\frac{\tau}{\mu_0}} \delta(\Omega' - \Omega'_0) \\
&= E_s e^{-\frac{\tau}{\mu_0}} \delta(\theta - \theta_0) \delta(\phi - \phi_0) \\
&= E_s e^{-\frac{\tau}{\mu_0}} \delta(\mu - \mu_0) \delta(\phi - \phi_0)
\end{aligned} \tag{A.28}$$

Then the first scattering event for the solar radiation is determined by the scattering phase function that describes how radiation traveling in direction Ω' is scattered into direction Ω as shown in Eq. (A.29) (*e.g.* the notation $p(\Omega', \Omega)$ refers to photons moving in direction Ω' and then scattering into direction Ω):

$$\begin{aligned}
S^*(\tau, \Omega) &= \frac{w_o}{4\pi} \int_{4\pi} p(\Omega', \Omega) E_s e^{-\frac{\tau}{\mu_0}} \delta(\Omega' - \Omega'_o) d\Omega' \\
&= \frac{w_o}{4\pi} p(\Omega', \Omega) E_s e^{-\frac{\tau}{\mu_0}}
\end{aligned} \tag{A.29}$$

For diffuse illumination, the single-scattered source term is shown in A.5, but set to zero in exchange for the boundary condition $L^-(\tau = 0) = L_n$. This leaves the source emittance term and the multiply scattered terms. The multiply scattered terms are replaced by a constant, diffuse, average specific intensity term, and a phase function of $p = 1$, reducing the total source function shown in A.27 to that shown in Eq. (A.30).

$$\begin{aligned}
S(\tau) &\approx \frac{w_o}{2} [L_{d_{ave}}^+(\tau) + L_{d_{ave}}^-(\tau)] + (1 - w_o)B \\
&= \frac{w_o L_n}{2E} (1 + \rho_\infty) [e^{\Gamma(\tau^* - \tau)} - \rho_\infty e^{-\Gamma(\tau^* - \tau)}]
\end{aligned} \tag{A.30}$$

The solution to the radiative transfer (the specific intensity) can be determined after integration of the approximate source function shown above, or Eqs. (A.31) and (A.32):

$$L^-(\tau, \mu) = \int_0^\tau \frac{S(\tau')}{\mu} e^{-\frac{(\tau - \tau')}{\mu}} d\tau' + L_n e^{-\frac{\tau}{\mu}} \tag{A.31}$$

$$L^+(\tau, \mu) = \int_\tau^{\tau^*} \frac{S(\tau')}{\mu} e^{-\frac{(\tau' - \tau)}{\mu}} d\tau' \tag{A.32}$$

The full result of integration is shown in Eqs. (A.33) and (A.34).

$$\begin{aligned}
L^-(\tau, \mu) &= \frac{L_n}{E} [T^-(\mu) e^{\Gamma(\tau^* - \tau)} - T^+(\mu) \rho_\infty^2 e^{-\Gamma(\tau^* - \tau)}] \\
&+ \frac{L_n}{E} [(1 - T^-(\mu)) e^{\Gamma\tau^* - \frac{\tau}{\mu}} - \rho_\infty^2 (1 - T^+(\mu)) e^{-\Gamma\tau^* - \frac{\tau}{\mu}}]
\end{aligned} \tag{A.33}$$

$$\begin{aligned}
L^+(\tau, \mu) &= \frac{L_n \rho_\infty}{E} \left[T^+(\mu) e^{\Gamma(\tau^* - \tau)} - T^-(\mu) e^{-\Gamma(\tau^* - \tau)} \right] \\
&+ \frac{L_n \rho_\infty}{E} \left[T^-(\mu) - T^+(\mu) \right] e^{-\frac{(\tau^* - \tau)}{\mu}}
\end{aligned} \tag{A.34}$$

where T is:

$$T^\pm = \frac{(1 \pm \Gamma \bar{\mu})}{(1 \pm \Gamma \mu)} \tag{A.35}$$

A.2.1 Uniform Illumination, Isotropic Scattering & a Lambertian-Reflecting Bottom Boundary

