

ABSTRACT

HYKES, JOSHUA MICHAEL. Radiation Source Mapping with Bayesian Inverse Methods. (Under the direction of Yousry Y. Azmy.)

We present a method to map the spectral and spatial distributions of radioactive sources using a small number of detectors. Locating and identifying radioactive materials is important for border monitoring, accounting for special nuclear material in processing facilities, and in clean-up operations. Most methods to analyze these problems make restrictive assumptions about the distribution of the source. In contrast, the source-mapping method presented here allows an arbitrary three-dimensional distribution in space and a flexible group and gamma peak distribution in energy. To apply the method, the system's geometry and materials must be known. A probabilistic Bayesian approach is used to solve the resulting inverse problem (IP) since the system of equations is ill-posed. The probabilistic approach also provides estimates of the confidence in the final source map prediction. A set of adjoint flux, discrete ordinates solutions, obtained in this work by the Denovo code, are required to efficiently compute detector responses from a candidate source distribution. These adjoint fluxes are then used to form the linear model to map the state space to the response space. The test for the method is simultaneously locating a set of ^{137}Cs and ^{60}Co gamma sources in an empty room. This test problem is solved using synthetic measurements generated by a Monte Carlo (MCNP) model and using experimental measurements that we collected for this purpose. With the synthetic data, the predicted source distributions identified the locations of the sources to within tens of centimeters, in a room with an approximately four-by-four meter floor plan. Most of the predicted source intensities were within a factor of ten of their true value. The chi-square value of the predicted source was within a factor of five from the expected value based on the number of measurements employed. With a favorable uniform initial guess, the predicted source map was nearly identical to the true distribution, and the source intensities agreed within the predicted uncertainty. Using experimental data, the mapping was more difficult due to laboratory limitations. However, by supplanting 14 flawed measurements (out of 69 total) with synthetic data, the proof-of-principle source mapping was nearly as accurate as the synthetic-only prediction.

© Copyright 2012 by Joshua Michael Hykes

Some Rights Reserved



This work is licensed under a
Creative Commons Attribution-NonCommercial-
NoDerivs 3.0 United States License.

To view a copy of this license, visit
<http://creativecommons.org/licenses/by-nc-nd/3.0/us/>
or send a letter to
Creative Commons
444 Castro Street, Suite 900
Mountain View, California, 94041, USA.

Radiation Source Mapping with Bayesian Inverse Methods

by
Joshua Michael Hykes

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Nuclear Engineering

Raleigh, North Carolina

2012

APPROVED BY:

Dmitry Y. Anistratov

John Mattingly

C.T. Kelly

Yousry Y. Azmy
Chair of Advisory Committee

DEDICATION

☛ To Beth ☛

BIOGRAPHY

Joshua Hykes graduated from the Pennsylvania State University with a bachelor's degree in nuclear engineering in 2007. At North Carolina State University, he earned a master's degree in 2009 and in 2012 a doctoral degree, both in nuclear engineering.

ACKNOWLEDGEMENTS

Primary credit goes to my advisor, Dr. Yousry Azmy, who is both an effective and caring guide. Dr. Anistratov offered several practical suggestions in this work and pointed me toward related work in other fields. Dr. Cacuci introduced me to many of the concepts relating to sensitivity and uncertainty analysis. Dr. Mattingly's expertise in radioactive sources and detector systems was especially valuable, without which I would still be conducting the experiment described in this thesis. Dr. Kelly's class on optimization turned out to be more useful than I expected for this project.

Gerald Wicks assisted me in finding a laboratory room and radioactive sources. I appreciate the help of Dr. Gardner and his graduate students (Jiaxin Wang, Zhijian Wang, Wesley Holmes, and Adan Calderon). They are experts in understanding and obtaining detector response functions for the 2×2 sodium iodide detector.

I also appreciate the encouragement from and camaraderie with my fellow Nuclear Computational Science Group members: Dan Gill, Joe Zerr, Sebastian Schunert, Sean O'Brien, Sameer Vhora, and Brian Powell. Without them, my time in graduate school may have been slightly shorter than it was, but it would have certainly been drearier.

Finally, I must recognize the support of my family. In more ways than one, I would not be here without my parents. My wife demonstrated hope, patience, and encouragement as the last stages of the research and writing seemed to never end.

This work was supported by the DOE Computational Science Graduate Fellowship, provided under grant DE-FG02-97ER25308. Additional support came from the United States Nuclear Regulatory Commission in the form of a Graduate Fellowship in Nuclear Engineering, and from the Henry/Greebler and Dannels Scholarships from the American Nuclear Society.

TABLE OF CONTENTS

LIST OF TABLES	vi
LIST OF FIGURES	viii
Chapter 1 Introduction	1
Chapter 2 Solving inverse problems	8
Chapter 3 Inverse radiation transport review	33
Chapter 4 Source mapping methods	49
Chapter 5 Source mapping in a controlled environment	73
Chapter 6 Conclusions and future work	109
Appendices	123
Appendix A Responses and the adjoint flux	124
Appendix B Gamma cross section uncertainties	134
Appendix C Source mapping experiment	136
Appendix D Detector efficiencies	160
Appendix E Radiation source reconstruction with known geometry and materials using the adjoint	168
Appendix F Mathematical notes	181

LIST OF TABLES

Table 4.1	Parameters for source inversion.	59
Table 5.1	The energies E and group index i for the 16 energy group structure used for this test problem. The structure is based on the fine-group gamma shielding cross section library from SCALE. ¹¹ Energies, in keV, are the upper limit of the respective group. The groups above ^{60}Co 's highest peak (1.332 MeV) and below 100 keV are omitted.	78
Table 5.2	Comparing the predicted source emissions against the true intensities. The intensities are in total emissions per gamma line per second. They are summed over all spatial cells. The true source intensities are computed with the activities from Table c.4 and the yields from Table c.2. Unless noted otherwise, the initial guess was set to ten, $\vec{\alpha}_{\text{init}} = 10$	88
Table 5.3	Comparing the predicted source emissions per energy per second summed over space against the true values for the carpet test problem.	108
Table A.1	Description of the problem for numerical-adjoint test. The spatial domain is a 5 cm cube, with $x, y, z \in [0, 5]$. Vacuum boundary conditions are applied, thus eliminating the need for the bilinear concomitant.	126
Table A.2	Comparison of relative inner product difference ϵ of Equation A.4 for various TORT and Denovo spatial discretizations and flux convergence criteria.	127
Table A.3	Numerically-adjoint spatial discretizations in TORT and Denovo in the test.	128
Table B.1	Estimated uncertainties in photoelectric cross section in EPDL97 (assumed to be 1σ).	135
Table C.1	Detector components used in experiment. All components are branded Ortec except the NaI crystal.	140
Table C.2	^{137}Cs and ^{60}Co radioisotope properties. Gamma energies and branching fractions are from NNDC. ¹⁵ Half-lives are from NIST. ¹¹⁶	141
Table C.3	The ^{137}Cs and ^{60}Co sources used in this experiment, with their specified calibration date and activity. Most of the sources are undated.	142

Table c.4	Estimated sources activities based on computed activities of dated sources 11, 21, and 24 and on the measured peak area ratios. For each undated ^{137}Cs source, there is one peak and one dated ^{137}Cs reference source, so we have only one activity estimate. For each undated ^{60}Co source, there are two dated ^{60}Co reference sources and two peaks, yielding four activity estimates. The reported activity (and the derived 1 μCi date) are based on a weighted average of the four estimates.	145
Table c.5	The elemental composition of materials used in the 2144 model. The densities and compositions are from the Standard Composition manual of SCALE. ¹¹ Integer values of composition are the atomic abundance per molecule, while real values indicate atomic weight percent.	148
Table c.6	Source locations for ^{60}Co . All dimensions are in cm. The y location of all sources was 9.5 cm, and 90.2 cm for z . All the ^{137}Cs sources (IDs 11-16) were located at (440, 5, 1) cm.	152
Table c.7	Detector locations. All dimensions are in cm.	152
Table c.8	Gamma particle energies present in background. These were used in energy calibration. The peak energies are from the National Nuclear Data Center. ¹⁵	156
Table c.9	The normalized χ^2 -value for the four peak fits in Figure c.10. The expected normal distribution around one for the normalized χ^2/n_{dof} is also provided. All of the fits except for the ^{40}K peak fit fall less than two standard deviations from the mean.	159
Table D.1	Peak and total intrinsic efficiencies for the bare and collimated detectors, as well as the mean chord length, as computed by MCNP. The one-sigma uncertainties are derived from the Monte Carlo statistical errors. For reference, the mean chord length predicted by the Dirac formula is 3.387 cm.	165
Table E.1	Description of the source-reconstruction problem. All faces had vacuum boundary conditions.	173
Table E.2	Comparison of number of transport solves needed for the adjoint method and the forward Levenberg-Marquardt iterative method. TS stands for the number of transport solves, and Its. is the number of iterations.	180

LIST OF FIGURES

Figure 2.1	Direct problem.	9
Figure 2.2	Inverse problem—causation.	10
Figure 2.3	Inverse problem—model identification.	10
Figure 5.1	The true distribution of the ^{137}Cs and ^{60}Co sources, as represented on the Cartesian computational mesh. The ^{137}Cs sources are grouped together in a point source, and the ^{60}Co sources are aligned in a row parallel to the x -axis. The natural logarithm of the source intensity is plotted, that is, $\log_e(q_{pijk}/q_0)$, where q_0 is a unit scalar needed to make the argument of the logarithm dimensionless and q_{pijk} is the source in peak p in spatial cell (i, j, k) . The units of q_{pijk} and q_0 are particles emitted per second. In reality, the blue cells have no sources, but they are shown here with small source intensities to aid comparison with the results presented below.	75
Figure 5.2	A comparison of the measured and computed responses for the source mapping with photopeak synthetic responses.	84
Figure 5.3	Volume rendering of the mean of the posterior distribution of the source using the photopeak synthetic responses. The source prediction has suppressed source regions in cells around the primary source cells. This results in dark areas near the sources visible in the insets.	85
Figure 5.4	Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using the photopeak synthetic responses. The natural logarithm of the source intensity is plotted, that is, $\log_e(q_{pijk}/q_0)$, where q_0 is a unit scalar needed to make the argument of the logarithm dimensionless. The units of q_{pijk} and q_0 are particles emitted per second.	86
Figure 5.4	Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using the photopeak synthetic responses. (continued)	87
Figure 5.5	Change in chi-square and the number of function evaluations with respect to the intensity of the uniform initial guess for optimization.	89
Figure 5.6	The local minimum reached when using $\vec{\alpha}_{\text{init}} = -6$. The slices in the x - y plane are the mean and standard deviation of the posterior distribution of the 661.7 keV emissions from ^{137}Cs using the uncollided, photopeak responses. The source is primarily predicted in cells close to detector locations.	90

Figure 5.7	Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using the photopeak synthetic responses with a favorable initial guess, $\vec{a}_{\text{init}} = 1$	91
Figure 5.8	Change in source intensity summed over space with respect to the initial guess for optimization. The upper, cyan horizontal line corresponds to the true value of the 662 keV peak, while the lower, blue line corresponds to the intensity of the two ^{60}Co peaks. The shaded area around a line represents the 1σ confidence interval about the mean.	92
Figure 5.9	Measured and computed responses for the source mapping with synthetic photopeak and continuum responses. For indices where no “computed with true source” data point is present, the value is off the bottom of the scale. See Figure 5.2 for an explanation of the absolute and relative error plots.	94
Figure 5.10	Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using both the photopeak and continuum synthetic responses. The natural logarithm of the source intensity is plotted, that is, $\log_e(q_{p_{ijk}}/q_0)$, where q_0 is a unit scalar needed to make the argument of the logarithm dimensionless. The units of $q_{p_{ijk}}$ and q_0 are particles emitted per second.	95
Figure 5.10	Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using both the photopeak and continuum synthetic responses. (continued)	96
Figure 5.10	Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using both the photopeak and continuum synthetic responses. (continued)	97
Figure 5.11	Measured and modeled scalar flux for the experiment in 2144 Burlington. If an experiment value is missing from the plot, its value is zero. The model in this case is MCNP.	99

Figure 5.12	Measured and computed responses for the source mapping with experimental data. For indices in the top subplot where a data point is missing, the value is off the bottom of the scale. See Figure 5.2 for an explanation of the absolute and relative error plots. The Denovo-computed response using the true source distribution for the 662 keV peak at location 2 looking in the +z direction is about 100 times too large. This is because the ^{137}Cs source is located within the fuzzy region around $\cos \theta = \mu = 0$ modeled in Denovo (see section 4.5.3). This inaccuracy is also present in Figure 5.2 and Figure 5.9A, but it is obscured there by the larger number of responses.	101
Figure 5.13	Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using the photopeak responses of the experimental data. The natural logarithm of the source intensity is plotted, that is, $\log_e(q_{pijk}/q_0)$, where q_0 is a unit scalar needed to make the argument of the logarithm dimensionless. The units of q_{pijk} and q_0 are particles emitted per second.	102
Figure 5.14	The true distribution of the ^{137}Cs and ^{60}Co sources for the carpet source problem. In reality, the blue cells have no sources, but they are shown here with small source intensities to aid comparison with the results presented below.	104
Figure 5.15	Measured and computed responses for the source mapping with the carpet source. For indices where no “computed with true source” data point is present, the value is off the bottom of the scale. See Figure 5.2 for an explanation of the absolute and relative error plots.	105
Figure 5.16	Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using the photopeak responses for the floor-distributed source. The natural logarithm of the source intensity is plotted, that is, $\log_e(q_{pijk}/q_0)$, where q_0 is a unit scalar needed to make the argument of the logarithm dimensionless. The units of q_{pijk} and q_0 are particles emitted per second.	106
Figure 5.16	Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using the photopeak responses for the floor-distributed source. (continued)	107
Figure A.1	Demonstrating the non-uniform mesh for the test problem. This is a cut of the x - y plane.	131
Figure A.2	Accuracy of adjoint as discrete solution degrades.	132

Figure C.1	A schematic of the detector components.	140
Figure C.2	The six ^{137}Cs sources and five ^{60}Co disc sources. Each source is labeled with the source id used throughout this experiment.	143
Figure C.3	Slices from the x - y plane of the 2144 BEL room layout.	146
Figure C.4	Pictures of the 2144 room.	147
Figure C.5	The count times for two counting sets. Each histogram bin is one hour wide.	149
Figure C.6	The source and detector configuration for the source-strength counts.	150
Figure C.7	The ^{137}Cs and ^{60}Co source locations.	151
Figure C.8	The 2π collimated detector. The lead bricks are $2 \times 4 \times 8$ inches in dimension. The blocks with a hollowed cylinder have an outside side-length of 4 inches and a height of 2 inches.	153
Figure C.9	The IDs of the detector measurements that were taken.	154
Figure C.10	A representative background spectrum taken with the NaI detector.	159
Figure D.1	Lead collimation around NaI detector crystal.	164
Figure D.2	Measured and modeled photopeak responses. No ^{137}Cs and ^{60}Co measurements were taken from location 5. The experiment value is zero where there is no blue experiment point accompanying a green model (predicted) value.	167
Figure E.1	Reconstructed source using 81 detector locations.	173
Figure E.2	Reconstructed sources using adaptive procedure and “aligned” detectors. Each filled circle or x-mark denotes one source unknown. The detectors are located at the source unknowns.	176
Figure E.3	Reconstructed sources using adaptive procedure with evenly-spaced detectors. The plus marks indicate the x position of the detectors. The vertical positioning of the pluses on the graph is meaningless.	177
Figure E.4	Reconstructed sources using adaptive procedure with Monte Carlo forward flux.	178
Figure E.4	Reconstructed sources using adaptive procedure with Monte Carlo forward flux. (continued)	179
Figure F.1	Two-dimensional representation of a mapping T . Here $Tx_1 = y_1$ and $Tx_2 = y_2$	183

INTRODUCTION



Before describing our new approach towards resolving the radiation source mapping problem, we first give a full description of the problem which we intend to solve. This is the goal of this chapter, as well as discussing potential applications of the developed methods. Also included is an outline of the remainder of this thesis.

1.1 DEFINING THE PROBLEM

The goal of this work is to build a map of the distribution of radioactive sources in a structure using limited detector measurements. Because there could be a number of variations on this problem, depending largely on what information is assumed known or unknown, we spend the following sections giving an account of the approach we have chosen. These details are delineated according to the known information, the unknown information, assumptions which we accept in order to solve the problem, and the external data and tools available.

1.1.1 *Knowns*

First we discuss the known information. Most significantly, the geometric configuration of the structure is sufficiently known. This structure might be a room, building, roadway, or other environment with suspected radiation sources present. The knowledge of the setting should include all significant structures and equipment in the space, whether walls and floors or machinery and piping. For instance, we need to know the location and thickness of all the walls in the space. Components should be included in the description of the configuration space if they could significantly affect the radiation fields around the detector or the sources. Engineering drawings, blueprints, or a CAD file would typically provide this data. Although outside the scope

of this work, engineers are exerting significant effort in methods to automatically build maps of an unknown building.^{1,105} This is often necessary for autonomous robots. Thus, the lack of prior knowledge of a structure may not eliminate the applicability of the methods here described, as long as an automated geometry mapping tool is available.

Closely related to the geometric configuration, the elemental composition and density of each material in the space is needed so that the macroscopic cross sections can be computed. However, precise data may not be required or necessary; this data could be estimated for most standard materials (such as concrete block walls or steel pipes).

The other information we require is radiation detector measurements. Here we focus on gamma (photon) detection, mainly for the ability to measure the energy spectrum of the gammas. In addition, high-energy photons are able to pass through thick materials, potentially increasing the information which our detector collects. Although neutrons are also highly penetrating, we lack the ability to measure the energy distribution of neutrons and so must typically settle for gross-counts, summing over all energies.

The multichannel analyzer (MCA) records the energy deposited by photons in the detector. This distribution comprises both photopeaks and the smooth continuum. The photopeaks are produced by uncollided photons (that is, photons which were emitted from the source and have made their first interaction within the detector) that deposit all their energy in the detector. The continuum portion of the spectrum is caused by scattered photons and by photons which deposit only a fraction of their energy in the detector. Analyzing the photopeaks is much simpler than analyzing the full spectrum. In this work, we desire the flexibility to analyze the photopeaks alone, or the continuum along with the photopeaks. Thus, capturing the effects of scattered photons is necessary.

We assume that collimated detectors are available. While some of the detectors may look in all directions, other detectors may be shielded for a more focused view. These collimated detectors presumably increase the overall information content of a set of detector signatures.

1.1.2 *Unknowns*

We are primarily interested in the distribution of radioactive sources, in both energy and space. The location and energy spectra of the sources are initially unknown to us. The goal is to estimate the source map $q(\vec{r}, E)$ over the domain of interest. We are not attempting to find the location and strength of a point source or of several point sources. Although the point source model is often sufficient, in this work we desire the ability to estimate distributed sources. Imagine, for instance, estimating the source density in contaminated water covering the floor of a basement room of a nuclear power plant after an accident. For measurements taken near the water, a point source approximation will be considerably lacking. As this example illustrates, one may not know exactly where a source is located, but it may be obvious that certain areas can be ruled out, for instance, the air above the contaminated water. Our method should allow one to exclude these regions without diminishing the generality of the new approach.

In this work, more emphasis is placed on the spatial dependence of the distribution. However, that is not to say that the energy dependence is ignored. Algorithms to reconstruct the sources' energy spectrum, usually with the goal of predicting the radionuclides present, have received much study over the past few decades.³⁴ Identifying radionuclides is outside the scope of this project.

While we are narrowing the scope of this project, we should note two further choices that limit the reach of the developed approach. First, we study situations in which the radioactive source does not change the material properties of the system, most importantly the macroscopic cross sections. This means that the problem of estimating the activity of a significant mass of a chunk of highly-enriched uranium is outside the scope of this project, since the uranium will self-attenuate and scatter some of the emitted gamma rays. On the other hand, these methods would be applicable to situations where low concentrations of a radionuclide are dissolved in water, or for a thin dusting of a source on the surface of walls, floor, or ground.