The two-stream approximation described in the preceding section will now be extended to handle reflective (benthic) boundary conditions. For a Lambertian reflecting boundary, $L_d^+(\tau^*)$ can be expressed according to Eq. (A.36), where r is the reflectance:

$$\begin{aligned}
L_{d_{ave}}^+(\tau^*) &= r F^- \\
&= r \int_- L_{d_{ave}}^-(\tau^*) \cos \bar{\theta} d\Omega \\
&= r \bar{\mu} L_{d_{ave}}^-(\tau^*) \int_- d\Omega \\
&= r 2\pi \bar{\mu} L_{d_{ave}}^-(\tau^*)
\end{aligned} \tag{A.36}$$

For isotropic scattering and in the absence of thermal medium emittance, the boundary conditions in Eqs. (A.20) and (A.21) can be re-expressed as follows:

$$L_{d_{ave}}^-(\tau = 0) = \rho_\infty A + D = L_n \tag{A.37}$$

$$L_{d_{ave}}^+(\tau = \tau^*) = A e^{\Gamma \tau^*} + \rho_\infty D e^{-\Gamma \tau^*} = r 2\pi \bar{\mu} L_{d_{ave}}^-(\tau^*) \tag{A.38}$$

where $L_{d_{ave}}^-(\tau^*) = \rho_\infty A e^{\Gamma \tau^*} + D e^{-\Gamma \tau^*}$. If the constants A and D are isolated, the solution for the half range intensity can be solved for in exactly the same manner as that described above for a non-reflecting bottom boundary condition.

$$A = \frac{L_n [\rho_\infty e^{-\Gamma \tau^*} - r 2\pi \bar{\mu} e^{-\Gamma \tau^*}]}{-e^{\Gamma \tau^*} + \rho_\infty^2 e^{-\Gamma \tau^*} - r 2\pi \bar{\mu} \rho_\infty [e^{-\Gamma \tau^*} - e^{\Gamma \tau^*}]} \tag{A.39}$$

$$D = \frac{Ln [e^{\Gamma\tau^*} - r2\pi\bar{\mu}\rho_\infty e^{\Gamma\tau^*}]}{e^{\Gamma\tau^*} - \rho_\infty^2 e^{-\Gamma\tau^*} - r2\pi\bar{\mu}\rho_\infty [e^{\Gamma\tau^*} - e^{-\Gamma\tau^*}]} \quad (\text{A.40})$$

For an overhead camera, viewing a reflective bottom boundary within a water volume, the solution to the sensor reaching radiance can be solved in a sequential manner. First, the downwelling radiance striking the bottom boundary at τ^* is:

$$\begin{aligned} L^-(\tau^*, \mu) &= L^-(\tau = 0)e^{-\frac{\tau^*}{\mu}} + \int_0^{\tau^*} \frac{S(\tau')}{\mu} e^{-\frac{(\tau^* - \tau')}{\mu}} d\tau' \\ &= L_n e^{-\frac{\tau^*}{\mu}} \\ &+ \frac{w_o}{2} \int_0^{\tau^*} \left[A(1 + \rho_\infty) e^{\left[\Gamma\tau' - \frac{(\tau^* - \tau')}{\mu} \right]} d\tau' \right] \\ &+ \frac{w_o}{2} \int_0^{\tau^*} \left[D(\rho_\infty + 1) e^{\left[-\Gamma\tau' - \frac{(\tau^* - \tau')}{\mu} \right]} d\tau' \right] \end{aligned} \quad (\text{A.41})$$