Second, we assume that the system is effectively time-independent. Thus, all significant objects, especially the source itself, are stationary. Furthermore, over the time span of the measurements, the source should be of constant magnitude. Thus, these methods are not designed for systems where the radiation source is composed of radionuclides with short half-lives.

1.1.3 Available data and tools

Attempting to solve this problem is only possible through the use of a wealth of prior data and tools to predict the detector response given a particular physical system. These tools and data are briefly described below.

Microscopic cross section libraries To predict the interaction of gamma rays with matter, and the flux of gammas throughout the system, macroscopic cross sections are required. These macroscopic cross sections are computed from the atomic number density of the nuclides constituting a material and these nuclides' microscopic cross section. Thankfully, the microscopic data for gamma rays is tabulated by `ENDF`¹⁶ and `EPDL97`²⁰ for all materials and energies of interest here. As described in Appendix B, one area in which the `EPDL97` lacks is an explicit, quantitative estimate of the uncertainty of these gamma cross sections.

Neutral-particle radiation transport codes Radiation transport simulation codes require as input (at a minimum) a description of the geometry, material cross sections, and radioactive source description, and they produce an estimate of the particle flux throughout the spatial and energy domain. These codes, executed in adjoint mode, are also able to compute the sensitivity of measurements to various parameters. The level of sophistication, efficiency, and fidelity of these codes is what allows one to undertake an inverse problem such as we have outlined here. Throughout this work, both deterministic and Monte Carlo codes are used, primarily `Denovo`²⁷ and `MCNP`.⁸⁰

Detector energy calibration While this should be part of a basic detection experiment, we mention here the importance of knowing the mapping between `MCA` bin number and deposited energy. This is typically performed using a calibration source with multiple peaks of known energy.

Detector response function The detector response function (`DRF`) provides a mapping between the photon flux incident on a detector and the response of that detector as measured by the `MCA`. The `DRF` is highly dependent on the type of detector, and even varies between detectors of the same type. Thankfully, considerable effort has been expended in computing these `DRFs` for our direct use.³⁵

1.1.4 *Computational resources and constraints*

As the goal of this project is to develop a proof-of-principle of the proposed approach, meeting stringent operational constraints is not of primary concern at this stage. While some inverse transport problems require that the calculation complete in a second or two,⁷⁷ we are under no such constraints here. In addition, we take the liberty to use high-performance computing resources, where applicable, that are typically unavailable in the field. Obviously these are important issues that must be resolved in future work if the developed method is to find acceptance among practitioners charged with locating and identifying radiation sources.

1.2 POTENTIAL APPLICATIONS

The prototypical application for the methods developed here is a room with a number of radioisotope source distributions – emitting neutrons, gamma rays, or both. We have physical access to the room to make detector measurements and to make note of the geometric shape and material composition of the room, whether in its original condition or having been modified by malicious acts or via an accident. Normally this could be done directly by a technician, but a robot could also perform the task remotely if exposure to the radiation field poses health risks. After obtaining the detector readings and the information about the geometry and materials configuration, an analyst would feed these as input to the proposed methods. The final result would be an estimate of the locations, activities, and energy spectra of the sources in the room and an estimate of the confidence level of the results.

One could imagine a number of related applications, some of which would require modifications or extensions to the methods here proposed.

- Estimating a radiation source map for a building could be useful in a number of situations. Of course, it would be useful for rescue or cleanup following the dispersal of nuclear materials throughout the building. In addition, preemptive scanning of a building or public area could help decision makers in the event of a radiological emergency to better judge between abnormal and normal radiation source levels.
- Inspectors of nuclear facilities currently get estimates of the quantity of nuclear materials caught in pipes using specialized instruments.⁹⁵ In some instances,

these “hold up” measurements might be improved with a better estimate of where the source material is lodged.

- Detectors carried by airplanes or helicopters are used to map the radiation source on the ground, for instance, in the aftermath of the Fukushima Daiichi accident. To see the results of one such survey, see the NNSA dose maps on page 17 of Ref. [62].
- Instead of using radiation portal monitors at border crossings to detect sources in only their own lane of traffic, an array of detectors feeding an inversion algorithm could give more detailed information about sources in several lanes.
- On a larger scale, the radiation field in a facility could be retroactively reconstructed using the readings of the detectors worn by their workers.
- This method could potentially be extended to detect a missing fuel rod in an assembly.

This is not an exhaustive list of potential end applications. The methods that we develop in this work are general enough to be applied in a variety of other settings.

1.3 OUTLINE OF THIS DOCUMENT

The rest of this thesis is structured as follows.

- Chapter 2 is a broad introduction to inverse problems, with no emphasis on the particular radiation mapping problem at hand. Inverse problems are defined, some history and mathematical notes are provided, and Bayesian and discrete methods for solving them are discussed.
- Chapter 3 discusses inverse problems in radiation transport. First a brief review of numerical methods for forward and adjoint radiation transport is provided, followed by a review of existing inverse problems in the field of radiation transport or nuclear engineering.
- We develop our radiation source mapping methods in chapter 4, discuss how these fit with the theory from chapter 2, and explain the necessary numerical techniques.
- Chapter 5 contains the numerical results of the test problem conducted in Burlington Laboratories room 2144 using both synthetic and experimental measurements.

- Conclusions are provided in chapter 6.
- The appendices contain supporting material on the numerical and experimental aspects of this work. We verify the correctness of the adjoint-computed responses in Appendix A. Appendix B explains the current availability of uncertainties for photon cross section data. Appendix C details the experiment conducted in 2144 Burlington. Appendix D describes the process of converting detector peak counts to uncollided flux, via the peak intrinsic efficiency. Finally, Appendix F reviews some functional analysis and operator theory which is assumed in chapter 2.

Before proceeding with the inverse radiation transport problem, we review the history and theory of inverse problems.

Tartaglia and cannonballs

In the 1530s, Niccolò Tartaglia, a self-taught mathematician from Brescia, Italy, busied himself with an important task for a nearby nobleman and benefactor – understanding, and thus predicting, the flight of a cannonball.³⁸ The cannon was a vital means of defense against surrounding city-states, but also against the powerful invaders from the east, the Ottoman Empire. In a military confrontation, the challenge for the gunners was to determine the inclination angle at which they should aim the cannon to hit the target. Although Tartaglia was unable to determine the correct laws of motion, discovered by Galileo a century later, his flawed calculations were still useful for the gunners.

Along the way, he made an important insight about gunnery: Tartaglia developed a method to hit the same spot with two different firing angles. In other words, the solution to the aiming problem was not unique. Unknowingly, he had stumbled upon an important characteristic of inverse problems.

2.1 WHAT IS AN INVERSE PROBLEM?

Much of the mental exercise we have come to call *reasoning* is the study of cause and effect. An event has a cause or causes and an effect or effects. We will package all the causes into one quantity which will be called the *input* to the system, and all the effects will be lumped into the *output*. For Tartaglia, the input was the size of the gunpowder charge, the size and weight of the cannonball, and the aiming angle.

The output was the distance the cannonball flew, or the location where it landed. In addition to identifying the input and output, a third task is developing a *model* that connects the input to the output. In the gunnery example, a model of the cannonball's motion was proposed by Galileo and Newton many years after Tartaglia.

In any scientific or engineering problem, one or part of the three components – input, model, or output – is missing. In a typical science class assignment, the missing component is the output. The input and system model are given, and the student proceeds to calculate the output. In this context, this process is called the *direct* or *forward* problem, and is illustrated in Figure 2.1. Finding the range of a cannonball

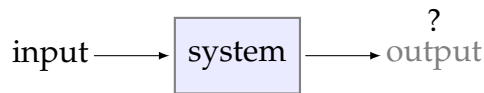


Figure 2.1: Direct problem.

is an example of a direct problem. We compute the horizontal range of flight r of a cannonball shot in a vacuum at an initial speed v_0 and an angle of elevation θ , assuming that the cannon and target are at the same altitude. The inputs are v_0 and θ , and the output is r . Integrating $\vec{F} = m\vec{a}$, the solution to this direct problem is

$$r(\theta, v_0) = \frac{v_0^2}{g} \sin 2\theta \quad , \quad (2.1)$$

where g is the constant acceleration due to gravity. Notice that for each set of inputs $\{\theta, v_0\}$, there is one and only one solution r .

Even if most classroom problems are direct, it is not true that most problems in science are direct. In many challenging situations, the output is available but either the input or model is missing. This leads to two types of *inverse* problems (IPs). First, for a *causation* problem, one estimates an unknown input from a given model and output. Figure 2.2 shows this.

One causation inverse problem for Tartaglia's cannonball is to compute the aiming angle θ for a given range r (assuming that the initial speed v_0 is fixed). This was the problem facing the gunners. Solving Equation 2.1 for θ , the two solutions to this

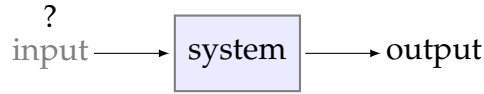


Figure 2.2: Inverse problem—causation.

inverse problem are then

$$\theta(r) = \frac{1}{2} \arcsin \frac{rg}{v_0^2} ,$$

or

$$\theta(r) = \frac{\pi}{2} - \frac{1}{2} \arcsin \frac{rg}{v_0^2} .$$

Just as Tartaglia had surmised, there are in general two angles for each range. In addition, only certain values for the range r are allowed in this expression. Otherwise, there may be no solution. For instance, the desired range r must be less than or equal to the maximum range v_0^2/g . Having multiple solutions (mathematicians call this *non-uniqueness*) or no solutions (called *non-existence*) is a recurring difficulty in resolving inverse problems.

The second type of inverse problem, *model identification*, seeks to find a suitable model using only known input and output, as shown in Figure 2.3. This is the problem that Galileo and Newton solved for ideal cannonball flight, showing that the path is a parabola. The result of this problem is an equation or rules mapping the input to output, in this case, Newton's laws of motion.




Figure 2.3: Inverse problem—model identification.

2.1.1 Historical examples of inverse problems

Groetsch in chapter one of his introductory book *Inverse Problems*³⁸ provides a number of helpful historical examples of IPs. His first example comes from the first paragraph of Book VII in Plato's *Republic*⁶⁰:

And now, I said, let me show in a figure how far our nature is enlightened or unenlightened: — Behold! human beings living in a underground den, which has a mouth open towards the light and reaching all along the den; here they have been from their childhood, and have their legs and necks chained so that they cannot move, and can only see before them, being prevented by the chains from turning round their heads. Above and behind them a fire is blazing at a distance, and between the fire and the prisoners there is a raised way; and you will see, if you look, a low wall built along the way, like the screen which marionette players have in front of them, over which they show the puppets.

In this scene of shadow puppets, the audience sees only the shadow and not the puppet. The prisoners, with some intuition and imagination, could make a reasonable guess for much of the outside activity, based only on the shadows on the cave wall. However, their captors would not struggle to find means to hide certain details from the unfortunate prisoners. This is a classic causation problem, where the outsiders' body positions are the input, the system is the projection of light from the fire onto the cave wall, and the output is the shadows.

For each configuration of captors between the fire and the cave wall, there is one shadow image. However, for a given shadow image, there are an infinite number of configurations of captors that would produce such an image. Again we see one of the severe challenges of inverse problems – a nonunique solution. That the reconstruction of the real world is not unique is a consequence of the dimensionality of the input and output. The real world has three spatial dimensions, while the prisoners are viewing a two-dimensional snapshot. The projection from three dimensions to two destroys much of the information.

Groetsch offers a number of examples of other historical inverse problems, including Archimedes' test of the purity of the king's gold crown, Johannes Kepler's three laws of planetary motion based on the measurements of Tycho Brahe, the unsuccessful and successful attempts to estimate the age of the earth by Joly, Kelvin, and Holmes, and, more recently, the computed tomography medical scan invented by Cormack and Hounsfield. There is also an abundance of examples in modern use. Although examples of inverse problems can be found in any field of research, geophysics has the lion's share. The problem of determining the age of the earth has already been

stated. For many years, measuring the earth's mass was a problem of interest. With normal methods of measuring precluded, Cavendish in 1798 used indirect methods to estimate the mean density of the earth as 5.4 g cm^{-3} , only slightly less than 5.5 g cm^{-3} , the commonly accepted density at present.³⁸

After obtaining a global mean density, it was only natural to look for the internal structure of the earth. A key tool in this search is the seismograph, which measures seismic waves. R.D. Oldham used the seismograph readings from around the globe during and after earthquakes to infer the existence of the earth's core. Prospectors have great interest in local geologic structure as well, as they search for oil or valuable minerals. Books about inverse problems authored by researchers in geophysics, for example, works by Menke⁸¹ and Tarantola,¹¹² are evidence of the field's progress in posing and solving inverse problems.

2.1.2 *Well-posed versus ill-posed*

In addition to the physical basis for inverse problems, reasoning from effect to cause, there is a closely related dichotomy of well-posed and ill-posed problems. Ill-posed problems lack one or more desirable properties of well-posed problems: existence and uniqueness of a solution, and continuity of the operator. Bertrand in 1889 described a probability riddle in geometry with more than one answer as ill-posed.⁶ However, Hadamard in 1902 more clearly identified these three characteristics.⁴⁰ At the time, he believed that ill-posed problems were a mathematical curiosity without any physical significance. In this observation, he was mistaken; as we will see, many physically significant inverse problems are indeed ill posed. Each of the three requirements formalizes our intuitions about how solutions should behave.

First, existence of a solution is paramount. Without a solution, there is little reason to proceed. The problem statement must admit at least one solution. For example, an overdetermined system of linear equations, for which the number of equations exceeds the number of unknowns, has in general no solution. It is also common for nonlinear problems to lack a solution. If no solution exists, the problem must be restated, perhaps by loosening some constraint, to admit at least one solution. In the case of an overdetermined linear system, the standard approach is the least squares method. In any case, without a solution, one has no where to go.

In other instances, there may be many solutions. For the example of linear systems,

an underdetermined system has infinitely many solutions, since the number of unknowns exceeds the number of linearly-independent equations. To pick from the set of possible solutions, more information is needed. In some ways this is better than having no solution, but it is typically not clear which solution of the many to choose.

Finally, small changes in the input data should only change the solution by a small amount. This property is known as stability, and it is related to the continuity property. If the solution is stable, then the inverse operator is continuous. Without stability, the solution can exhibit unphysical behavior. Since inverse problems are typically conducted with uncertain measured data, a lack of stability means that the solution may change significantly with only small measurement errors. Much research has focused on eliminating or reducing this instability, most of it going under the banner of *regularization*.

The concepts of existence, uniqueness, and continuity are all mathematical terms. In the deterministic solution of IPs, a number of mathematical concepts about operators and their inverses are helpful in analysing and resolving the problems. The next section presents a few of these concepts that are relevant to this work.

NOW WE RESTATE the conditions for a well-posed problem in more precise terms. (For a review of operator theory, see Appendix F.) Consider the problem $Tx = y$, where $y \in Y$ is given and the goal is to determine $x \in X$. T is a mapping from X to Y , that is, $T : X \rightarrow Y$. In this case, this problem is well-posed if these conditions are satisfied:⁵¹

1. Existence: for all $y \in Y$, there exists a solution $x \in X$.
2. Uniqueness: for all $y \in Y$, there is only one x such that $Tx = y$.
3. Continuity: for every $\epsilon > 0$, there is a $\delta > 0$ such that

$$\|T^{-1}y - T^{-1}y^*\|_X < \epsilon \quad \text{for all } y^* \text{ such that } \|y - y^*\|_Y < \delta \quad ,$$

where $\|\circ\|_X$ is the norm in X and $\|\circ\|_Y$ is the norm in Y .

Continuity is related to a stability or robustness condition, but it is a weak condition for well-posedness, in the sense that an inverse operator can be continuous but ill-conditioned, making the solution unstable.

Condition 1 for existence is equivalent to requiring that the mapping T is onto or surjective. Condition 2 for uniqueness is equivalent to requiring T to be one-to-one or

injective. If T is both onto and one-to-one, then it is known as bijective. A bijective mapping has an inverse. However, without the continuity or stability requirement, applying this inverse in practice may be unwise. Appendix F.2 defines the operator terms with more description.

Understanding the null space of an operator is important when finding the inverse of an operator. The null space of an operator T is defined as

$$\mathcal{N}(T) = \{x \mid Tx = 0\} \quad . \quad (2.2)$$

The null space is a subset of the domain $\mathcal{D}(T)$ of T , $\mathcal{N}(T) \subset \mathcal{D}(T) \subset X$. If the null space is equivalent to the domain, $\mathcal{N}(T) = \mathcal{D}(T)$, then T is the zero operator.

The conditions for a well-posed problem can be restated in terms of the null space and range of the operator. If X and Y are Hilbert spaces, then $Tx = y$ is well-posed if⁵¹

$$Y = \mathcal{R}(T), \quad \mathcal{N}(T) = \{0\}, \quad \mathcal{R}(T) = \overline{\mathcal{R}(T)} \quad . \quad (2.3)$$

Here \overline{M} is the closure of M , which is the set containing all the elements of M and the accumulation points (or limit points) of M . If M is a subset of X , $M \subset X$, then an element $x_0 \in X$ is an accumulation point of M if every neighborhood around x_0 contains at least one point $z \in M$ distinct from x_0 .⁶⁴

2.1.3 Fredholm integral equation

In many IP applications, the problem can be expressed mathematically as

$$\int_a^b K(s, t) f(t) dt = g(s) \quad . \quad (2.4)$$

This is known as the Fredholm integral equation of the first kind (IFK).¹²⁰ The kernel $K(s, t)$ and the right-hand side function $g(s)$ are known on the square $a \leq s, t \leq b$. The function $f(t)$ is to be determined on the interval $[a, b]$. In a geophysics example, $K(s, t)$ could be the known, or at least assumed, mathematical model describing wave transmission, $g(s)$ the measurements taken on the surface, and $f(t)$ some measure of the composition of the earth (density, for example). The integral operator accomplishes the projection of the distribution onto the measurement, in a manner analogous to the light creating shadows on Plato's cave wall by being projected from the fire.

The properties of the integral operator and the Fredholm equation make finding solutions of Equation 2.4 difficult, if not impossible. In fact, Wing makes a point of calling the determination of $f(t)$ *resolving* the equation instead of *solving* it, stressing the inherent difficulty in solving these types of equations.¹²⁰

Wing provides a simple and helpful tutorial on the inherent difficulties of the Fredholm integral equation of the first kind. It is helpful to define the operator

$$\mathcal{K} \circ = \int_a^b K(s, t) \circ dt \quad . \quad (2.5)$$

He takes as his first example in §5.2 a seemingly simple kernel,

$$K(s, t) = \cos(st), \quad s, t \in [-1, 1] \quad . \quad (2.6)$$

If we choose an odd function for f , such that $f(t) = -f(-t)$, then, since $K(s, t)$ is even, we have

$$\mathcal{K}f = \int_{-1}^1 \cos(st)f(t)dt = 0 \quad . \quad (2.7)$$

This should give us pause. We would hope, that for a particular set of measurements g , we could find a unique f that produces that g . Unfortunately, this uniqueness wish is not fulfilled. There are an infinite number of functions f which satisfy $g(s) = 0$ for this kernel. In this case, the null space $\mathcal{N}(\mathcal{K})$ comprises all odd functions. As mentioned above, the null space issue is common in inverse problems – it is important to consider in both continuous and discrete problems. For instance, underdetermined linear systems have a null space that is greater than $\{0\}$.

Another worry with IFKS is their potential to be ill-conditioned or unstable. The difficulty is that often the inverse operator is unbounded. For an operator $T : \mathcal{D}(T) \rightarrow Y$, the operator is *bounded* if⁶⁴

$$\|Tx\| \leq c\|x\| \quad \forall x \in \mathcal{D}(T) \quad . \quad (2.8)$$

Integral operators are bounded, while differentiation operators are unbounded [64, §2.7]. This is not surprising, as integrals tend to smooth or average the variations in their operand functions, while derivatives tend to sharpen them. Unfortunately for the IFK, to invert an integral equation, it is necessary to apply some form of a derivative operator. Thus, this inverse operator is prone to be unbounded. In this context, an

operator is bounded if and only if it is continuous. Continuity is one of the conditions for a well-posed problem, so we see that the inverse IFK is, in general, ill-posed.

2.1.4 Singular value expansion

An important tool in the analysis of IFK is the *singular value expansion* (SVE). Any square-integrable kernel* has an expansion⁴²

$$K(s, t) = \sum_{i=1}^{\infty} \sigma_i u_i(s) v_i(t) \quad . \quad (2.9)$$

The functions u_i and v_i are called the left and right singular functions of the kernel. The $\{u_i\}$ and $\{v_i\}$ form two orthonormal sets, that is,

$$\langle u_i, u_j \rangle = \langle v_i, v_j \rangle = \delta_{ij} \quad \forall i, j \in \mathbb{N} \quad , \quad (2.10)$$

where δ_{ij} is the Kronecker delta and the inner product is defined as

$$\langle f, g \rangle \equiv \int_a^b f(t) g(t) dt \quad .$$

As i increases, the singular functions oscillate more frequently. The singular values σ_i are non-negative real numbers arranged in descending order: $\sigma_1 \geq \sigma_2 \geq \dots \geq 0$. The kernel is called *degenerate* if there are only finitely-many singular values greater than zero.

With the SVE, the fundamental relations of the IFK are

$$\mathcal{K}v_i = \sigma_i u_i \quad , \quad \mathcal{K}^+ u_i = \sigma_i v_i \quad , \quad (2.11)$$

where the integral operator \mathcal{K} is [42, §2.5]

$$[\mathcal{K}f](s) = \int_a^b K(s, t) f(t) dt$$

and \mathcal{K}^+ is the adjoint of \mathcal{K} [64, §3.9-1].

*A kernel is square integrable if $\int_a^b dt \int_a^b ds K(s, t)^2$ is finite.

Expanding f and g in the singular function bases, we have

$$f(t) = \sum_{i=1}^{\infty} \langle v_i, f \rangle v_i(t) \quad , \quad (2.12)$$

and

$$g(s) = \sum_{i=1}^{\infty} \langle u_i, g \rangle u_i(s) \quad . \quad (2.13)$$

Substituting these expansions and the fundamental relation into Equation 2.4, the immediate result is

$$\sum_{i=1}^{\infty} \sigma_i \langle v_i, f \rangle u_i(s) = \sum_{i=1}^{\infty} \langle u_i, g \rangle u_i(s) \quad . \quad (2.14)$$

If the singular values are zero for $i > n$, then a solution exists if and only if $\langle u_i, g \rangle u_i(s) = 0$ for $i > n$. If all singular values are nonzero, then⁴²

$$f(t) = \sum_{i=1}^{\infty} \frac{\langle u_i, g \rangle}{\sigma_i} v_i(t) \quad , \quad (2.15)$$

This solution is only square integrable if

$$\int_a^b dt f(t)^2 = \sum_{i=1}^{\infty} \left(\frac{\langle u_i, g \rangle}{\sigma_i} \right)^2 < \infty \quad . \quad (2.16)$$

This is known as the *Picard condition*. It is a condition on the smoothness of the right-hand side $g(s)$ of Equation 2.4. Thus, a sufficiently non-smooth g can destroy the IFK solution given by Equation 2.15 by making it unbounded, even in the case of continuous and exact data. When g contains errors, this condition becomes even more difficult to satisfy.

2.1.5 Discretization of the IFK

To solve the continuous IFK Equation 2.4 numerically, the equation must be discretized. Two methods are often used to accomplish this, quadrature or expansion methods.

Quadrature methods Here we compute an approximation to $f(t)$ at specified points t_1, t_2, \dots, t_n . Using a numerical quadrature to evaluate the integral in the IFK, we

obtain⁴²

$$g(s) = \int_a^b K(s,t)f(t)dt = \sum_{j=1}^n \omega_j K(s,t_j)f(t_j) \quad . \quad (2.17)$$

Recognizing that in general this cannot be satisfied at all points s , we require this equation to be satisfied at a selected set of points s_1, s_2, \dots, s_m . Then the discrete equations are

$$\sum_{j=1}^n \omega_j K(s_i, t_j) f_j = g_i \quad \text{for } i = 1, \dots, m \quad , \quad (2.18)$$

where $f_j = f(t_j)$ and $g_i = g(s_i)$. This is an $m \times n$ linear system of equations in the discrete unknowns f_j .

Expansion methods The second class of discretization methods approximates f and g in the span of orthogonal basis functions,

$$\tilde{f}(t) = \sum_{j=1}^n f_j x_j(t) \quad , \quad \tilde{g}(s) = \sum_{j=1}^m g_j y_j(s) \quad , \quad (2.19)$$

where $\{f_j\}$ and $\{g_j\}$ are scalars and $\{x_j(t)\}$ and $\{y_j(s)\}$ are the basis functions. Substituting these expansions into the IFK,

$$\sum_{i=1}^m g_i y_i(s) \approx \int_a^b K(s,t) \sum_{j=1}^n f_j x_j(t) dt \quad , \quad (2.20)$$

we see that the discrepancy in this relation will be caused by errors outside of the span $\{y_i(s)\}$. Thus, the inner product of Equation 2.20 with each $y_i(s)$ must be exact:⁴²

$$\langle y_i, g \rangle = \sum_{j=1}^n f_j \left\langle y_i, \int_a^b K(s,t) x_j(t) dt \right\rangle \quad \text{for } i = 1, \dots, m \quad . \quad (2.21)$$

This is an $m \times n$ linear system for f_j , where the elements of A and \vec{b} are computed with the appropriate inner product.

2.2 PROBABILISTIC BAYESIAN METHODS

Given the existence, uniqueness, and stability difficulties in solving IPs, naïvely inverting an operator is a poor strategy in general. Instead, we can treat IPs in a probabilistic sense. Probabilistic methods take as input a range of possible values and the probability of each of the values. In return, the probabilistic method gives a range of possible values for the unknowns. In order to describe these methods, we first review some basics of probability theory.

A proposition a has a probability of being true denoted as $p(a)$ and assumed to be mapped to a real number. The probability that the proposition is false, written as “not a ” or \bar{a} , is $p(\bar{a})$. To begin, we only treat propositions which must be true or false. (Later we will consider probabilities that describe the value of a numerical parameter.) Thus, $p(a) + p(\bar{a})$ represents certainty. If we use the convenient convention that unity represents certainty, then

$$p(a) + p(\bar{a}) = 1 \quad . \quad (2.22)$$

Also, if a is known with certainty to be false, then $p(a) = 0$.

Usually we are interested in describing how one event a affects another event b . This motivates a number of ways to combine the probabilities of observing a and b . One of the most basic ways to combine probabilities is the *Boolean AND*, which is not so different than the general English usage of the conjunction “and.” For our two events a and b , the Boolean AND combination, written $p(a, b)$, is the probability that both a and b are true. This is also called *joint* probability and sometimes written as $p(a \cap b)$. For the Boolean AND, the order of the propositions does not matter, that is, $p(a, b) = p(b, a)$.

A second combined probability is called the *conditional* probability. A conditional probability gives the probability of one proposition given (or conditioned upon) the truthfulness of another proposition. The conditional probability of a on b is the probability of a being true if b is true, denoted $p(a | b)$. In this case, the ordering of a and b does matter, meaning that in general $p(a | b) \neq p(b | a)$.

Relating the joint and conditional probabilities turns out to be important. Using a Venn diagram and the example of an urn with colored balls, it is not difficult to develop the so-called *product rule*,¹⁰⁷

$$p(a, b) = p(a | b)p(b) \quad . \quad (2.23)$$

For a more rigorous derivation, many cite Cox, who arrived at this relation and Equation 2.22 as consequences of preserving logical consistency.¹⁹

Since a and b can be interchanged in Equation 2.23, it follows that

$$p(a | b, c) = \frac{p(b | a, c)p(a | c)}{p(b | c)} . \quad (2.24)$$

This is known as *Bayes' theorem*. Since in practice we are not concerned with $p(b | c)$, we can omit the factor in the denominator, leaving the proportionality expression

$$p(a | b, c) \propto p(b | a, c)p(a | c) . \quad (2.25)$$

We have added the proposition c as a catch-all for background information. Since Bayes' theorem is used frequently, the individual factors have been given names:

- $p(a | c)$ is the *prior*, the information we have about a before considering b ;
- $p(b | a, c)$ is the *likelihood*, which quantifies how likely b is for a given a ;
- $p(a | b, c)$ is the *posterior*, the updated probability for a , accounting for the information in b .

Bayes' theorem is often used in discerning the accuracy of a physical model after receiving new data. In this case, the letters can be replaced with more meaningful terms:

$$p(\text{model} | \text{data}, \text{old data}) \propto p(\text{data} | \text{model}, \text{old data})p(\text{model} | \text{old data}) . \quad (2.26)$$

Proponents of Bayesian inference draw an analogy between this approach and human reasoning in everyday circumstances.⁵⁶ A typical example is the decision to bring an umbrella on a walk. A person relies on both theory and observation (rain comes from dark clouds), and past experience (rain is less common in summer). On a particularly clear day in August, the likelihood probability of rain would be small for lack of clouds, and the prior would be small since it's summer. This is plausible reasoning, in contrast to deductive reasoning.

INSTEAD OF REASONING ABOUT PROPOSITIONS, in this work we are more concerned with estimating the numerical value of quantities, both discrete and continuous. The probability of a continuous variable is most generally expressed as a probability

density function (PDF). For a given quantity x , we define the PDF $f(x)$ such that

$$p(x_1 \leq x < x_2) = \int_{x_1}^{x_2} f(x) dx \quad . \quad (2.27)$$

For convenience, we will write the PDF with the same symbol as we have used for probabilities, $p(x)$. The domain of $p(x)$ is the interval on which x is defined. In general, this is the real line \mathbb{R} , although other intervals can be specified, here denoted as X .

In applications, it is usually impossible to work with the PDF in closed form. One common method to keep some but not all information of the PDF is to work with the moments of the distribution. The zeroth moment is the normalization constraint

$$1 = \int_X p(x) dx \quad . \quad (2.28)$$

The first moment, the mean, provides more useful information. It is defined as

$$\langle x \rangle = \int_X xp(x) dx \quad . \quad (2.29)$$

The second moment about the mean, called the variance, is

$$\langle (x - \langle x \rangle)^2 \rangle = \int_X (x - \langle x \rangle)^2 p(x) dx \quad . \quad (2.30)$$

With higher moments, one could extract more information, but generally the mean and variance are sufficient.

If there are several quantities of interest, say x_1 to x_n , then a multivariate PDF $p(x_1, \dots, x_n)$ is needed to describe the system. Using vector notation, we place all the quantities in a vector $\vec{x} = [x_1, \dots, x_n]^T$ and the PDF is $p(\vec{x})$. The vector \vec{x} lives in the space X_n . The multivariate mean vector is then

$$\langle \vec{x} \rangle = \int_{X_n} \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} p(\vec{x}) dx_1 \cdots dx_n = \int_{X_n} \vec{x} p(\vec{x}) d\vec{x} \quad . \quad (2.31)$$

In addition to the variance of each component

$$\langle (x_i - \langle x_i \rangle)^2 \rangle = \int_{X_n} (x_i - \langle x_i \rangle)^2 p(\vec{x}) d\vec{x} \quad , \quad (2.32)$$

there is also the *covariance* of each component to another component

$$\langle (x_i - \langle x_i \rangle)(x_j - \langle x_j \rangle) \rangle = \int_{X_n} (x_i - \langle x_i \rangle)(x_j - \langle x_j \rangle) p(\vec{x}) d\vec{x} \quad . \quad (2.33)$$

The variances and covariances can be organized in the *covariance matrix*

$$C_x = \langle (\vec{x} - \langle \vec{x} \rangle)(\vec{x} - \langle \vec{x} \rangle)^T \rangle = \int_{X_n} (\vec{x} - \langle \vec{x} \rangle)(\vec{x} - \langle \vec{x} \rangle)^T p(\vec{x}) d\vec{x} \quad . \quad (2.34)$$

2.2.1 Maximum entropy

In applying Bayes theorem, it can be difficult to select the prior PDF that best describes the current state of knowledge. Often when deciding what prior PDF to use, one has a few pieces of information about the distribution, but not enough to fully constrain the PDF. For example, the mean of the distribution may be given, but nothing else. At other times, no information except the interval of possible values is available. How should the PDF be assigned?

This question was answered by the mathematician Claude Shannon in 1948 with the concept of maximum entropy.¹⁰⁴ He derived the form for entropy starting from three conditions for a quantification of uncertainty with discrete probabilities p_1 to p_n :⁵⁵

1. The quantification is continuous with respect to p_i .
2. For a uniform PDF, the quantification increases with increasing n .
3. The quantification is the same for equivalent groupings of p_i .

Since this quantity was identical to entropy in thermodynamics, this quantity has become known as *information entropy*:

$$S = - \sum_i p_i \log_e p_i \quad (2.35)$$

for the discrete case or

$$S = - \int dx p(x) \log_e p(x) \quad (2.36)$$

for continuous data.

To determine the PDF of maximum entropy that fulfills the constraints, Lagrange multipliers are used. If no constraints are given, the maximum entropy PDF is constant

or uniform. When means and covariances are given as constraints, the maximum entropy PDF is a multivariate Gaussian¹⁰⁷

$$p(\vec{x}) = \frac{1}{\sqrt{|2\pi C_x|}} \exp \left[-\frac{1}{2} (x - \langle x \rangle)^T C_x^{-1} (x - \langle x \rangle) \right] \quad (2.37)$$

where $|C_x|$ is the determinant of the covariance matrix of \vec{x} .

2.2.2 Data assimilation procedure

Applying Bayes' theorem for general multivariate PDFs is complicated by the high dimensionality of the space X_n . One way to deal with this is to track only certain moments of the PDFs. Since many types of data and processes can be described by the normal distribution, it makes sense to incorporate the first and second moments, that is, the mean and variance, respectively. Cacuci and Ionescu-Bujor have presented the equations for the Bayesian inference method using the mean and variance.¹³ The following is a summary and time-independent simplification of the data assimilation theory and equations they present.

Assume we have an n_a -vector $\vec{\alpha}$ of model parameters and an n_r -vector \vec{r} containing system responses. With this notation, $\vec{\alpha}$ contains the model inputs of interest and \vec{r} the outputs. Since these values are not known exactly, we consider them as random variables. The parameters have a mean or nominal value of $\vec{\alpha}^0$ and covariance matrix C_α , defined as $\vec{\alpha}^0 = \langle \vec{\alpha} \rangle$ and

$$C_\alpha = \left\langle \left(\vec{\alpha} - \vec{\alpha}^0 \right) \left(\vec{\alpha} - \vec{\alpha}^0 \right)^T \right\rangle , \quad (2.38)$$

respectively.

The responses are both measured and computed. The nominal values for \vec{r} are the measured responses, denoted \vec{r}_m . These measurements have uncertainties summarized in the measurement covariance matrix

$$C_m = \left\langle \left(\vec{r} - \vec{r}_m \right) \left(\vec{r} - \vec{r}_m \right)^T \right\rangle . \quad (2.39)$$

If there are correlations between parameters and responses, these are captured by the

parameter-response covariance matrix

$$C_{\alpha r} = \left\langle \left(\vec{\alpha} - \vec{\alpha}^0 \right) \left(\vec{r} - \vec{r}_m \right)^\top \right\rangle . \quad (2.40)$$

The response vector can be computed via a map $R : \mathbb{R}^{n_a} \rightarrow \mathbb{R}^{n_r}$. For a given $\vec{\alpha}$, the response is $\vec{r} = R(\vec{\alpha})$. The sensitivities of the map R form the matrix S , whose (i, j) element is

$$[S(\vec{\alpha})]_{i,j} = \frac{\partial R_i(\vec{\alpha})}{\partial \alpha_j} . \quad (2.41)$$

For a linearization of the response map,

$$\vec{r} = R(\vec{\alpha}) \approx R(\vec{\alpha}^0) + S(\vec{\alpha}^0) \left(\vec{\alpha} - \vec{\alpha}^0 \right) , \quad (2.42)$$

the expectation value of the response is $\langle \vec{r} \rangle = R(\vec{\alpha}^0)$. The covariance matrix of this computed response is then

$$C_{rc}(\vec{\alpha}^0) \equiv \left\langle \left(\vec{r} - \langle \vec{r} \rangle \right) \left(\vec{r} - \langle \vec{r} \rangle \right)^\top \right\rangle = S(\vec{\alpha}^0) C_\alpha S(\vec{\alpha}^0)^\top . \quad (2.43)$$

With the mean and covariance information, the maximum entropy principle suggests that the most objective PDF for the posterior is the multivariate normal distribution [14, p. 2006]

$$p(\vec{z} | C) = \frac{1}{\sqrt{|2\pi C|}} \exp \left[-\frac{1}{2} Q(\vec{z}) \right] . \quad (2.44)$$

The distribution is a Gaussian if the response map $R(\vec{\alpha})$ is linear. If R is nonlinear, then p will be approximately Gaussian around the minima of functional

$$Q(\vec{z}) \equiv \vec{z}^\top C^{-1} \vec{z} ,$$

appearing in the argument of the exponential in Equation 2.44.

The vector of deviations \vec{z} is defined as

$$\vec{z} \equiv \begin{bmatrix} \vec{\alpha} - \vec{\alpha}^0 \\ \vec{r} - \vec{r}_m \end{bmatrix} = \begin{bmatrix} \vec{z}_\alpha \\ \vec{z}_r \end{bmatrix} ,$$

and the super covariance matrix C is

$$C \equiv \begin{bmatrix} C_{\alpha} & C_{\alpha r} \\ C_{\alpha r}^T & C_m \end{bmatrix} .$$

The functional $Q(\vec{z})$ is similar to a traditional chi-squared metric, but it is generalized to include covariance information.

It is convenient to define an n_r -vector of deviations

$$\vec{d} \equiv R(\vec{\alpha}^0) - \vec{r}_m . \quad (2.45)$$

The covariance matrix of this \vec{d} is

$$C_d(\vec{\alpha}^0) \equiv \langle \vec{d} \vec{d}^T \rangle = C_{rc}(\vec{\alpha}^0) - C_{\alpha r}^T S(\vec{\alpha}^0)^T - S(\vec{\alpha}^0) C_{\alpha r} + C_m . \quad (2.46)$$

We also need the parameter-deviation covariance matrix

$$C_{\alpha d}(\vec{\alpha}^0) \equiv \langle (\vec{\alpha} - \vec{\alpha}^0) \vec{d}^T \rangle = C_{\alpha r} - C_{\alpha} S(\vec{\alpha}^0)^T \quad (2.47)$$

and the response-deviation covariance matrix

$$C_{rd}(\vec{\alpha}^0) \equiv \langle (\vec{r} - \vec{r}_m) \vec{d}^T \rangle = C_m - C_{\alpha r}^T S(\vec{\alpha}^0)^T . \quad (2.48)$$

Putting all this moment information in Bayes' theorem, the most probable values, or the "best estimate," for the parameters is

$$\vec{\alpha}^{be} = \vec{\alpha}^0 + \left(C_{\alpha r} - C_{\alpha} S(\vec{\alpha}^0)^T \right) C_d(\vec{\alpha}^0)^{-1} \vec{d} \quad (2.49)$$

and

$$\vec{r}^{be} = \vec{r}_m + \left(C_m - C_{\alpha r}^T S(\vec{\alpha}^0)^T \right) C_d(\vec{\alpha}^0)^{-1} \vec{d} \quad (2.50)$$

for the responses.