After all constant values are removed from within the integral, integration of an exponential persists. The solution is then:

$$\begin{aligned} L^-(\tau^*, \mu) &= L_n e^{-\frac{\tau^*}{\mu}} \\ &+ \frac{w_o}{2\mu(\Gamma + \frac{1}{\mu})} A(1 + \rho_\infty) e^{-\frac{\tau^*}{\mu}} \left[e^{\left(\Gamma + \frac{1}{\mu}\right)\tau^*} - 1 \right] \\ &+ \frac{w_o}{-2\mu(\Gamma - \frac{1}{\mu})} D(\rho_\infty + 1) e^{-\frac{\tau^*}{\mu}} \left[e^{-\left(\Gamma - \frac{1}{\mu}\right)\tau^*} - 1 \right] \end{aligned} \quad (\text{A.42})$$

Eq. (A.42) describes the downwelling radiation that has been attenuated along transit through the water volume, combined within the in-scattered source term along the first leg of the trip through the water volume (top to bottom boundary - see Figure A.2). If the result of Eq. (A.42) is then reflected from the bottom boundary and attenuated on the second leg through the water volume (bottom back to top boundary), the result will describe the sensor reaching radiance that has been downwelled, attenuated (along leg-one and leg-two), in-scattered (along leg-one) and reflected (as illustrated by the solid line path in Figure A.2).

The remaining radiance to be determined is that which is in-scattered and attenuated along leg-two of the trip (or the radiance emerging from the top of the slab). This can be found as a result of integration of Eq. (A.43) as shown below:

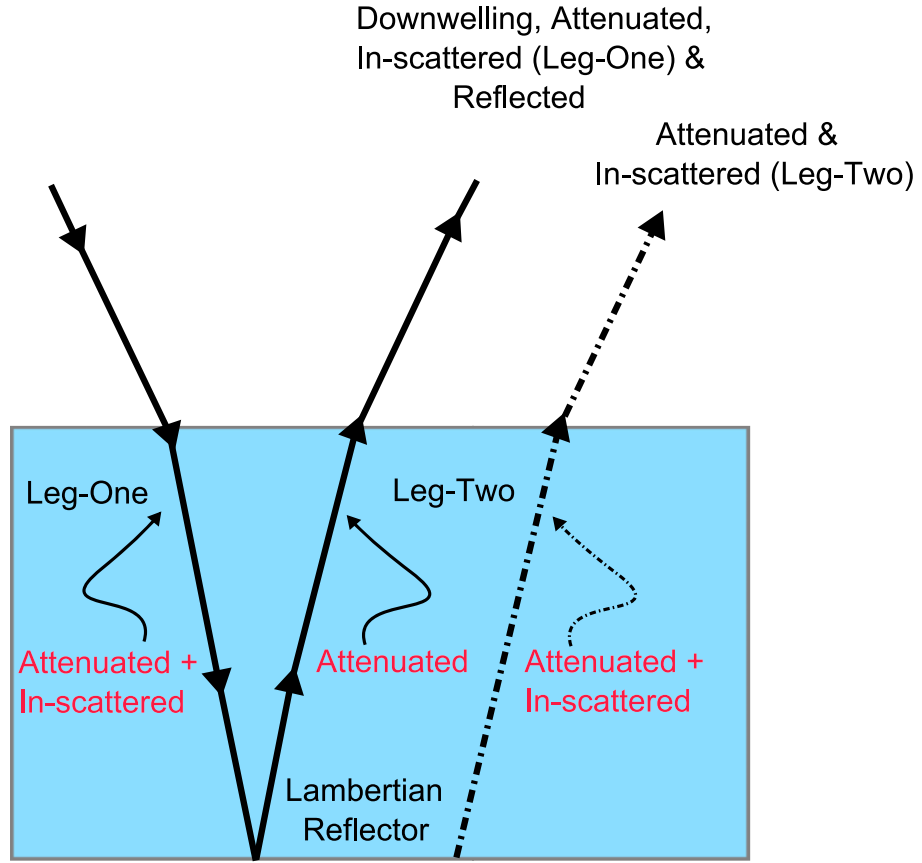


Figure A.2: The path described by the solid lines denotes sensor reaching radiance that has been downwelled, attenuated (along leg-one and leg-two), in-scattered (along leg-one) and reflected. The path described by the dashed line is the sensor reaching radiance that has been attenuated and up-scattered along leg-two.