The best estimates for the parameter, response, and parameter-response covariance

matrices are

$$\mathbf{C}_\alpha^{be} = \mathbf{C}_\alpha - \mathbf{C}_{\alpha d} \left(\vec{\alpha}^0 \right) \mathbf{C}_d \left(\vec{\alpha}^0 \right)^{-1} \mathbf{C}_{\alpha d} \left(\vec{\alpha}^0 \right)^\top, \quad (2.51)$$

$$\mathbf{C}_r^{be} = \mathbf{C}_m - \mathbf{C}_{rd} \left(\vec{\alpha}^0 \right) \mathbf{C}_d \left(\vec{\alpha}^0 \right)^{-1} \mathbf{C}_{rd} \left(\vec{\alpha}^0 \right)^\top, \quad \text{and} \quad (2.52)$$

$$\mathbf{C}_{\alpha r}^{be} = \mathbf{C}_{\alpha r} - \mathbf{C}_{\alpha d} \left(\vec{\alpha}^0 \right) \mathbf{C}_d \left(\vec{\alpha}^0 \right)^{-1} \mathbf{C}_{rd} \left(\vec{\alpha}^0 \right)^\top. \quad (2.53)$$

Finally, a commonly-used indicator of consistency is the χ^2 value, which quantifies the deviations of the measurements from the computed predictions. If the predictions are too far from the measurements (relative to the variance of each quantity), then the χ^2 value will be high and the conclusion is that some of the measurements are in error or the model is faulty. In terms of the vectors and matrices already defined, the χ^2 value is

$$\chi^2 = \vec{d}^\top \mathbf{C}_d \left(\vec{\alpha}^0 \right)^{-1} \vec{d}. \quad (2.54)$$

This χ^2 has n_r , the number of responses, degrees of freedom. The χ^2/n_r -distribution has mean one and variance $2/n_r$, so it is often more helpful to report the χ^2 value as χ^2/n_r . Then, if χ^2/n_r is close to one, the measurements and predictions are consistent. If not, there is reason to be suspicious. For $\chi^2/n_r < 1$, the agreement is better than would be expected, whereas the agreement is worse than expected if $\chi^2/n_r > 1$.

2.2.3 Solving nonlinear problems

If $R(\vec{\alpha})$ is nonlinear, then the posterior is no longer Gaussian, so it cannot be maximized in one step. Instead, we must find other methods to maximize the posterior distribution $p(\vec{z} \mid \mathbf{C})$ or minimize $Q(\vec{z})$, these two being equivalent. There are many ways to find the extremal points of these functionals. One method which has achieved significant success is Markov Chain Monte Carlo [112, §2.3.5], which uses the Metropolis-Hastings algorithm.^{43,82,83} Monte Carlo methods are most appropriate for highly nonlinear models, where the nonlinearities cause many local maxima in the posterior. Unfortunately, Monte Carlo methods tend to be slower, so we have not pursued these methods here.

For models that are mildly nonlinear, an iterative approach with a series of linearizations is possible. At a local minimum, the shape of the functional can be approximated with a (local) covariance matrix. If there are a handful of local minima, each could be

visited and approximated with its own covariance matrix. Tarantola discusses this in §3.2.3.¹¹² In taking this approach, we recognize that the posterior is not a multivariate Gaussian; however, we approximate it as Gaussian in the neighborhood of the local minimum.

A common approach to minimize $Q(\vec{z})$ is to use the linear update process in an iterative fashion, for which the posterior means and covariances are passed to the next iteration as the prior. However, as Rodgers points out [97, §5.6.2], this is not a prudent approach because it confuses the role of the iteration state and the prior distribution. In an iterated linear solution, the prior distribution may be forgotten as the iteration progresses.

Instead, we minimize Q using a Newton optimization method. Additionally, this takes advantage of the first and second derivative information of the functional, which Monte Carlo methods generally ignore. For Newton's method, the functional's gradient $\vec{\nabla}_\alpha Q(\vec{z}_k)$ and Hessian $\nabla_\alpha^2 Q(\vec{z}_k)$ are required. The Newton update step is

$$\vec{\alpha}_{k+1} = \vec{\alpha}_k - \lambda_k \left(\nabla_\alpha^2 Q(\vec{z}_k) \right)^{-1} \vec{\nabla}_\alpha Q(\vec{z}_k) \quad . \quad (2.55)$$

The line search parameter $\lambda_k \in (0, 1]$ adjusts the step length. The initial iterate $\vec{\alpha}_0$ is often set to the mean of the prior $\vec{\alpha}^0$. At each iteration $\vec{\alpha}_k$, the sensitivities are recomputed.

To compute the Hessian and gradient, it is helpful to partition the inverse covariance matrix as

$$\mathbf{C}^{-1} = \begin{bmatrix} \mathbf{C}_1 & \mathbf{C}_3 \\ \mathbf{C}_3^T & \mathbf{C}_2 \end{bmatrix} \quad , \quad (2.56)$$

where the three submatrices are

$$\begin{aligned} \mathbf{C}_1 &= \mathbf{S}_c^{-1} \quad , \\ \mathbf{C}_2 &= \mathbf{C}_m^{-1} + \mathbf{C}_m^{-1} \mathbf{C}_{\alpha r}^T \mathbf{S}_c^{-1} \mathbf{C}_{\alpha r} \mathbf{C}_m^{-1} \quad , \\ \mathbf{C}_3 &= -\mathbf{S}_c^{-1} \mathbf{C}_{\alpha r} \mathbf{C}_m^{-1} \quad . \end{aligned}$$

The Schur complement⁸⁵ is $\mathbf{S}_c = \mathbf{C}_\alpha - \mathbf{C}_{\alpha r} \mathbf{C}_m^{-1} \mathbf{C}_{\alpha r}^T$. These expressions can be verified by performing the blockwise multiplication of $\mathbf{C}\mathbf{C}^{-1}$ to show that $\mathbf{C}\mathbf{C}^{-1} = \mathbf{I}$, where \mathbf{I} is the identity matrix.

In terms of these three block matrices, the functional Q is

$$Q(\vec{z}) = \vec{z}_\alpha^T C_1 \vec{z}_\alpha + 2\vec{z}_\alpha^T C_3 \vec{z}_r + \vec{z}_r^T C_2 \vec{z}_r \quad . \quad (2.57)$$

The gradient of Q is then

$$\vec{\nabla}_\alpha Q(\vec{z}) = \left. \frac{\partial Q}{\partial \vec{\alpha}} \right|_{\vec{z}} = C_1 \vec{z}_\alpha + S^T C_2 \vec{z}_r + C_3 \vec{z}_r + S^T C_3^T \vec{z}_\alpha \quad . \quad (2.58)$$

Taking the partial derivative of the gradient with respect to $\vec{\alpha}$, the Hessian is

$$\nabla_\alpha^2 Q(\vec{z}) = \left. \frac{\partial^2 Q}{\partial \vec{\alpha}^2} \right|_{\vec{z}} = C_1 + S^T C_2 S + C_3 S + S^T C_3^T + \frac{\partial S}{\partial \vec{\alpha}} C_2 \vec{z}_r + \frac{\partial S}{\partial \vec{\alpha}} C_3^T \vec{z}_\alpha \quad . \quad (2.59)$$

In the context of nonlinear least squares problems, the typical approach is to ignore the second-order derivatives of \vec{R} (the $\frac{\partial S}{\partial \vec{\alpha}}$ factor), resulting in the Gauss-Newton method [61, §2.4]. This approach is called a quasi-Newton method by Tarantola [112, §3.4.4]. At the optimal point $\vec{\alpha}^*$, the best-estimate covariance matrix C^{be} can be computed with Cacuci's formulas, provided the sensitivities are evaluated at $\vec{\alpha}^*$.

2.3 DISCRETE INVERSE THEORY

In parallel with the Bayesian approach, many researchers have developed methods and tools to compensate for the ill-posed nature of IPS. Born of necessity and based on the analysis of the instabilities of IPS, as demonstrated by the SVE, researchers have developed a number of methods to resolve discrete IFK problems. These methods change the problem slightly in an effort to preserve useful parts of the solution while minimizing the bad parts. They are primarily intended for linear systems,

$$A\vec{x} = \vec{b} \quad , \quad A \in \mathbb{R}^{m \times n} \quad , \quad \vec{x} \in \mathbb{R}^n \quad , \quad \vec{b} \in \mathbb{R}^m \quad , \quad (2.60)$$

where \vec{x} is unknown and A and \vec{b} are given.

These concepts and methods are described in a number of works, for example, books by Hansen^{41,42}, Vogel¹¹⁷, Santamarina¹⁰¹, and Groetsch.³⁹

As mentioned above, the IFK is ill-conditioned because the inverse operator

is unbounded and not continuous. When the integral equation is discretized, the operators become finite dimensional. Every finite-dimensional operator is bounded.⁶⁴ This is a favorable consequence of the discretization. However, as the discretization level or order is refined (n and m increase), the ill-conditioning of the problem worsens and the discrete problem better approximates the poor properties of the continuous one. This is manifested by an increase in the condition number of A , $\kappa(A) = \|A\| \|A\|^{-1}$.⁵² Thus the solution can actually degrade as the mesh is refined, behavior that is opposite to conventional wisdom in the numerical solution of PDES.

2.3.1 Regularization

Since IPs are ill-conditioned, one approach to the resolution of an IP is to replace it with a nearby problem that is well-conditioned. The process of modifying a problem to improve its conditioning is known as *regularization*. A wide variety of regularization techniques have been developed, including truncated singular value decomposition, Tikhonov regularization, and iterative methods. Usually regularization methods force the solution to be smooth.⁴²

These methods are best analyzed in terms of the singular value decomposition (SVD), the discrete analog of the SVE. The SVD of a matrix $A \in \mathbb{R}^{m \times n}$ is

$$A = U \Sigma V^T \quad . \quad (2.61)$$

The diagonal matrix $\Sigma \in \mathbb{R}^{m \times n}$ comprises the singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{\min\{m,n\}} \geq 0$. The unitary matrices $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ contain the right and left singular vectors \vec{u}_i and \vec{v}_i of A .⁸⁵

In terms of the SVD, the generalized inverse of A , called the pseudoinverse, is

$$A^\dagger = \sum_{i=1}^r \frac{\vec{v}_i \vec{u}_i^T}{\sigma_i} \quad ,$$

where r is the rank[†] of Σ .

Of interest for discrete inverse problems is the *discrete Picard condition*, which states that $|\vec{u}_i^T \vec{b}|$ must decay faster with increasing i than σ_i for as long as σ_i are not polluted with round-off error. The discrete Picard condition should be satisfied to

[†]The rank of a matrix is the number of linearly independent rows or columns, whichever is larger.

obtain a useful solution. Clearly, this is the discrete analogue of the Picard condition expressed in Equation 2.16. A noisy right-hand-side \vec{b} prematurely ends the decay of the coefficients $|\vec{u}_i^T \vec{b}|$.

Truncated svd The truncated singular value decomposition (TSVD) regularization method attempts to preserve the right singular vectors \vec{v}_i associated with the indices satisfying the discrete Picard condition, while excluding the singular vectors and values associated with higher i indices.⁴² Unlike other methods that gradually filter the higher singular values, TSVD excludes all singular values with indices greater than k , so that the solution is

$$A^+ \vec{b} = \sum_{i=1}^k \frac{\vec{u}_i^T \vec{b}}{\sigma_i} \vec{v}_i \quad ,$$

The goal is to include the information coming from the useful signal in \vec{b} while excluding the noisy information. Unlike the common computation of the pseudoinverse (such as `pinv` in Matlab) that determines the index k based on the magnitude of the singular values σ_i , in TSVD for IPS, this decision is governed by the rate of decay of $\vec{u}_i^T \vec{b}$.

Tikhonov regularization This method is usually written in terms of the least squares problem

$$\vec{x}_\lambda = \arg \min_{\vec{x}} \|\mathbf{A}\vec{x} - \vec{b}\|_2^2 + \lambda^2 \|\vec{x}\|_2^2 \quad , \quad (2.62)$$

where the regularization parameter λ allows the user to selectively filter portions of the noise. In terms of the svd, Tikhonov regularization is

$$\vec{x}_\lambda = \sum_{i=1}^n \frac{\sigma_i^2}{\sigma_i^2 + \lambda^2} \frac{\vec{u}_i^T \vec{b}}{\sigma_i} \vec{v}_i \quad .$$

In contrast to the TSVD approach, Tikhonov regularization gradually filters higher singular values with the filter factor $\frac{\sigma_i^2}{\sigma_i^2 + \lambda^2}$.

Choosing the value of the regularization parameter λ is an important part of Tikhonov regularization. Various strategies exist for this, including the Morozov discrepancy principle, the L-curve, and generalized cross validation.⁴²

Iterative methods For large systems, `rsvd` or Tikhonov regularization may become intractable. In this case, iterative methods are often the only choice. Useful iterative methods for discrete `IPs` exhibit *semiconvergence*, that is, the iterates \vec{x}_k tend to get closer to the exact solution \vec{x}_{exact} for a while, but then in further iterations the iterates diverge from the desired solution and instead converge to the unregularized solution $A^{-1}\vec{b}$.⁴²

Landweber iteration is a well-known method,

$$\vec{x}_{k+1} = \vec{x}_k + \omega A^T (\vec{b} - A\vec{x}_k) \quad , \quad (2.63)$$

where the weight ω is a relaxation parameter. This method also has a decaying filter factor in the `svd` expression of the solution. Although the Landweber method can be competitive with other methods in certain applications, for example, tomography, it is in general much slower to converge than other methods. A related method with regularizing properties, Kaczmarz's method, uses a similar iterative equation but only updates one component of \vec{x} per iteration by applying one row of A .

One of the most powerful regularizing iterative methods is Conjugate Gradient Least Squares (`CGLS`), a modified conjugate gradient method for the normal equations $A^T A \vec{x} = A^T \vec{b}$. The `CGLS` solution minimizes the error

$$\vec{x}_k = \arg \min_x \|A\vec{x} - \vec{b}\|_2 \quad (2.64)$$

such that

$$\vec{x} \in \text{span}\{A^T \vec{b}, (A^T A)A^T \vec{b}, \dots, (A^T A)^{k-1}A^T \vec{b}\} \quad . \quad (2.65)$$

The `CGLS` method does exhibit semiconvergence, and it also has a filtered `svd` solution, although the filter factors are a bit more complicated than the other methods presented here.⁴²

2.4 INVERSE CRIMES

Before concluding this chapter, we must mention a practice that is often cautioned against – the inverse crime – that involves a common shortcut taken by method developers. The ultimate goal in an inverse problem is to accurately estimate system parameters based on physical measurements using a numerical model of the sys-

tem. However, while developing and implementing methods for inverse problems, researchers commonly avoid cumbersome experiments by using a mathematical model to compute the expected values of the measurements. While this practice is not inherently wrong, using the same model for the “measurements” and the inverse procedure can make the procedure appear to perform extraordinarily, and often unrealistically, well, even if the method would fail for actual measurements.

In the inverse problem community, the term *inverse crime* describes the use of the identical numerical model for both the forward and inverse solutions.⁴² Although success of a method for an inverse-crime problem is a necessary condition for success using real measurements, it is not a sufficient condition. Thus, while we do commit a number of inverse crimes to test the initial viability of the method, we also include tests that do not use the same forward and inverse models to test our approach in configurations free of inverse crimes. Finally, note that the severest inverse crime is committed by using the same discretization mesh for the measurements and the inverse solution (in our case, the mesh comprises the set of discrete ordinates in the angular quadrature, energy groups, and spatial mesh cells), while using different meshes with the same numerical model is merely an “inverse misdemeanor.”

The solution of inverse radiation transport problems is not new. It is natural to wonder what information can be inferred from the radiation emitted from an object. In this chapter, we first review the conservation equation for neutral particles moving in a medium. Next we discuss the numerical methods commonly used to solve this conservation equation. This model for the transport of radiation is a necessary building block for any successful inverse transport solution. In the second part of this chapter, we review some of the prior inverse radiation transport problems that appear in the literature.

3.1 RADIATION TRANSPORT EQUATIONS

To solve an inverse problem, one should have a model of the observed system. In our case, this mathematical model is the radiation transport equation, namely the time-independent linear Boltzmann transport equation for the conservation of neutral particles in non-multiplying media,⁷²

$$\hat{\Omega} \cdot \vec{\nabla} \psi(\vec{x}, E, \hat{\Omega}) + \sigma(\vec{x}, E) \psi(\vec{x}, E, \hat{\Omega}) = \int dE' \int d\hat{\Omega}' \sigma_s(\vec{x}; E', \hat{\Omega}' \rightarrow E, \hat{\Omega}) \psi(\vec{x}, E', \hat{\Omega}') + q(\vec{x}, E, \hat{\Omega}) \quad , \quad (3.1)$$

for

$$x \in V \quad , \quad \hat{\Omega} \in 4\pi \quad , \quad E \in [0, \infty) \quad ,$$

and with boundary conditions

$$\psi(\vec{x}, E, \hat{\Omega}) = \psi_0(\vec{x}, E, \hat{\Omega}) \quad \text{for } \vec{x} \in \partial V \text{ and } \hat{\Omega} \cdot \hat{n} < 0 \quad ,$$

where V is the volume of interest, ∂V is the boundary, \hat{n} is the surface normal pointing out of V , and ψ_0 is the prescribed incoming flux.

The quantities in this equation are

$$\begin{aligned} \vec{x} &\equiv \text{position vector in three dimensions,} \\ E &\equiv \text{particle energy,} \\ \hat{\Omega} &\equiv \text{unit direction vector along the particle's} \\ &\quad \text{direction of motion in three dimensions,} \\ \sigma(\vec{x}, E) &\equiv \text{macroscopic total cross section,} \\ \sigma_s(\vec{x}; E', \hat{\Omega}' \rightarrow E, \hat{\Omega}) &\equiv \text{double differential macroscopic scattering} \\ &\quad \text{cross section,} \\ \psi(\vec{x}, E, \hat{\Omega}) &\equiv \text{angular flux, and} \\ q(\vec{x}, E, \hat{\Omega}) &\equiv \text{fixed source of particles distributed} \\ &\quad \text{per unit time and volume.} \end{aligned}$$

The macroscopic cross sections represent the probability of interaction per unit path-length traveled, any interaction in the case of the total cross section σ and a scattering interaction from a certain incident state $(E', \hat{\Omega}')$ to another scattered state $(E, \hat{\Omega})$ for the scattering cross section $\sigma_s(\vec{x}; E', \hat{\Omega}' \rightarrow E, \hat{\Omega})$. The source term $q \neq 0$ in Equation 3.1 and boundary conditions $\psi_0 \neq 0$ make the equation inhomogeneous. The angular flux is defined as

$$\begin{aligned} \frac{1}{v} \psi(\vec{x}, E, \hat{\Omega}) d\vec{x} d\hat{\Omega} dE &\equiv \text{the number of particles} \\ &\quad \text{in the volume } d\vec{x} \text{ centered at } \vec{x} \\ &\quad \text{in the energy interval } dE \text{ centered at } E \\ &\quad \text{in the solid angle } d\hat{\Omega} \text{ centered at } \hat{\Omega} \text{ ,} \end{aligned}$$

where v is the particle speed. The angle-integrated angular flux, called the *scalar flux*,

$$\phi(\vec{x}, E) = \int_{4\pi} d\hat{\Omega} \psi(\vec{x}, E, \hat{\Omega}) \text{ ,} \quad (3.2)$$

is important in computing reaction rates. For example, a detector response could be

computed as¹⁰⁶

$$r = \int_0^\infty dE \int_V d\vec{x} \sigma_d(\vec{x}, E) \phi(\vec{x}, E) \quad , \quad (3.3)$$

where $\sigma_d(\vec{x}, E)$ is the detector response function (DRF). However, for a full-spectrum response, the DRF is $\sigma_d(\vec{x}, E', E)$, where E' is the energy deposited in the detector.³⁵ Then the energy dependent response is

$$r(E') = \int_0^\infty dE \int_V d\vec{x} \sigma_d(\vec{x}, E', E) \phi(\vec{x}, E) \quad . \quad (3.4)$$

3.1.1 Adjoint transport equation

Whether for efficiently computing detector responses or for computing sensitivities, the adjoint transport equation is important. The equation can be derived in a number of ways. One could begin with the definition of the Hilbert-adjoint operator and use integration by parts to work out the adjoint transport operator, the approach taken in the text by Lewis and Miller.⁷² Bell and Glasstone derive the equation as a conservation of importance, which is one physical interpretation of the adjoint.⁴ A simple approach is to apply rules for derivative and integral terms: the sign of odd-order derivatives flips and the sign stays the same for even-order derivatives, and the kernels of integral terms have their arguments transposed.

Deriving the adjoint operator from the definition of the adjoint is preferable to make the boundary conditions apparent. In Equation F.4, if x and y are from the same space, then the adjoint T^+ to T is defined as

$$\langle Tx, y \rangle = \langle x, T^+y \rangle \quad . \quad (3.5)$$

It is important to define the inner product. Since we are usually interested in a reaction rate (as in Equation 3.3), a natural choice for the inner product is

$$\langle f, g \rangle = \int_{4\pi} d\hat{\Omega} \int_0^\infty dE \int_V dV f(\vec{x}, E, \hat{\Omega}) g(\vec{x}, E, \hat{\Omega}) \quad . \quad (3.6)$$

From Equation 3.1, the forward transport operator is

$$L\psi = \hat{\Omega} \cdot \vec{\nabla} \psi(\vec{x}, E, \hat{\Omega}) + \sigma(\vec{x}, E) \psi(\vec{x}, E, \hat{\Omega}) - \int dE' \int d\hat{\Omega}' \sigma_s(\vec{x}; E', \hat{\Omega}' \rightarrow E, \hat{\Omega}) \psi(\vec{x}, E', \hat{\Omega}') \quad , \quad (3.7)$$

Since the inner product is linear in the first and the second arguments for real numbers, we can treat each of the terms in L in turn. For the streaming operator,

$$L_1\psi = \hat{\Omega} \cdot \vec{\nabla} \psi(\vec{x}, E, \hat{\Omega}) \quad , \quad (3.8)$$

we compute the adjoint operator with an arbitrary function ξ :

$$\begin{aligned} \langle L_1\psi, \xi \rangle &= \int_{4\pi} d\hat{\Omega} \int_0^\infty dE \int_V dV \left(\hat{\Omega} \cdot \vec{\nabla} \psi(\vec{x}, E, \hat{\Omega}) \right) \xi(\vec{x}, E, \hat{\Omega}) \\ &= \int_{4\pi} d\hat{\Omega} \int_0^\infty dE \hat{\Omega} \cdot \int_V dV \left(\vec{\nabla} \psi(\vec{x}, E, \hat{\Omega}) \right) \xi(\vec{x}, E, \hat{\Omega}) \\ &= \int_{4\pi} d\hat{\Omega} \int_0^\infty dE \hat{\Omega} \cdot \left[\int_{\partial V} dS \hat{n} \psi(\vec{x}, E, \hat{\Omega}) \xi(\vec{x}, E, \hat{\Omega}) \right. \\ &\quad \left. - \int_V dV \psi(\vec{x}, E, \hat{\Omega}) \vec{\nabla} \xi(\vec{x}, E, \hat{\Omega}) \right] \\ &= \langle \psi, -\hat{\Omega} \cdot \vec{\nabla} \xi \rangle \\ &\quad + \int_{4\pi} d\hat{\Omega} \int_0^\infty dE \int_{\partial V} dS \hat{\Omega} \cdot \hat{n} \psi(\vec{x}, E, \hat{\Omega}) \xi(\vec{x}, E, \hat{\Omega}) \quad . \end{aligned}$$

Here we have used integration by parts to switch the gradient from acting upon ψ to ξ .^{*} We see in this expression that a boundary term is left over.

^{*}Integration by parts in multiple dimensions for a closed volume V with surface ∂V is

$$\int_V f \vec{\nabla} g dV = \int_{\partial V} \hat{n} f g dS - \int_V g \vec{\nabla} f dV \quad .$$

This expression is a consequence of the divergence theorem and the chain rule for derivatives. The divergence theorem for a vector-valued function \vec{F} is¹¹¹

$$\int_V \vec{\nabla} \cdot \vec{F} dV = \int_{\partial V} \vec{F} \cdot \hat{n} dS \quad .$$

If $\vec{F} = f\vec{c}$, with \vec{c} a constant vector, then the divergence theorem reduces to¹¹⁹

$$\int_V \vec{\nabla} f dV = \int_{\partial V} \hat{n} f dS \quad . \quad (3.9)$$

The second term in L is much simpler:

$$L_2\psi = \sigma(\vec{x}, E)\psi(\vec{x}, E, \hat{\Omega}) \quad ,$$

where σ is a real-valued function. The adjoint of this term is trivial, $L_2^+ = \sigma$, since

$$\langle L_2\psi, \xi \rangle = \langle \sigma\psi, \xi \rangle = \langle \psi, \sigma\xi \rangle \quad .$$

The third term in the forward operator is

$$L_3\psi = - \int dE' \int d\hat{\Omega}' \sigma_s(\vec{x}; E', \hat{\Omega}' \rightarrow E, \hat{\Omega})\psi(\vec{x}, E', \hat{\Omega}') \quad .$$

The adjoint operator for L_3 is then

$$\begin{aligned} \langle L_3\psi, \xi \rangle &= - \int_V dV \int_{4\pi} d\hat{\Omega} \int_0^\infty dE \\ &\quad \times \int_0^\infty dE' \int_{4\pi} d\hat{\Omega}' \sigma_s(\vec{x}; E', \hat{\Omega}' \rightarrow E, \hat{\Omega})\psi(\vec{x}, E', \hat{\Omega}')\xi(\vec{x}, E, \hat{\Omega}) \\ &= - \int_V dV \int_{4\pi} d\hat{\Omega} \int_0^\infty dE \psi(\vec{x}, E, \hat{\Omega}) \\ &\quad \times \int_0^\infty dE' \int_{4\pi} d\hat{\Omega}' \sigma_s(\vec{x}; E, \hat{\Omega} \rightarrow E', \hat{\Omega}')\xi(\vec{x}, E', \hat{\Omega}') \quad , \end{aligned}$$

where the dummy integration variables for energy and direction have been interchanged to extract ψ . Notice that the arguments to the scattering kernel have thus been flipped. Thus the adjoint to L_3 is

$$L_3^+\xi = - \int dE' \int d\hat{\Omega}' \sigma_s(\vec{x}; E, \hat{\Omega} \rightarrow E', \hat{\Omega}')\xi(\vec{x}, E', \hat{\Omega}') \quad .$$

Combining these expressions, the formal adjoint operator is

$$\begin{aligned} L^+\xi &= -\hat{\Omega} \cdot \vec{\nabla}\xi(\vec{x}, E, \hat{\Omega}) + \sigma(\vec{x}, E)\xi(\vec{x}, E, \hat{\Omega}) \\ &\quad - \int dE' \int d\hat{\Omega}' \sigma_s(\vec{x}; E, \hat{\Omega} \rightarrow E', \hat{\Omega}')\xi(\vec{x}, E', \hat{\Omega}') \quad . \quad (3.10) \end{aligned}$$

The chain rule gives

$$\vec{\nabla}(fg) = f\vec{\nabla}g + g\vec{\nabla}f \quad .$$

Integrating this and applying Equation 3.9 gives the equation for integration by parts.

Note the word *formal*. This is not the general adjoint operator. In operator notation, the inner product can be written

$$\langle L\psi, \tilde{\zeta} \rangle = \langle \psi, L^+\tilde{\zeta} \rangle + P[\psi, \tilde{\zeta}] \quad , \quad (3.11)$$

where the bilinear concomitant, evaluated on the surface of the volume V , is

$$P[\psi, \tilde{\zeta}] = \int_{4\pi} d\hat{\Omega} \int_0^\infty dE \int_{\partial V} dS \hat{\Omega} \cdot \hat{n} \psi(\vec{x}, E, \hat{\Omega}) \tilde{\zeta}(\vec{x}, E, \hat{\Omega}) \quad . \quad (3.12)$$

With vacuum boundary conditions, the concomitant P can vanish. The forward vacuum boundary conditions are

$$\psi(\vec{x}, E, \hat{\Omega}) = 0 \quad \text{for} \quad \vec{x} \in \partial V \quad \text{and} \quad \hat{\Omega} \cdot \hat{n} < 0 \quad , \quad (3.13)$$

which states that no particles enter the volume from outside. Then we impose on $\tilde{\zeta}$ the boundary condition

$$\tilde{\zeta}(\vec{x}, E, \hat{\Omega}) = 0 \quad \text{for} \quad \vec{x} \in \partial V \quad \text{and} \quad \hat{\Omega} \cdot \hat{n} > 0 \quad . \quad (3.14)$$

Interpreting $\tilde{\zeta}$ as an importance, this condition requires that particles leaving the volume have no importance, which makes intuitive sense. With these conditions, the integral over angle in the bilinear concomitant equals zero, and thus $P[\psi, \tilde{\zeta}] = 0$.

Now we can see a use for the adjoint equation. Consider the source-driven forward problem, written in operator notation as

$$L\psi = q \quad . \quad (3.15)$$

Taking the inner product of Equation 3.15 with the arbitrary function $\tilde{\zeta}$, we have

$$\langle L\psi, \tilde{\zeta} \rangle = \langle q, \tilde{\zeta} \rangle \quad . \quad (3.16)$$

Substituting for $\langle L\psi, \tilde{\zeta} \rangle$ from Equation 3.11 into this expression, we obtain

$$\langle \psi, L^+\tilde{\zeta} \rangle = \langle q, \tilde{\zeta} \rangle - P[\psi, \tilde{\zeta}] \quad . \quad (3.17)$$

The left hand side of Equation 3.17 is the inner product of the forward flux with some

unspecified function. Since we are interested in responses of the form $\langle \psi, \sigma_d \rangle$, with σ_d a detector response function, it is reasonable to require that

$$L^+ \psi^+ = \sigma_d \quad , \quad (3.18)$$

where we have switched from ζ to ψ^+ since the function is no longer arbitrary. Then, a response can be computed with the adjoint as

$$\langle \psi, \sigma_d \rangle = \langle q, \psi^+ \rangle - P[\psi, \psi^+] \quad . \quad (3.19)$$

Finally, if the appropriate vacuum boundary conditions on the forward and adjoint problem are imposed, then the response is simply

$$\langle \psi, \sigma_d \rangle = \langle q, \psi^+ \rangle \quad . \quad (3.20)$$

In summary, with boundary conditions specified in Equation 3.13 and Equation 3.14, the adjoint equation to Equation 3.1 is

$$-\hat{\Omega} \cdot \vec{\nabla} \psi^+(\vec{x}, E, \hat{\Omega}) + \sigma(\vec{x}, E) \psi^+(\vec{x}, E, \hat{\Omega}) = \int dE' \int d\hat{\Omega}' \sigma_s(\vec{x}; E, \hat{\Omega} \rightarrow E', \hat{\Omega}') \psi^+(\vec{x}, E', \hat{\Omega}') + q^+(\vec{x}, E, \hat{\Omega}) \quad , \quad (3.21)$$

where commonly $q^+ = \sigma_d$.

3.1.2 Numerical methods in transport

For most practical applications, the forward and adjoint radiation transport equations are solved numerically. There are a variety of numerical approaches; here we briefly outline the deterministic methods used in this work. The energy, angle, and spatial variables must be discretized.

Energy We use the standard multigroup approximation, in which all energy-dependent quantities are integrated over a set of smaller ranges of energies.²⁴ For each energy group, the group cross sections are considered group averages, while the group flux is a total flux, that is, the flux integrated over the energy interval.

Angle A common method to discretize the angular variable is the discrete ordinates approach.⁷² The single transport equation is replaced with a set of equations, one for each angle in a predetermined quadrature set. Thus, the angular flux is only known in certain directions. The scalar flux, as well as higher angular moments, is obtained by a quadrature summation instead of an integration of the angular flux. The discrete ordinates approach can be susceptible to unphysical oscillations in the solution, called *ray effects*. These can be reduced by increasing the number of discrete directions in the employed quadrature. However, with a localized radiation source, ray effects can be noticeable at large distances from the source even with a high-order quadrature.

A common strategy to reduce the effect of these oscillations is the first-collision source method.²⁷ In this method, the angular flux is divided into two components:

$$\psi = \psi_{\text{unc}} + \psi_{\text{col}} \quad ,$$

the uncollided component ψ_{unc} and the collided component ψ_{col} . A similar relation exists for the scalar flux. The uncollided flux is that portion of the flux that has been emitted by a source but has experienced no subsequent collisions. The collided flux, on the other hand, has had one or more interactions with the host medium. The advantage of this approach is the simplicity of the expression for ψ_{unc} , which is a closed-form equation involving material and geometric attenuation. The uncollided flux can be computed accurately at a point using ray tracing. This eliminates ray effects in the uncollided flux. Then the uncollided flux is multiplied by the scattering cross section in each cell to give the so-called first-collision source, that is, the scattering source of particles that come from the first scatter. This first-collision source is given as the source term to the discrete-ordinates solver to compute the collided flux ψ_{col} . The collided flux will still contain ray effects, but they will be less noticeable because the radiation source is distributed throughout the space.

Space There are a variety of spatial discretization methods in the literature, including finite difference, finite volume, and finite element schemes.²⁷ Some of these are linear methods, while others add non-linear factors to increase the order of convergence to the exact solution.

Discrete matrix equations The discretized transport equation can be written in matrix form. Evans presents a compact notation for formulating these equations.²⁷ We follow his notation here. Using the multigroup energy treatment, discrete ordinate for angle, and Legendre moments to approximate scattering, the discrete transport equation can be written as

$$L\vec{\psi} = MS\vec{\phi} + \vec{q}_a \quad . \quad (3.22)$$

The discretizations are made with the following parameters:

- n_g = number of groups,
- n_t = number of Legendre scattering moments,
- n_o = number of discrete ordinate directions,
- n_c = number of spatial cells, and
- n_e = number of unknowns per cells.

If

$$a = n_g \cdot n_o \cdot n_c \cdot n_e \quad , \quad f = n_g \cdot n_t \cdot n_c \cdot n_e \quad ,$$

then the size of Equation 3.22

$$(a \times a)(a \times 1) = (a \times f)(f \times f)(f \times 1) + (a \times 1) \quad .$$

The flux angular moments can be computed from the angular flux by a quadrature rule

$$\vec{\phi} = D\vec{\psi} \quad ,$$

where D has size $f \times a$ and $\vec{\phi}$ has size f .

For our purposes, the angular flux is not of primary concern. Rather, we are primarily concerned with the flux moments. Thus, the transport matrix equation can be rearranged as

$$(DLD^{-1} - S)\vec{\phi} = \vec{q}_s$$

where $\vec{q}_s = D\vec{q}_a$ with size f . This particular form requires that we assume a Galerkin quadrature such that $D = M^{-1}$.⁸⁹ Setting $L_s = DLM - S$, the equations can be written

$$L_s\vec{\phi} = \vec{q}_s \quad . \quad (3.23)$$

The adjoint equations would then be $L_s^+ \vec{\phi}^+ = \vec{q}_s^+$.

3.2 INVERSE RADIATION TRANSPORT PROBLEMS

Starting in the late 1960s, work on radiation transport IPs has progressed steadily. The simplest problems, in which scattering is ignored, were examined first. Example problems using this approach are computed tomography,⁴⁵ the reconstruction of absorption coefficients based on active measurements taken on the exterior, and the estimation of a one-dimensional spatial change in atmospheric parameters.¹⁸ These methods assume that the medium is optically thin.

Many other IPs have since come along, in a wide range of applications. In oil-well logging, the material composition deep underground is inferred based on radiation measurements.² By measuring the energy spectrum of radiation scattered back to a detector from a radioactive source inserted in a borehole, engineers can estimate the hydrogen density or bulk density of the surrounding rock formation. In marine biology, bioluminescence, the emission of light by living organisms, can be computed as a function of depth in the ocean using an inverse method.¹²² A common IP is finding the elemental composition of a material.⁴⁴ Closely related is material identification, where a common goal is to detect illicit nuclear materials.⁷⁷

Given the wide range of IP applications and methods, there are a number of ways to classify and sort the methods. Algorithmically, the most important distinction is analytic or closed-form solutions versus discrete or implicit methods. In addition, the particular application necessitates a choice from the following options:

- Time-dependent versus time-independent. Are the system and measurements steady-state or not?
- Solving for material properties, internal sources, incoming fluxes, or slab thickness. What is known and unknown?
- Active versus passive. Is the measured radiation coming from the system, or is an outside source used for interrogation?
- Interior versus exterior sensing. Are the detector measurements taken within the medium, or are they required only at the outside surface?

These alternatives clearly have a large effect on the problem, sometimes making the difference between a simple solution and an intractable one. Each of these must be

considered when facing an inverse problem.

From a methods perspective, the analytic versus discrete distinction is the most important. Analytic inverse methods seek a closed-form solution using manipulations of the transport equation. In contrast, discrete methods begin with the discretized equations of the forward problem and borrow techniques from numerical linear algebra, applied math, and optimization to solve the inverse problem. (McCormick chooses a similar, but not identical, taxonomy: explicit and implicit methods.⁷⁸) In the early radiation transport IP literature, analytic methods are most common until the middle of the 1990s, when discrete solutions began to attract attention. This work employs the discrete approach, and thus most of the proceeding historical review focuses on the discrete approach. However, a brief summary of key analytic solutions is first presented.

3.2.1 *Analytic solutions*

Analytic IP methods are concerned with exact solutions for the unknown. These methods are reviewed by McCormick.⁷⁸ Most of these solutions are limited to one dimension, usually slab or plane-parallel geometry, since general solutions for the radiation transport equation are only available in one dimension. The problems are mostly one group, although a few deal with the complications raised by a dependence on energy. Most early methods require both interior and exterior detectors, while later work eliminated the need for interior detectors in some cases. Many analytic solutions were for material properties, densities, cross sections, and scattering moments. However, a few papers present solutions for an internal radiation source.

In one of the early papers, Larsen computes the internal angular fluxes and sources for a half-space plane-parallel one-speed problem if the incident and outgoing angular fluxes are known.⁷⁰ Note that no interior measurements are required. Zweifel builds on this solution with what he calls the “canonical inverse problem” – determining the incoming angular flux given an outgoing flux. Under certain conditions, his solution reduces to Larsen’s solution for the internal source. However, in commenting on this solution, McCormick doubts its usefulness:⁷⁹

The successful implementation of this source estimation technique is even more questionable than the estimation of a canonical inverse problem incident distribution.

This is a reference to the equation's ill-conditioning, which is a general property of problems with only exterior detectors.

This ill-conditioning can be lessened by taking measurements within the medium. This approach was taken in a unique application in bioluminescence in the ocean.¹²² Researchers were interested in the variation of bioluminescent sources as a function of depth below the ocean surface. The computation was complicated by moonlight incident on the surface. They developed analytic equations based on a balance of photons in each layer. Then they were able to compute the amount of bioluminescence in a given layer using measurements of photon irradiance and scalar irradiance (current and scalar flux in NE parlance) on the surface and at selected depths.

Since this thesis employs the adjoint flux, we point out a few papers that make use of the adjoint in an analytic IP solution.^{68,69,98} These papers presented solutions for material properties, not internal sources. The adjoint fluxes aided in the manipulation of the transport equation to achieve the authors' desired result.

Sanchez and McCormick answer the question of uniqueness for the linear inverse source problem, where the system geometry and cross sections are known but the source term is unknown.^{99,100} They consider invasive problems, in which detectors are placed in the interior and exterior of the system, as well as noninvasive problems, where the detectors are only on the outside of the system. When the angular flux is known everywhere in the interior of the domain, the invasive problem has a unique solution, given by applying the transport operator to the known flux distribution. However, they offer a family of possible solutions for the noninvasive solution, proving that the noninvasive solution is not unique. The nonuniqueness is a property of the linear transport equation; the effect of measurement error is ignored in their arguments.

In general, analytic solutions are not possible for three-dimensional, general geometry, multigroup problems. This lack of applicability, combined with the growing power and capacity of computing resources, prompted research into discrete methods for IPs.

3.2.2 *Discrete methods*

Instead of seeking closed-form solutions, discrete methods generally assume that the differential equations are solved numerically. Within the discrete approach, researchers

in the past decade have tackled a variety of problems. Inverse problems have been considered with many different combinations of the set of alternatives presented above. The particular combination places certain limits on the solution methods available. A few of the relevant inverse problems are listed here.