$$\begin{aligned}
 L^+(0, \mu) &= \underbrace{L_v^+(\tau^*, \mu) e^{-\frac{\tau^*}{\mu}}}_{\text{downwelling \& reflected}} + \int_0^{\tau^*} \frac{S(\tau')}{\mu} e^{-\frac{\tau'}{\mu}} d\tau' \\
 &= \underbrace{L_v^+(\tau^*, \mu) e^{-\frac{\tau^*}{\mu}}}_{\text{downwelling \& reflected}} \\
 &\quad + \frac{w_o}{2} \int_0^{\tau^*} \left[A(1 + \rho_\infty) e^{\left[\Gamma \tau' - \frac{\tau'}{\mu} \right]} d\tau' \right]
 \end{aligned}$$

$$+ \frac{w_o}{2} \int_0^{\tau^*} \left[D(\rho_\infty + 1) e^{\left[-\Gamma\tau' - \frac{\tau'}{\mu}\right]} d\tau' \right] \quad (\text{A.43})$$

The result of integrating the source function is:

$$\begin{aligned} \int_0^{\tau^*} \frac{S(\tau')}{\mu} e^{-\frac{\tau'}{\mu}} d\tau' &= \frac{w_o A(1 + \rho_\infty)}{2\mu(\Gamma - \frac{1}{\mu})} \left[e^{\tau^*(\Gamma - \frac{1}{\mu})} - 1 \right] \\ &+ \frac{w_o D(\rho_\infty + 1)}{-2\mu(\Gamma + \frac{1}{\mu})} \left[e^{-\tau^*(\Gamma + \frac{1}{\mu})} - 1 \right] \end{aligned} \quad (\text{A.44})$$

The sensor reaching radiance is then:

$$\begin{aligned} L_{\text{sensor}}(\mu) &= L_{\text{dave}}^+(0, \mu) \underbrace{[\tau_F(\mu)]^2}_{\text{Fresnel Transmittance}} \\ &+ L_{\text{reflected}}(\mu) \end{aligned} \quad (\text{A.45})$$

where τ_F refers to the Fresnel transmittance factor that governs the passage of radiance between two materials that differ in refractive index (the air-water interface), and $L_{\text{reflected}}$ refers to the radiance reflected from the air-water interface. Note that the downwelling-reflected radiance field passes through the air-water interface twice and the Fresnel transmittance function must be applied during each passage.

A.2.2 Anisotropic Two-Stream Approximation

If the phase function is approximated using Legendre polynomials, but retaining only the first two moments, the back and forward scattering functions shown in Eqs. (7.10) and (7.11) become:

$$b = \frac{1}{2}p(-\bar{\mu}, \bar{\mu}) = \frac{1}{2}(1 - 3g\bar{\mu}^2) \quad (\text{A.46})$$

$$1 - b = \frac{1}{2}p(\bar{\mu}, \bar{\mu}) = \frac{1}{2}(1 + 3g\bar{\mu}^2) \quad (\text{A.47})$$

Using the same methodology described in Section A.2, differentiation and substitution can be carried out to determine the solution for anisotropic scattering using the half-range source functions shown below:

$$S^+(\tau) = w_o(1 - b)L_{\text{dave}}^+(\tau) + w_obL_{\text{dave}}^-(\tau) + S^{*+}(\tau) \quad (\text{A.48})$$

$$S^-(\tau) = w_o(1 - b)L_{\text{dave}}^-(\tau) + w_obL_{\text{dave}}^+(\tau) + S^{*-}(\tau) \quad (\text{A.49})$$


Substitution of the expression for the diffuse half-range intensities into the source functions, followed by integration provides a two-stream approximate solution for anisotropic scattering, in a manner analogous to that shown for isotropic scattering in Section A.2. For more complicated situations, the user can expand the phase function using a greater number of terms. A suitable approximation or discretization of the phase function is often achieved using Legendre polynomial expansion, and the reader is referred to Zender (2008), Thomas and Stamnes (1999), Marshak and Davis (2005), and MacRobert (1967) for a general discussion. Furthermore, the specific intensity can be approximated using more than the average value (the two-stream approximation) and more than the first two moments (the Eddington approximation) of the radiance distribution. This continued expansion of the specific intensity and discretization of the phase function shows the progressive link between approximate and numerical solutions of the RTE.