Norton is one of the earliest to suggest the use of the adjoint flux for IPs.^{91,92} His method assumes that the problem can be written in terms of an error functional that returns the difference between measured and computed responses. He proposes a general method using the adjoint to compute Fréchet derivatives of this error functional. For discrete problems, the derivatives are gradients, which can be used in any gradient-based optimization procedure: steepest descent, conjugate gradient, or quasi-Newton. He also derives these expressions using the Lagrange multiplier. Note that the use of the adjoint in computing the Fréchet derivative is the same as that used in the sensitivity analysis community.¹²

Favorite, Bledsoe, and Ketcheson, in various co-author combinations, have considered geometry and materials (partially) unknown in a series of papers. Their methods have primarily used gradient-based minimization procedures. Building on the work of Norton, Favorite laid the foundation of their methods by giving the expressions to compute the gradient of the error functional using adjoints and the Schwinger variational functional.³² They assume that modeling the uncollided flux is sufficient to describe the detector responses, since the spectral lines can be resolved by a high purity germanium detector. Thus, instead of multigroup calculations, they do uncoupled calculations on multiple discrete energy “lines.” The first test problem they use is a sphere with a central source and surrounding layers of shielding.^{29-31,33} With few unknowns and a reasonable initial guess, they achieved success in resolving the inverse problem under the above assumptions and limitations.

In a more recent series of papers, Favorite, Bledsoe, and Ketcheson estimate some of the dimension and composition parameters of a shielded uranium cylinder, assuming the other parameters are known.^{8,9} Instead of relying only on gradient-based methods, Bledsoe and Favorite also use global stochastic optimization methods. Since they use only the uncollided flux, they also explore the use of ray tracing methods.²⁸

Mattingly and Mitchell tackled a one-dimensional spherical problem in which the material layer thicknesses were unknown but the materials were known, or could be inferred from the peaks present in the high-resolution detector reading.⁷⁷ The unknowns in the solution were the layer thicknesses, although the more fundamental

unknowns seemed to be the total mass of each material present. To more tightly constrain the solution, Mattingly and his coauthors include more physics than others, using neutron multiplicity counting in addition to the more conventional gamma spectroscopy.⁷⁶ For the nonlinear minimization, they use the Levenberg-Marquardt⁶¹ procedure, where they compute the needed derivatives using a finite-difference approximation.

Researchers at Texas A&M have published a number of papers on neutron tomography – where an external active neutron source is applied on one side of the object and detectors are placed on the opposite side. In contrast with computed tomography in medical applications, these methods are focused on objects that are optically thick, making the inclusion of scattering effects essential. Scipolo illuminated a two-dimensional area of unknown cross sections and used a gradient-based method to minimize an objective functional.¹⁰² Sternat and Ragusa use the adjoint flux solution (what they call the “duality principle”) to compute the gradients necessary for a Newton-type optimization.¹¹⁰ They present results for a one-group, two-dimensional problem. Wu and Adams suggest using a change-of-variables to constrain the unknowns.¹²¹ For example, they use a logarithmic transformation to ensure that the cross section is always non-negative. This allows the use of unconstrained optimization, whose theory and implementation is much simpler than that of constrained optimization.

Working on border security applications, Miller and Charlton seek to triangulate and estimate the activity of a point source in two dimensions near a portal monitor at a border crossing.^{86,87} They assume that the position of automobiles and trucks could be obtained with image recognition software and that the material cross sections of the automobiles and surroundings could be reasonably estimated based on a database. Thus, the problem is similar to the one considered here, with the geometry and cross sections known. However, there is a significant difference between their problem and ours: they focus on identifying one point source, while we do not begin with that assumption. In their test problem with three detectors, the iterative procedure was sensitive to the initial guess, becoming stuck in local minima in many cases.

Jarman, et al. build on the work of Miller and Charlton by applying a probabilistic inverse method to the portal monitoring problem.⁵⁴ They solve for the posterior probability distribution for the location and source strength of a point source. They compute the posterior distribution directly at discrete values of each model parameter. This ap-

proach is tractable for their problem because they have 4 unknown model parameters: three spatial coordinates and one source strength. Their results show the usefulness of the Bayesian approach for estimating the uncertainty in the computed mean location. In addition, they demonstrate that accounting for scattered radiation significantly improves the estimated source location and intensity for their test problem, in which the source is located in a medium-density absorber.

Finally, we note that Bayesian inverse methods have been used for decades within the nuclear engineering community to estimate or adjust cross section data. Jesse in §1.6 of his dissertation reviews the use of these methods.⁵⁸ He calls the method the Generalized Linear Least Squares algorithm. Although the quantities to be estimated are different than what we study here, the underlying methods are quite similar.

3.3 MORE INFORMATIVE DETECTORS

The discrete methods discussed thus far assume that the detector is unable to distinguish the incidence direction of incoming particles. While clever algorithms can infer a surprising amount of information using only isotropic detectors, direction-aware detectors can provide much richer and sensitive data for use in inversion procedures. Many of these detector systems share features with cameras.

3.3.1 *Compton camera*

A Compton camera is made of two parallel planes of semiconductor detectors, often CdZnTe, abbreviated CZT.^{23,123} Each plane is made of an array of crystals. In Ref. [118], each crystal is approximately a cube with 20 mm sides, and each plane is composed of 3×3 crystals. The crystals are three-dimensional position sensitive. That is, through electronic measurements, the location of a photon interaction within the crystal can be obtained. If a photon happens to interact via Compton scatter in both the front plane and then scatter or be absorbed within the back plane, the direction of the incident particle can be inferred to within the surface of one or two cones. This computation is based on the location of interaction, the energy deposited in the interaction, and the conservation equations for energy and momentum.

3.3.2 *Coded-aperture imaging*

A coded aperture imager is similar to a pinhole camera, in which light enters the front plane through a small opening and is recorded on the back plane. With this arrangement, given a location on the back plane, there is only one direction in which a photon could travel to arrive at this location. The same principle can be applied to radiation detection, using a shielding material in the front plane to block all photons except the few entering through the narrow hole. However, radiation detectors function with many fewer particles available to detect than for a normal camera. For a pinhole detector, the number of particles which reach the back plane would usually be unacceptably low. To avoid these losses, a middle ground is sought in which, instead of only a small opening in the front plane, a significant portion of the front plane is left open.^{21,124} This preserves some of the ability to determine photon trajectories while improving the detection efficiency.

In the preceding chapters, we have presented background on inverse problems in general, as well as specific details for radiation transport applications. This chapter describes the methods implemented in this work to solve the radiation source mapping challenge. Since we have focused on the Bayesian approach, the various components needed for this method are presented in the following sections. However, we first show how this problem can be written as a first-kind Fredholm integral equation.

4.1 FREDHOLM INTEGRAL EQUATION

Many inverse problems can be phrased as Fredholm integral equations of the first kind (Equation 2.4). This is true for this problem as well. Recall that the measured data is the detector response. We could in theory measure this everywhere in the spatial domain. The response is also a function of energy. Thus, we propose $r(\vec{x}', E')$ for the continuous detector response [106, §5.4]. If the detector response function is $\sigma_d(E', E)$, then the response can be computed as³⁵

$$r(\vec{x}', E') = \int dE \sigma_d(E', E) \phi(\vec{x}', E) \quad , \quad (4.1)$$

where E is the energy of the incident radiation and E' is the energy deposited in the detector. Here we assume that a detector could be placed at every point in the domain. In this equation, r and σ_d are known and ϕ is unknown. However, we are here interested in the source, not ultimately the flux field that it generates. To make this clear, we can substitute into Equation 4.1 from $L_s \phi = q$ (the continuous operator analog to Equation 3.23). Thus, one way to write the Fredholm integral equation with

the source q is

$$r(\vec{x}', E') = \int dE \sigma_d(E', E) L_s^{-1}(\vec{x}', E) q(\vec{x}', E) \quad ,$$

where the detector response function is the kernel.

Alternately, these equations can be reformulated with the assistance of the adjoint flux (see section 3.1.1). However, now we need a family of adjoint fluxes, $\phi^+(\vec{x}, E \rightarrow \vec{x}', E')$, that are solutions to

$$L_s^+(\vec{x}, E) \phi^+(\vec{x}, E \rightarrow \vec{x}', E') = \sigma_d(\vec{x}', E'; \vec{x}, E) \quad ,$$

where the detector response function now depends on space. The DRF will be discussed in more detail in section 4.4. The arrow is used in the argument of ϕ^+ to emphasize the similarity with Green's function methods [24, §5.4(a)].

Thus, the Fredholm equation in terms of the adjoint ϕ^+ is

$$r(\vec{x}', E') = \int dE \int d\vec{x} \phi^+(\vec{x}, E \rightarrow \vec{x}', E') q(\vec{x}, E) \quad . \quad (4.2)$$

The adjoint flux is now the kernel. An unfortunate consequence of using the adjoint flux is the additional integral over space.

In practice, we only have measurements of $r(\vec{x}', E')$ at a few points in (\vec{x}', E') space, limiting the resolution of the reconstruction of q . In addition, ϕ^+ is estimated with numerical methods. Thus, in practice, the continuous functions r , ϕ^+ , and q must be represented by discrete quantities, r and q as vectors and ϕ^+ as a matrix. The details will be provided below, but the general form of the resulting matrix equation is

$$\vec{r} = \Phi^+ \vec{q} \quad . \quad (4.3)$$

The matrix Φ^+ is typically not square. We can apply the Bayesian inference method to this equation.

4.2 MODEL PARAMETERS

In regression, least squares, and Bayesian inference, we have two categories for items to quantify: model parameters and system responses. This section discusses model parameters and the next handles system responses. Model parameters are the “input”

pictured in Figure 2.1. Following Cacuci's notation, the model parameters are packaged into the vector $\vec{\alpha}$. While $\vec{\alpha}$ could include all data going into our model, often we only include the most interesting or sensitive parameters.

In our case, the most important model parameters describe the radiation source distribution in space and energy, while we assume weak dependence on angle. In addition, other model parameters may include system dimensions, material composition and geometry, position of detectors, and many others. The values of these state variables will be known with varying degrees of confidence. For example, we have little initial information about the source distribution (perhaps we can bracket the range of possible source intensities within an interval spanning ten orders of magnitude). On the other hand, the dimensions of walls and floors are known with low uncertainty, assuming that the geometrical configuration is given, say from floor plans or construction blueprints. While all of these parameters could be included, usually only those inputs which have significant uncertainty need be included. Given an infinite budget, all the parameters could be included in $\vec{\alpha}$. However, using engineering judgement, the exclusion of parameters with low uncertainties should not affect the final results.

4.2.1 *Basis of the source*

One could imagine numerous ways in which to represent the source distribution in energy and space. The two primary means to do this were discussed in section 2.1.5: quadrature and expansion methods. Since the adjoint flux will be computed on a Cartesian mesh in space and a multigroup energy structure, a basis to represent the source should be compatible with this discretization. The spatial and energy discretizations are discussed below.

Energy basis Radiation sources can produce radiation with discrete energies or with a continuous spectrum. We wish to have the ability to differentiate between them and track them individually. First, nuclear or electronic transitions often result in the emission of gamma or x-rays with energies in a small set of defined values. The gamma energies associated with the decay of radionuclides are known and tabulated.¹¹⁵ From a measured detector spectrum, the gamma line energies should be easily discernible. We set n_p (p for peak) as the number of gamma energies emitted. The source unknown

q_p^{peak} for a particular discrete line p is simply the number of particles emitted per unit time at that energy.

The second type of radiation is emitted with a continuous spectrum. Bremsstrahlung photons, commonly produced by an x-ray tube used in medical applications, have a range of energies from nearly zero to a maximum energy corresponding to the initial energy of the incident beam of charged particles [106, §3.6.2]. Neutrons born from fission are another example of a continuous-spectrum radiation [67, §3-4]. To preserve compatibility with the multigroup calculation, we bin the continuous source into the n_g energy groups. If group g has energy boundaries E_g to E_{g-1} , the source in this group from a continuous source $q(E)$ is

$$q_g^{\text{cont}} = \int_{E_g}^{E_{g-1}} dE q(E) \quad ,$$

where q_g^{cont} is the number of particles emitted from a continuous-spectrum source within group g per unit time.

Sometimes it is useful to combine the discrete and continuous radiation sources into one total source. We sum this according to the same multigroup structure. To perform this sum, we must know to which group each peak belongs. If the energy of peak p is E_p , then p is in group g if $E_p \in [E_g, E_{g-1})$. We call this mapping $g(p)$, which returns the corresponding energy group g given a peak index p . Many or no peaks may fall within a particular group. Using the peak-to-group mapping, the total source in group g from both discrete and continuous spectrum is

$$q_g^{\text{total}} = \int_{E_g}^{E_{g-1}} dE q(E) + \sum_{p \in S_g} q_p^{\text{peak}} \quad , \quad (4.4)$$

where $S_g = \{p \mid g(p) = g\}$. For convenience, let $n_{\tilde{g}} = n_g + n_p$.

Spatial basis Following the spatial discretization used by the adjoint, we split the spatial domain into n_c Cartesian cells. Then the source unknown is the cell-averaged source. In this work we only track the zeroth spatial moment of source, although higher moments could theoretically be included if so desired. Although this basis is an obvious choice, it is not difficult to use others. We could form a basis on a coarse spatial mesh, thus constraining the source to be constant over groups of finer mesh

cells. Or one could coarsen some cells but not others, leading to an adaptive basis (see Appendix E). If the source distribution is believed to be smooth over the domain, one could use a polynomial or Fourier basis, assuming that it yields non-negative source intensities over the entire domain. In these ways, the selection of the basis can incorporate any prior knowledge about features of the source distribution.

Taking the Cartesian product of the energy and space discretizations, we can form $n_{\bar{g}} \cdot n_c = n_{\bar{g}c}$ basis vectors. For many alternate bases, there is a linear map $P : \mathbb{R}^s \rightarrow \mathbb{R}^{n_{\bar{g}c}}$ such that

$$\vec{q} = P\vec{x} \quad , \quad (4.5)$$

where s is the dimension of the smaller basis and \vec{x} is an s -vector containing the expansion coefficients of this basis. This projection transforms Equation 4.3 from

$$\vec{r} = \Phi^+\vec{q} \quad , \quad (4.6)$$

to

$$\vec{r} = \Phi^+P\vec{x} = A\vec{x} \quad , \quad (4.7)$$

where $A = \Phi^+P$. By tuning s such that it is smaller or equal to the size of \vec{r} , the system of equations is overdetermined or well-determined. This strategy is suggested by Santamarina and Fratta on the basis of Ockham's Razor, that is, to favor a simple model over a complex one [101, §9.5.3]. This principle, also known as the rule of parsimony, is: "Plurality must not be posited without necessity."

One simple modification of the natural basis is to exclude certain energy-space cells which do not have any radiation source. For a variety of reasons, we may know that certain spatial or energy regions contain no source. For instance, it is reasonable to only treat energies at and below the maximum radiation energy detected. The algorithm could accommodate higher groups for which no source would be computed, but this would be a waste of time and effort. In a similar manner, certain spatial regions may be ruled out and excluded from \vec{a} . For example, at a border crossings, radioactive sources are sought as they are carried by cars or trucks. Algorithms to map the sources in these vehicles could include as unknowns the source everywhere, both inside and outside the vehicles. However, it would be more efficient to only include those cells which are located inside the monitored vehicles. In the test problem in this work, locating sources in a room, we have assumed that the sources are not suspended

in air and so have eliminated the cells which only contain air. The linear projection operator P in this case is particularly simple; it is an identity matrix missing select columns. The number of reduced spatial cells is represented by n'_c .

Packaging source unknowns in a vector $Q \in \mathbb{R}^{n_g \times n_c}$ is an array containing the expansion coefficients of the source intensities per group or peak and per spatial cell. This array is flattened into a vector,

$$\vec{q} = \text{flat}(Q) \quad .$$

The flat function unravels the array with multiple dimensions into a one-dimensional array. The inverse of the flat function is reshape, which requires a second argument giving the shape of the matrix. The particulars of the flat and reshape functions are not of importance here, as long as they are performed consistently. Returning to Equation 4.4, we can rewrite this sum of the peak and continuous sources $\vec{q}^{\text{total}} \in \mathbb{R}^{n_{gc}}$ as

$$\vec{q}^{\text{total}} = G \vec{q}^{\text{peak}} + \vec{q}^{\text{cont}} \quad .$$

The matrix $G \in \mathbb{R}^{n_{gc} \times n_{pc}}$ sums the discrete peaks into their energy group. The above equations are valid for the reduced spatial basis if n'_c is substituted for n_c .

Here we pause to note that the dimension of the natural basis, n_{gc} , or even the dimension of the reduced basis n'_{gc} , will typically be much larger than the number of measurements available. As mentioned above, the system of equations represented by Equation 4.3 or Equation 4.7 will be severely underdetermined since the number of unknowns is much larger than the number of measurements (and equations). For least-squares problems, this poses a serious dilemma, since an underdetermined system admits an infinite number of solutions. However, the Bayesian inference approach sidesteps this issue by accounting for uncertainties. That is, the method ranks the various admissible solutions in terms of their probability of occurrence given the prior knowledge and the measured data. Thus, the method allows for the solution of underdetermined, well-determined, or overdetermined problems. However, there is no free lunch. Although one can solve a severely underdetermined system using Bayesian inference, it is likely that the low information content of the measurements will be reflected in little or modest reductions in uncertainties in the final solution.

4.2.2 Logarithmic transformation of source intensities

The most direct means of including the sources in $\vec{\alpha}$ is by setting $\vec{\alpha} = \vec{q}$. While using \vec{q} as the model parameters is logical, we instead use $\log_e(\vec{q}/q_0)$, where \log_e is the natural logarithm operating element-wise. The positive scalar value q_0 is some reference source value necessary to make the argument of the logarithm dimensionless. This logarithmic transform has two advantages. First, as noted by Tarantola [112, §1.2.4, §6.2], for positive scaling factors such as the source intensity, the logarithmic transform disallows unphysical negative source intensities. This is especially important in our context, because we describe the PDFs of the parameters with multivariate Gaussians. A Gaussian is non-zero everywhere, and so it assigns non-zero probability to negative values of the parameter – an unphysical consequence. However, a log-normal distribution assigns non-zero probability only to positive values of the parameter.

Secondly, this transform enables a much broader prior distribution, which is beneficial here because we initially have little idea of the source strength order of magnitude. Assigning a prior distribution which spans many orders of magnitude is much simpler for a log-normal distribution than for a normal distribution. This feature is exploited when we assign the prior distribution of the source intensities.

The log-normal distribution was first suggested by Jeffreys for positive parameters.⁵⁷ These positive parameters and the prior PDFs associated with them are known as Jeffreys parameters and Jeffreys priors. Others in the nuclear engineering community¹²¹ and in physics⁵³ have used similar techniques for pragmatic reasons.

Before leaving this topic, we should mention the disadvantages of the logarithmic transform. Conceptually, the biggest cost is the conversion of a linear set of equations to a nonlinear set. Usually one seeks to linearize a nonlinear problem, not to unlinearize a linear problem. Solving a nonlinear problem can be significantly more costly than a linear one. While the Jacobian of the linear system is constant, the Jacobian of the transformed system changes as $\vec{\alpha}$ changes. Also, the logarithmic transform can cause instabilities in the optimization method. The optimization is done in log-space, but a medium-sized step in log-space can produce floating point overflow or other breakdowns in linear-space. As a result, a constrained optimization method is used even though we introduced the logarithmic transform in part to remove the constraints. These issues will be discussed in greater detail in section 4.6.

4.2.3 *Other inputs as model parameters*

As stated above, all the system inputs could be included in the model parameter vector $\vec{\alpha}$. Beside the source map, this could include material compositions and densities, material cross sections, geometric dimensions, and characterizations of the detector system. Deciding among all of these possibilities could be a bewildering task. In a situation where the parameters of interest have a relative uncertainty close to the relative uncertainty of other inputs, the latter's effect may have significant influence on the inverse solution accuracy and trustworthiness. However, in the present case, the inputs other than the source distribution are assumed to have a relatively low uncertainty, especially when compared to the initial uncertainty in the source, and even in comparison to the desired uncertainty in the solution, which is presumably an order-of-magnitude estimate.

Furthermore, we have not included the material cross sections in $\vec{\alpha}$ for several reasons. First, the covariance matrices describing the errors in microscopic cross sections for photon interactions are unavailable in the current version of ENDF/B-VII.1.¹⁶ See Appendix B for select estimates of the errors in this data found in the source libraries. The other reason for excluding the microscopic cross section uncertainties is that their level of relative uncertainty (estimated to be a few percent, see Table B.1) is likely to be much lower than the uncertainty for material composition and densities. These quantities were estimated according to common building material values, but these estimates are likely in error by more than a few percent.

4.2.4 *Model parameter covariances*

While we could include some positive correlations in neighborhoods of spatial cells to promote smoother source solutions, we have not pursued this option in this work. The off-diagonal elements of C_α are set to zero.

4.3 RESPONSES

Having discussed model parameters, the first category of values, we now turn to the second category, that of measured responses. The raw detector data is a series of counts placed in uncalibrated bins by the multichannel analyzer (MCA). We assume that there are n_b MCA channels, where n_b typically is a few thousand, depending on the

hardware. Through the use of calibration sources or other noticeable peaks, an energy calibration is computed to map each channel edge to the photon's corresponding energy value. This raw spectrum consists of two components. First, we have full-energy photopeaks. These peaks represent the full absorption of an uncollided photon into the detector. These peaks are the easiest to analyze. Second, we have the detector continuum. This smoother portion of the measured spectrum is more difficult to analyze because the cause of these contributions may be (a) a scattered photon depositing all or part of its energy in the detector or (b) an uncollided photon depositing part of its energy in the detector. We desire the ability to use both the photopeaks and continuum of the detector response in the reconstruction of the source distribution.

There are n_p peaks. These peaks can be identified and analyzed automatically or by hand. In either case, the important data about the peak is the centroid energy and the peak area and accompanying Poisson variance. Photopeaks are convenient to analyze because the peak area corresponds directly to the number of particles (at the peak energy) incident on the detector and the detector's peak intrinsic efficiency. Computing the peak intrinsic efficiency is a relatively simple task involving the detector's shape, the source-to-detector distance and orientation, and the detector's energy-dependent linear attenuation coefficient. A Monte Carlo ray tracing simulation is sufficient to perform this calculation. It can be executed as a preprocess and the data saved for repeated use by the IP solution algorithm.

The continuum is more difficult to analyze. First, we must rebin the MCA finely-binned data into n_g coarser groups. This is because we do not want to perform a transport calculation with thousands of energy groups. Second, to analyze the continuum, we must have access to the full detector response function, $\sigma_d(x', E', E)$. Unfortunately, for a non-spherical detector, the detector response is a function of space, that is, the source-to-detector configuration (see section 4.4).

4.3.1 *Detector locations and directions*

To make a map of the source in energy and space, we need several different detector readings. The most basic method to get multiple readings is to move the detector to a different location, or install an array of detectors throughout an area. We will assume there are n_l detector locations. In addition, having a sense of direction with

the detectors can be useful. With a large budget, one could use a coded-aperture detector system or a Compton camera (see section 3.3). A cheaper and lower-resolution option is to use collimators to shield the detector from particles approaching the detector with certain directions. In the demonstration, we used uncollimated detectors with a 4π view as well as hemispherical 2π collimators to look left or right, up or down, and in or out. We assume that we have $n_a = 7$ directions at each location, including the uncollimated 4π reading. Thus, we have $n_l \cdot n_a$ detector spectra. For each spectrum, there are n_p photopeaks and n_g coarse energy groups. Altogether, this gives $n_l \cdot n_a \cdot (n_p + n_g)$ data elements, which is the size of \vec{r} .

The response vector \vec{r} has two sub-vectors, the peak and the continuum. The peak response vector is $\vec{r}^{\text{peak}} = \text{flat}(\mathbf{R}^{\text{peak}})$, where \mathbf{R}^{peak} has size $n_l \times n_a \times n_p$. Likewise, the continuum response vector is $\vec{r}^{\text{cont}} = \text{flat}(\mathbf{R}^{\text{cont}})$, where \mathbf{R}^{cont} has size $n_l \times n_a \times n_g$. Then the full response vector is

$$\vec{r} = \begin{bmatrix} \vec{r}^{\text{peak}} \\ \vec{r}^{\text{cont}} \end{bmatrix} .$$

4.3.2 Response errors and covariances

Each channel has an associated Poisson error. These errors need to be propagated through the peak summing and the collapsing of MCA bins into coarse multigroup structure. The error propagation is simple since these are essentially summations. Although there are efforts to measure the correlations between channels in a detector spectrum,¹¹⁴ this technology is not widely available. Thus, we treat the detector readings in different groups as uncorrelated, making the \mathbf{C}_m matrix diagonal. In addition, the $\vec{\alpha}$ - \vec{r} correlations are assumed to be zero.

4.4 DETECTOR RESPONSE FUNCTION

A detector response function translates the radiation field at a detector to the response of that detector. The detector response can be either an energy-integrated count or a spectrum produced by an MCA. The DRF is the kernel in Equation 4.1. However, there are many variations of the DRF, some more complex and accurate and others simpler but approximate. The detector response depends on the size and shape of the detector,

Table 4.1: Parameters for source inversion.

Index	Total	Description
c	n_c	Spatial cells in adjoint calculation
c'	n'_c	Spatial cells in reduced source basis
g	n_g	Energy groups in adjoint calculation
b	n_b	MCA energy channels
p	n_p	Photopeaks
\tilde{g}	$n_{\tilde{g}}$	Photopeaks and energy groups
l	n_l	Detector locations
a	n_a	Detector directions

so a more general representation of the DRF for a detector occupying volume V_d is

$$r(E') = \int_{V_d} dV \int_0^\infty dE \sigma_d(\vec{x}, E', E) \phi(\vec{x}, E) \quad . \quad (4.8)$$

Here $\sigma_d(\vec{x}, E', E) dE'$ is the probability per path length that a particle at position \vec{x} with energy E will deposit an amount of energy in the interval $[E', E' + dE']$ into the detector. At this stage, the specific interactions within the detector are important to track. In current practice, these integrals are typically computed with Monte Carlo transport methods, often with empirical adjustments to match certain phenomena.^{35,88,93} The DRF kernel can be computed by estimating this integral for a sequence of values of E . In discrete form, σ_d is a matrix. While Monte Carlo methods are quite powerful for estimating Equation 4.8, they are usually quite slow. Thus, a common approach is to precompute the DRF once. Then the response for a particular flux spectrum can be computed with a matrix-vector product. All of the physics and complications of the detector are hidden within the DRF matrix. By doing this, the spatial dependence of the DRF is removed:

$$r(E') = \int_0^\infty dE \sigma_d(E', E) \int_{V_d} dV \phi(\vec{x}, E) \quad . \quad (4.9)$$

For a small, optically-thin detector, the flux would be attenuated very little inside the

detector, and so

$$\int_{V_d} dV \phi(\vec{x}, E) \approx V_d \phi(\vec{x}_0, E) \quad , \quad (\text{optically-thin detector}) \quad (4.10)$$

where \vec{x}_0 is a point within the detector. However, at 600 keV, 2 inches in a sodium-iodide (NaI) crystal is 1.53 mean free paths.* This means that using Equation 4.10 for a 2''-by-2'' detector would introduce significant errors. Instead, one could use the volume-averaged flux

$$\langle \phi(E) \rangle_{V_d} = \frac{1}{V_d} \int_{V_d} dV \phi(\vec{x}, E) \quad , \quad (4.11)$$

to evaluate the integral

$$\int_{V_d} dV \phi(\vec{x}, E) = V_d \langle \phi(E) \rangle_{V_d} \quad . \quad (4.12)$$

In computing the integral in Equation 4.11, we are concerned only whether particles entering the detector volume have an interaction. It is unnecessary to track particles after scatter events, since the DRF accounts for these scatter events. The integral depends on several influences:

1. the energy-dependent linear attenuation coefficient of the detector,
2. the shape and position of the detector, and
3. the shape and position of the radiation source or sources, and
4. the direction of approach of radiation as it enters the detector. The radiation may all come directly from the source, or it could have scattered and have an entirely different approach.

Items 2 through 4 are necessary for computing the distribution of chord lengths that radiation travels through the detector volume. The distribution of chord lengths is dependent on the source's view of the detector. If the detector was spherical, the only influence would be the source-to-detector distance. For non-spherical sources, the source-to-detector orientation is also important. For example, in this work we use cylindrical detector crystals. While controlled experiments often place the source at a

*From NIST's database, for NaI, $\mu_t/\rho = 8.225 \times 10^{-2} \text{ cm}^2/\text{g}$ at $E = 600 \text{ keV}$.⁵ The density of NaI is $\rho = 3.67 \text{ g/cm}^3$ [63, Table 8.3], making $\mu_t(600 \text{ keV}) = 0.3019 \text{ /cm}$. For $\Delta x = 2 \text{ in} = 5.08 \text{ cm}$, $\mu_t \Delta x = 1.53$.

fixed point along the axis of the detector cylinder, in this work the source may have any orientation with the detector.

With the stated assumptions, computing the integral in Equation 4.11 is tightly connected to computing the intrinsic total efficiency. The intrinsic total efficiency $\epsilon_{t,i}(E)$ is the probability of interaction within the detector for a particle with energy E that enters the detector (see Appendix D.1 or [63, §10.D]). Just as for $\langle\phi(E)\rangle_{V_d}$, the intrinsic total efficiency depends on the four factors given above. To make this explicit, we will write the efficiency as depending on the source location \vec{x} and detector location \vec{x}' , $\epsilon_{t,i}(E; \vec{x}, \vec{x}')$. The probability of interaction within the detector with the total attenuation coefficient $\mu_t(E)$ for given a particular track of length s is

$$p(s, E) = 1 - e^{-\mu_t(E)s} \quad .$$

However, there is in general no closed-form equation to compute the mean of this probability of all particle tracks through the detector. This necessitates a Monte Carlo ray tracing approach to estimate the value of the integrals.

Accounting for these dependencies, the DRF for a detector located at \vec{x}' and source at \vec{x} is

$$\sigma_d(\vec{x}', E'; \vec{x}, E) = \tilde{\sigma}_d(E', E) \epsilon_{t,i}(E; \vec{x}, \vec{x}') \quad . \quad (4.13)$$

Here $\tilde{\sigma}_d(E', E)$ is a normalized probability distribution function. For a particle with energy E that interacts in the detector, $\tilde{\sigma}_d$ describes the probability of depositing energy E' .

4.4.1 Simplified response functions

That the DRF is dependent on space complicates the solution algorithm, so we make a few approximations to simplify the problem. The first approximation is to neglect the spatial dependence of the detector efficiency,

$$\epsilon_{t,i}(E; \vec{x}, \vec{x}') \approx \epsilon_{t,i}(E) \quad .$$

For detectors that are nearly spherical, this is a decent approximation. For detectors with a high aspect ratio, this approximation can cause significant error. Using a 2'' \times 2'' detector, the error induced at lower energies (10s to 100s keV) can approach 50%.

However, at energies around 1 MeV, the error is less than 5%. Since the higher energy photons will likely carry more information than low-energy photons, and hence be more valuable for solving the IP, this approximation is not altogether unreasonable.

The second simplification is motivated by numerical simulations rather than physical reasons. If the DRF is assumed to be the identity operator, then the detector response is the flux spectrum. The identity DRF is

$$\sigma_d(E', E) = \delta(E - E') \quad ,$$

such that

$$r(E') = \int_0^\infty dE \delta(E - E') \phi(E) = \phi(E') \quad .$$

The advantage of this DRF is that it allows for the testing of the source mapping algorithm without concern for the particular physics of the detector.

4.4.2 Discrete response functions

The above discussion has focused on continuous DRFs. However, in practice they must be discretized. Analogous to energy binning performed by an MCA, the DRF can be computed in discrete energy bins. The DRF is a matrix, Σ_d , with the row index corresponding to the detector energy E' and the column index with the particle energy E . Each column is a discrete probability distribution function representing the likelihood of scoring in bin E' from a photon with incident energy in a particular E bin. If $\vec{\phi}$ is the flux spectrum at the detector, then the measured spectrum is

$$\vec{r} = \Sigma_d \vec{\phi} \quad .$$

As a practical matter, the energy bins for the MCA and the DRF will likely not be equal. In this case, either the detector spectra or the DRF must be rebinned according to the binning of the other. Knoll outlines the basic approach in rebinning a spectrum from one set of bin edges to another [63, §18.B]. We use a standard energy mesh and rebin both the MCA spectra and the DRF (using a 2-D rebinning) to this energy mesh.

4.5 SENSITIVITIES

Having provided the details concerning the source and the responses, we can now proceed to describe the form of Equation 4.3.

4.5.1 Uncollided photopeak sensitivities

The discrete form of the uncollided adjoint equation for the detector location l , detector direction a , and photopeak p is

$$\mathbb{L}_{s,unc}^+ \vec{\phi}_{lap}^+ = \vec{\sigma}_d^{lap} \quad . \quad (4.14)$$

The vectors $\vec{\phi}_{lap}^+$ and $\vec{\sigma}_d^{lap}$ have n_c elements each, and $\mathbb{L}_{s,unc}^+$ is a $n_c \times n_c$ matrix. The adjoint source vector $\vec{\sigma}_d^{lap}$ is zero everywhere except for the element corresponding to the spatial cell where the detector is located. This element is equal to the intrinsic peak efficiency at the energy of peak p , that is, the probability an uncollided particle at this energy that enters the detector will deposit all its energy within the detector. The adjoint transport operator matrix $\mathbb{L}_{s,unc}^+$ considered in this section only accounts for uncollided flux; it ignores the scattering of particles.

While the adjoint flux vector only contains one group, we place it in a vector of zeros to make it length n_{pc} . This larger vector is

$$\vec{\phi}_{lap,peak}^+ = \begin{bmatrix} \vec{0} \\ \vdots \\ \vec{0} \\ \vec{\phi}_{lap}^+ \\ \vec{0} \\ \vdots \\ \vec{0} \end{bmatrix} \quad ,$$

where the p -th zero vector is replaced by $\vec{\phi}_{lap}^+$. Each zero vector has length n_c , so $\vec{\phi}_{lap,peak}^+ \in \mathbb{R}^{n_{pc}}$.

From Equation 3.20, the uncollided peak p response due to source $\vec{q}^{\text{peak}} \in \mathbb{R}^{n_{pc}}$ is

$$r_{lap}^{\text{peak}} = \left(\vec{\phi}_{lap,\text{peak}}^+ \right)^T \vec{q}^{\text{peak}} . \quad (4.15)$$

Often we exclude certain spatial regions from the source unknowns using the projection matrix $P_{\text{peak}} \in \mathbb{R}^{n_{pc} \times n'_{pc}}$,

$$\vec{q}^{\text{peak}} = P_{\text{peak}} \vec{q}'^{\text{peak}} , \quad (4.16)$$

where $\vec{q}'^{\text{peak}} \in \mathbb{R}^{n'_{pc}}$. The matrix P_{peak} is formed from an identity matrix of size n_{pc} that is missing $n_{pc} - n'_{pc}$ columns. Because $P_{\text{peak}}^T P_{\text{peak}} = I_{n'_{pc}}$

$$P_{\text{peak}}^T \vec{q}^{\text{peak}} = \vec{q}'^{\text{peak}} .$$

Then the adjoint flux and source vectors are both shortened to n'_{pc} elements by substituting $\vec{q}^{\text{peak}} = P_{\text{peak}} \vec{q}'^{\text{peak}}$ into Equation 4.15,

$$\begin{aligned} r_{lap}^{\text{peak}} &= \left(\vec{\phi}_{lap,\text{peak}}^+ \right)^T P_{\text{peak}} \vec{q}'^{\text{peak}} \\ &= \left(P_{\text{peak}}^T \vec{\phi}_{lap,\text{peak}}^+ \right)^T \vec{q}'^{\text{peak}} \\ &= \left(\vec{\phi}_{lap,\text{peak}}^{+'} \right)^T \vec{q}'^{\text{peak}} . \end{aligned}$$

Thus we only need save the adjoint fluxes in $\vec{\phi}_{lap,\text{peak}}^{+'} \in \mathbb{R}^{n'_{pc}}$.

So far we have only discussed the solution to Equation 4.14 for a particular set of (l, a, p) indices. However, there are n_{lap} equations for the uncollided peaks, and we store the reduced adjoint $\vec{\phi}_{lap,\text{peak}}^{+'}$ for each as a row in matrix $S^{\text{peak}} \in \mathbb{R}^{n_{lap} \times n'_{pc}}$. (The order of the rows should follow the same convention as the flat and reshape functions introduced in section 4.2.1.) With this matrix, all the uncollided responses can be computed with the matrix-vector product:

$$\vec{r}^{\text{peak}} = S^{\text{peak}} \vec{q}'^{\text{peak}} .$$

4.5.2 Continuum sensitivities

The discrete form of the adjoint equation for the detector location l , detector direction a , and detector energy group g is

$$\mathbf{L}_s^+ \vec{\phi}_{lag}^+ = \vec{\sigma}_d^{lag} \quad . \quad (4.17)$$

The vectors $\vec{\phi}_{lag}^+$ and $\vec{\sigma}_d^{lag}$ have n_{gc} elements. \mathbf{L}_s^+ is a $n_{gc} \times n_{gc}$ matrix form of the transport operator, including uncollided and scattering physics. The adjoint source vector $\vec{\sigma}_d^{lag}$ is zero everywhere except for the n_g elements corresponding to the spatial cell where the detector is located. These elements contain the column of Σ_d corresponding to the energy of group g .

Applying the adjoint flux $\vec{\phi}_{lag}^+$ from Equation 4.17 gives the total response for detector location-direction-group (l, a, g) ,

$$r_{lag}^{\text{total}} = \left(\vec{\phi}_{lag}^+ \right)^T \vec{q}^{\text{total}} \quad .$$

However, we are interested in the two components of the total response,

$$r_{lag}^{\text{total}} = \sum_{p \in S_g} r_{lap}^{\text{peak}} + r_{lag}^{\text{cont}} \quad .$$

In the previous section, the peak response was computed. The expression for the continuum response is then

$$r_{lag}^{\text{cont}} = r_{lag}^{\text{total}} - \sum_{p \in S_g} r_{lap}^{\text{peak}} \quad ,$$

$$\left(\vec{\phi}_{lag, \text{cont}}^+ \right)^T \vec{q}^{\text{total}} = \left(\vec{\phi}_{lag}^+ \right)^T \vec{q}^{\text{total}} - \sum_{p \in S_g} \left(\vec{\phi}_{lap, \text{peak}}^+ \right)^T \vec{q}^{\text{peak}} \quad .$$

Substituting for \vec{q}^{total} from Equation 4.4, this equation in terms of the peak and

continuum sources is

$$\begin{aligned} \left(\vec{\phi}_{lag,cont}^+\right)^T \left(G \vec{q}^{peak} + \vec{q}^{cont}\right) &= \left(\vec{\phi}_{lag}^+\right)^T \left(G \vec{q}^{peak} + \vec{q}^{cont}\right) \\ &\quad - \sum_{p \in S_g} \left(\vec{\phi}_{lap,peak}^+\right)^T \vec{q}^{peak} \quad . \quad (4.18) \end{aligned}$$

Since \vec{q}^{peak} and \vec{q}^{cont} are independent variables, equating their terms separately from Equation 4.18 yields the two independent equations:

$$\left(\vec{\phi}_{lag,cont}^+\right)^T \vec{q}^{cont} = \left(\vec{\phi}_{lag}^+\right)^T \vec{q}^{cont} \quad , \quad (4.19)$$

and

$$\left(\vec{\phi}_{lag,cont}^+\right)^T G \vec{q}^{peak} = \left(\vec{\phi}_{lag}^+\right)^T G \vec{q}^{peak} - \sum_{p \in S_g} \left(\vec{\phi}_{lap,peak}^+\right)^T \vec{q}^{peak} \quad . \quad (4.20)$$

Since Equation 4.19 must be valid for an arbitrary (non-negative) \vec{q}^{cont} , we can equate the two adjoint fluxes

$$\vec{\phi}_{lag,cont}^+ = \vec{\phi}_{lag}^+ \quad . \quad (4.21)$$

This tells us that the sensitivities for the continuous spectrum source are the same as that for the total responses. This is unsurprising since the continuous-spectrum source can only contribute to the continuum response, not to the peak response.