As inferred above, the two-stream approximation is a fledgling step toward what is termed a deterministic numerical method, where deterministic or explicit refers to all non-statistical (or non-Monte Carlo) radiative transfer solutions (Marshak and Davis 2005). These deterministic models usually strive to solve the less generalized plane-parallel radiative transfer problem, and solve for the whole radiance field, which is explicitly described in “discrete” terms. This is in contrast to statistical Monte Carlo methods which estimate the desired radiative quantities with some level of confidence that varies with the number of photons/trajectories simulated, and the variance of the estimate (Marshak and Davis 2005).


The most well-known in-water radiative transfer solution is the invariant imbedding deterministic method described by Mobley (1994). However, Mobley (1994) notes that the benefits associated with invariant imbedding theory (*e.g.* speed) are achieved at the price of mathematical complexity, and that “invariant imbedding methods require a considerable amount of mathematical development in going from the RTE to its solution, and the associated computer programming is much more tedious”. Mobley (1994) further reports that although Monte Carlo methods can be understood “in terms of fundamental radiative processes and elementary probability theory”, the same cannot be said of other solutions to the RTE. This certainly seems true when examining the two-stream solution since the conceptual transparency that may be gained from using an analytical method will not always be paralleled by an equally transparent mathematical solution.

Appendix B


Trends in Variance




(a.) Coastal, Dark.




(b.) Bay, Dark.




(c.) Turbid, Dark.



(d.) Coastal, Bright.




(e.) Bay, Bright.




(f.) Turbid, Bright.


Figure B.1: Variance in radiance (W/m^2sr) and trend lines for HG ($g = 0.92$), with mean photon densities of 9000 (diamonds), 3000 (squares) and 1000 (triangles) $\frac{photons}{m^3}$. The dashed lines connect the observed data points, while the solid lines show the computed trend line.




(a.) Coastal, Dark.




(b.) Bay, Dark.




(c.) Turbid, Dark.



(d.) Coastal, Bright.




(e.) Bay, Bright.




(f.) Turbid, Bright.


Figure B.2: Variance in radiance (W/m^2sr) and trend lines for HG ($g = 0.92$), with mean photon densities of 9000 (diamonds), 3000 (squares) and 1000 (triangles) $\frac{photons}{m^3}$. The dashed lines connect the observed data points, while the solid lines show the computed trend line.




(a.) Coastal, Dark.




(b.) Bay, Dark.




(c.) Turbid, Dark.



(d.) Coastal, Bright.




(e.) Bay, Bright.




(f.) Turbid, Bright.


Figure B.3: Variance in radiance ($W/m^2 \cdot sr$) and trend lines for HG ($g = 0.92$), with mean photon densities of 9000 (diamonds), 3000 (squares) and 1000 (triangles) $\frac{photons}{m^3}$. The dashed lines connect the observed data points, while the solid lines show the computed trend line.




(a.) Coastal




(b.) Bay




(c.) Turbid



(d.) Coastal



(e.) Bay



(f.) Turbid

Figure B.4: Ratio of variance (HG , $g = 0.92$) as a function of mean photon density and search radius. Dark targets (dashed lines), with ratio 3000/9000 (diamonds), 1000/9000 (squares) and 1000/3000 (triangles). Bright targets (solid lines), with ratio 3000/9000 (Xs), 1000/9000 (asterisks) and 1000/3000 (circles). Traditional density estimation indicates that $\frac{\sigma_{N1}^2}{\sigma_{N2}^2} \propto \frac{N2}{N1}$.