Similarly, since Equation 4.20 must be valid for any \vec{q}^{peak} , the sensitivities of the continuum response on the peak sources are

$$\vec{\phi}_{lag,pc}^+ = G^T \vec{\phi}_{lag,cont}^+ = G^T \vec{\phi}_{lag}^+ - \sum_{p \in S_g} \vec{\phi}_{lap,peak}^+ \quad . \quad (4.22)$$

This tells us that the sensitivity of the continuum response to the peak sources is the same as the total response, with the exception of the removal of the photopeaks caused by uncollided particles. The multiplication by the G^T matrices removes groups which contain no peaks.

Again, certain spatial regions can be excluded from the source unknowns using the matrices P_{peak} and P_{cont} . Then the adjoint flux and source vectors are shortened to

n'_{pc} and n'_{gc} ,

$$\begin{aligned}
r_{lap}^{\text{cont}} &= \left(\vec{\phi}_{lag,pc}^+ \right)^\top \vec{q}^{\text{peak}} + \left(\vec{\phi}_{lag,cont}^+ \right)^\top \vec{q}^{\text{cont}} \\
&= \left(\vec{\phi}_{lag,pc}^+ \right)^\top P_{\text{peak}} \vec{q}'^{\text{peak}} + \left(\vec{\phi}_{lag,cont}^+ \right)^\top P_{\text{cont}} \vec{q}'^{\text{cont}} \\
&= \left(P_{\text{peak}}^\top \vec{\phi}_{lag,pc}^+ \right)^\top \vec{q}'^{\text{peak}} + \left(P_{\text{cont}}^\top \vec{\phi}_{lag,cont}^+ \right)^\top \vec{q}'^{\text{cont}} \\
&= \left(\vec{\phi}_{lag,pc}^{+'} \right)^\top \vec{q}'^{\text{peak}} + \left(\vec{\phi}_{lag,cont}^{+'} \right)^\top \vec{q}'^{\text{cont}} .
\end{aligned}$$

Similarly to the peak responses, we also store the reduced peak-source-to-continuum-response sensitivities $\vec{\phi}_{lag,pc}^{+'}$ for the (l, a, g) indices as a row in matrix $S^{\text{pc}} \in \mathbb{R}^{n_{lag} \times n'_{pc}}$. Likewise for the continuous-source-to-continuum-response sensitivities, we store $\vec{\phi}_{lag,cont}^{+'}$ for the (l, a, g) indices as a row in matrix $S^{\text{cont}} \in \mathbb{R}^{n_{lag} \times n'_{gc}}$. With this matrix, all the continuum responses can be computed with the matrix-vector product

$$\vec{r}^{\text{cont}} = S^{\text{pc}} \vec{q}'^{\text{peak}} + S^{\text{cont}} \vec{q}'^{\text{cont}} .$$

Finally, we can write the peak and continuum equations in one block matrix equation,

$$\begin{bmatrix} S^{\text{peak}} & 0 \\ S^{\text{pc}} & S^{\text{cont}} \end{bmatrix} \begin{bmatrix} \vec{q}'^{\text{peak}} \\ \vec{q}'^{\text{cont}} \end{bmatrix} = \begin{bmatrix} \vec{r}^{\text{peak}} \\ \vec{r}^{\text{cont}} \end{bmatrix} , \quad (4.23)$$

or simply

$$S \vec{q}' = \vec{r} .$$

If one knows at the start that there are only discrete-energy gamma sources, or alternately, that the source is only continuous spectrum, then this equation could be simplified accordingly.

An adjoint calculation, and hence an adjoint flux, is necessary for each element in the response vector \vec{r} . More precisely, for Equation 4.23 we would need as many adjoints as there are rows in the equation, that is n_{lag} .

4.5.3 Directional detector responses

The adjoint calculations to solve Equation 4.14 and Equation 4.17 are performed with Denovo.²⁷ Because the test problem has a large region of air where scattering is low

and the adjoint sources are localized, the discrete ordinates calculation is likely to suffer from ray effects. One method to reduce these effects is the first-collision-source method, where the uncollided flux is computed using a ray-tracing procedure. Then the scattering source from the uncollided flux is fed to the S_N solver. This is the method we used to compute all the adjoint fluxes.

In specifying the adjoint source for the uncollided flux solver, the location and DRF of a point detector is given. For the full 4π detector readings, this is sufficient control. However, for the collimated 2π measurements, we must provide additional source information. This could be done in two ways. One way might seek to accurately model the detector shielding by modeling the lead shield around the detector location. Instead of pursuing this route, we chose to specify an angularly-dependent point adjoint-source. Denovo has the ability to distribute the point source in n cosine bins, where the user must specify a reference direction, $n + 1$ bin edges, and the probability of each bin. For the 2π collimated detectors, we specified three cosine bins, with edges $\{-1, -\epsilon, \epsilon, 1\}$, with $0 < \epsilon \ll 1$. For the results in chapter 5, we set $\epsilon = 0.001$. The first or last cosine bin made the primary contribution. The middle bin was created to make the transition between forward and backward “fuzzy,” since collimation will produce a similar effect. Note that the simplified scalar flux forms of Equation 4.14 and Equation 4.17 are not sufficient to specify the source in this situation since it is not isotropic.

4.5.4 *Adjusting sensitivities for logarithmic transform*

The Bayesian inference algorithm depends on the sensitivities of the responses with respect to the model parameters. The adjoint flux gives the response sensitivities with respect to the source intensities. Since the model parameters are the logarithm of the source intensities, the adjoint fluxes must be adjusted to serve as the sensitivities.

To determine the adjustment, we examine the dependence of the i -th response $R(\vec{\alpha})_i$, which depends on the j -th model parameter $\vec{\alpha}_j$. In the logarithmically transformed system, this is equal to $\vec{\alpha}_j = \log_e(\vec{q}_j/q_0)$, where $q_0 > 0$ is a constant. The derivative of interest is

$$\frac{dR_i}{d\vec{\alpha}_j} = \frac{dR_i}{d \log_e(\vec{q}_j/q_0)} = \frac{dR_i}{d\vec{q}_j} \vec{q}_j \quad (4.24)$$

Thus, to convert the sensitivity with respect to the source strength \vec{q} to the model

parameter $\vec{\alpha}$, we must multiply the adjoint flux by the source strength, the unknown quantity in the IP solution process. The dependence of the sensitivities (or Jacobian) on the independent variables $\vec{\alpha}$ is the reason we are compelled to use nonlinear iterations. The (i, j) element of the sensitivity matrix is

$$[S_{\alpha}(\vec{\alpha})]_{ij} = \frac{dR_i}{d\vec{q}_j} \vec{q}_j = [S_{\alpha}]_{ij} q_0 e^{\vec{\alpha}_j} .$$

The mapping from state space to response space is

$$R(\vec{\alpha}) = S \vec{q} = S q_0 e^{\vec{\alpha}} ,$$

where the exponential $e^{\vec{x}}$ operates on the vector \vec{x} element-wise.

4.6 NONLINEAR ITERATION

Since $R(\vec{\alpha})$ is nonlinear due to the logarithmic transformation, the one-step linear process described in section 2.2.2 is insufficient, and a nonlinear optimization is need. As described in section 2.2.3, one approach to minimizing $Q(\vec{\alpha})$ is a Newton-type method. This is the method that we adopt for this work. These methods are able to use the gradient and Hessian information to more efficiently find a minimum. The equations developed in section 2.2.3 are more general than we require, so we describe here how these are simplified and computed.

In this case, C is block-diagonal since $C_{\alpha r}$ is zero. Then

$$C^{-1} = \begin{bmatrix} C_{\alpha}^{-1} & 0 \\ 0 & C_m^{-1} \end{bmatrix} , \quad (4.25)$$

and

$$Q(\vec{z}) = \vec{z}_{\alpha}^T C_{\alpha}^{-1} \vec{z}_{\alpha} + \vec{z}_r^T C_m^{-1} \vec{z}_r . \quad (4.26)$$

The gradient of Q is then

$$\vec{\nabla}_{\alpha} Q(\vec{z}) = C_{\alpha}^{-1} \vec{z}_{\alpha} + S^T C_m^{-1} \vec{z}_r . \quad (4.27)$$

and the Hessian is

$$\begin{aligned}\nabla_{\alpha}^2 Q(\vec{z}) &= C_{\alpha}^{-1} + S^T C_m^{-1} S + \frac{\partial S}{\partial \vec{\alpha}} C_m^{-1} \vec{z}_r \\ &\approx C_{\alpha}^{-1} + S^T C_m^{-1} S \quad ,\end{aligned}$$

where we have used the Gauss-Newton approximation that $\frac{\partial S}{\partial \vec{\alpha}} \approx 0$.

To efficiently compute the gradient, we solve the following two equations:

$$\begin{aligned}C_{\alpha} \vec{x}_{\alpha} &= \vec{z}_{\alpha} \quad , \\ C_m \vec{x}_r &= \vec{z}_r \quad .\end{aligned}$$

The gradient is then

$$\vec{\nabla}_{\alpha} Q(\vec{z}) = \vec{x}_{\alpha} + S^T \vec{x}_r \quad . \quad (4.28)$$

To efficiently compute the action of a Hessian on a vector \vec{p} , $\nabla_{\alpha}^2 Q(\vec{z}) \vec{p}$, we solve the following two equations:

$$\begin{aligned}C_{\alpha} \vec{x}_{\alpha} &= \vec{p} \quad , \\ C_m \vec{x}_r &= \vec{v} = S \vec{p} \quad .\end{aligned}$$

The matrix-vector product between the Hessian and \vec{p} is then

$$\nabla_{\alpha}^2 Q(\vec{z}) \vec{p} = \vec{x}_{\alpha} + S^T \vec{x}_r \quad . \quad (4.29)$$

In solving these linear systems, it is important to take advantage of the form of the coefficient matrix, especially when the number of source unknowns is large. In every problem tested, C_{α} and C_m were sparse. While the small number of responses utilized in our studies meant that storing C_m as a dense matrix is cheap, storing a dense C_{α} was often beyond the capacity of RAM or even hard-disk storage. Thus, C_{α} was stored as a sparse matrix, and sparse linear solvers were used for solving the C_{α} equations.

Bound-constrained optimization One side benefit of using the logarithmic transformation was the conversion of the optimization problem from a constrained to an unconstrained form (see section 4.2.2). In general, an unconstrained problem is simpler to solve than a constrained one. However, in our numerical experiments, using an

unconstrained solver frequently led to overflow errors. As seen in section 4.5.4, the exponential of the current $\vec{\alpha}_k$ iterate is needed to compute the sensitivities at each nonlinear step. Unfortunately, a moderately large step taken in the logarithmic space can easily cause overflow in floating point arithmetic. To avoid these large steps in logarithmic space, we use a bound-constrained Newton optimizer. These algorithms accept lower and upper inequality constraints on each component of the solution. These lower and upper bounds enclose a feasible set of points for the solution. While the iteration vector is within the feasible set, the iteration proceeds as a normal Newton method. In this case, the constraints are said to be inactive. However, when the next iterative step falls outside the feasible set, the violated constraints are then active and the iteration vector is modified to fall within or on the boundary of the feasible set. Some constraints can be active and others inactive at a particular iteration. See chapter 5 in Ref. [61] for a fuller description of constrained optimization.

It is reasonable to impose an upper bound because physical limits exist to the intensity of radiation that could be produced in a given volume. In the results presented in the following chapter, no lower bound was supplied, and the upper bound was $\alpha = 30$ or $q \approx 1 \times 10^{13}$ photons/cm²/s. This is well above any reasonable expectation for the source intensity in our laboratory experiment. For a different application, this value could be adjusted. In our experience, all of the constraints were inactive at convergence.

We use the `fmin_tnc` routine from the Scipy⁵⁹ numerical library for Python.⁷⁴ This function is based on the TNBC Fortran routine written by Nash⁹⁰ and available on Netlib. This routine uses a quasi-Newton method, only accepting the gradient of the cost functional as input, and building an approximation to the Hessian.

Local minima For a well behaved functional, Newton methods generally are successful in finding a minimum. Unfortunately, there is no guarantee that this minimum is a global minimum; often it is only a local minimum. Indeed, in the test problems investigated in this work, we often found that the optimization routine found a local minimum. This was the case when the initial iteration was set to the mean of the prior distribution. However, when a much larger initial guess was supplied, the optimizer found the global minimum, or at least a minimum much nearer to the global minimum. More details of the initial guess' effect on the identified minimum will be discussed in the next chapter.

4.7 COMPUTING THE POSTERIOR COVARIANCE

After the best model parameter values $\vec{\alpha}^*$ are found by the optimization routine, the posterior covariance matrices can be computed using Equations 2.51 to 2.53. We are primarily interested in the covariance data for the model parameters $\vec{\alpha}$ contained in C_α^{be} . One modification of Equation 2.51 is necessary; the covariance and sensitivity matrices must be evaluated at $\vec{\alpha}^*$ [112, p. 69–70]. Thus, the best-estimate for the $\vec{\alpha}$ covariance matrix is

$$C_\alpha^{be} = C_\alpha - C_{\alpha d}(\vec{\alpha}^*) C_d(\vec{\alpha}^*)^{-1} C_{\alpha d}(\vec{\alpha}^*)^T$$

where $S_\alpha = S_\alpha(\vec{\alpha}^*)$. While C_α is sparse, C_α^{be} is dense. Given the large dimensions of C_α , only certain elements of C_α^{be} are computed, usually only the diagonal.

As stated in chapter 1, the application of the method considered in this thesis is to map the radiation source spectral and spatial distribution in a room. We conducted an experiment in NC State University's Burlington Engineering Laboratories room 2144 using small ^{137}Cs and ^{60}Co gamma sources, measuring the radiation using a sodium iodide detector. This experiment is more fully described in Appendix c. From the measured data, we use the source mapping method described in chapter 4 to estimate the location and intensity of the sources by energy.

Before presenting the results of the source mapping, we describe the layout of the room and sources, and the details of the numerical simulations of the room configuration.

5.1 DESCRIPTION OF PROBLEM

The following sections provide a summary of the experiment; Appendix c has a more thorough description.

Room layout The setting for this experiment was 2144 Burlington Engineering Labs at NC State University, which was formerly a laboratory and had been unoccupied in recent years. The room is roughly a cube with side length of 4 meters. The walls are concrete block. There is a pillar on the wall next to the windows. The room is mostly empty, with the exception of the counter top and cabinets on the east wall. A fuller description of the room is given in Appendix c.2.3, which also includes the material description. A Cartesian coordinate system is used to model the space, with the origin located at the northwest bottom corner of the room. The x -axis points south and the y -axis points east. Figure c.3 shows the room's components with the overlaid


Cartesian coordinates.

Source placement For this experiment, we have available some low-activity gamma sources from the nuclear engineering student laboratory. These sources were ^{137}Cs and ^{60}Co . ^{137}Cs produces one peak at 662 keV and ^{60}Co produces two gamma energies, 1173 and 1332 keV. Table c.2 gives further information about these sources, and Table c.4 provides the activities of each of the sources. The six ^{137}Cs sources had a combined activity of about 3.5 μCi , while the older shorter-lived ^{60}Co sources summed to less than 1 μCi . Since each source had a low activity, these eleven sources were combined to increase the detector count rates and decrease the time needed per measurement.


The ^{137}Cs sources were placed in a small box in the southwest corner of the room. We chose a point source configuration to give the source mapping method an easier challenge at first. The five ^{60}Co sources were arranged in a line along the bottom of the chalkboard on the west wall. Unfortunately, one of the ^{60}Co sources was significantly stronger than the others, making the effect of the line less pronounced. The locations and intensities are pictured in Figure 5.1. The coordinate locations are given in Table c.6.

Safety concerns limited our ability to take measurements with a more widely distributed source. While a spread-out source would be a more fitting challenge for the developed methods, it is not practical to handle a radioactive source of this nature in the laboratory setting to which we had access.

Detectors We used a sodium iodide scintillation detector coupled with an MCA to capture the gamma spectra at six different locations around the room. In addition to bare detectors that look equally in all directions, we used lead collimation to confine detection to specific directions. Specifically, we arranged the collimator so that the detector had a 2π hemispherical view. At each location, seven spectra were possible – one 4π and six 2π views ($\pm x$, $\pm y$, and $\pm z$). Due to time constraints, not all views at each location were measured. The measurement data is contained in the electronic archive attached to this thesis. The uncollided portion of the data can be seen in Figure D.2.



(a) ^{137}Cs 661.7 keV emissions.



(b) 1173 keV and ^{60}Co 1332 keV emissions.

Figure 5.1: The true distribution of the ^{137}Cs and ^{60}Co sources, as represented on the Cartesian computational mesh. The ^{137}Cs sources are grouped together in a point source, and the ^{60}Co sources are aligned in a row parallel to the x -axis. The natural logarithm of the source intensity is plotted, that is, $\log_e(q_{pijk}/q_0)$, where q_0 is a unit scalar needed to make the argument of the logarithm dimensionless and q_{pijk} is the source in peak p in spatial cell (i, j, k) . The units of q_{pijk} and q_0 are particles emitted per second. In reality, the blue cells have no sources, but they are shown here with small source intensities to aid comparison with the results presented below.

5.1.1 *Permutations of the source mapping problem*

Even though the ^{137}Cs and ^{60}Co source distribution does not change, we solve this problem in a variety of different ways to gain confidence in the mapping algorithm and establish best practices. They are broken into two main sections.

1. The simpler problem uses simulated “measurements.” The flux at the detector is computed using MCNP. Beginning with synthetic data allows us to use more detector locations, since we are not constrained by long counting times. It also eliminates the biases and errors found in measured data. While the use of synthetic data does not fully test the method, our use of it is not a blatant inverse crime because the “measurements” are computed with Monte Carlo and the forward model with a deterministic transport code. However, It could be labeled an inverse misdemeanor. We perform two tests with synthetic data. The first uses only the uncollided or photopeak information from the detector spectrum. Second, we use the full spectrum, both photopeaks and scattered radiation. In both cases, the responses are the scalar flux at the detector. We could apply the detector efficiency to both sides of the equation, but this would not change the solution.
2. The second test uses the experimentally measured data. This is more difficult than the synthetic data because the measurements are limited. Furthermore, the errors in this data are much larger than the model-mismatch in the previous trial. Because of the difficulty in peak identification in a noisy spectrum, some of the data points that we would expect to be present are not detected. For the measured data, we only map the source with the uncollided photopeaks, ignoring the continuum spectrum. We do this for a number of reasons. First, the uncollided flux dominates the spectrum since most of the room is air. Second, we began but did not complete the integration of the full NaI detector response function into the method. The responses in this case are still the scalar flux at the detector. We convert the measured count rate into flux using the process described in Appendix D.

5.2 TRANSPORT SIMULATION DETAILS

As mentioned in the previous section, we used both Monte Carlo and deterministic radiation transport codes in this test problem. The Monte Carlo code `MCNP` simulated the synthetic “measurements” and allowed a comparison against the measured data. The deterministic transport code `Denovo` provided the forward model used in the inversion method. The following sections give some of the details of these calculations.

5.2.1 *MCNP simulation of the experiment*

For the ^{137}Cs and ^{60}Co source distribution, we executed two separate `MCNP` simulations. The first was for the 4π bare detector responses. The six detectors were modeled using the `F5` point detector tallies. All six tallies were included in one simulation. The second simulation was for the collimated 2π detectors. There would be many ways to implement such a tally. We chose to include a 10 cm lead cube centered at each detector location. Then we added point detector tallies at the center of each face of the cube, separated from the cube by 0.5 mm. The amount of lead around each tally point is roughly equal to the amount used in the experiment. Each tally used the standard energy grid for the fine bins, from 100 to 2900 keV (see Appendix C.4.2). The number of particle histories was greater than 10^6 , which made the relative error of the significant tallies less than 5%. The `MCNP` input file is included in the electronic archive attached to this thesis.

The scalar fluxes computed by `MCNP` at the various locations are shown in Figures 5.2, 5.9 and 5.11.

5.2.2 *Denovo adjoint discrete ordinates model*

`Denovo`,²⁷ the multigroup discrete ordinates code from Oak Ridge National Laboratory, performed the adjoint simulations needed for the source mapping inversion. A library of adjoint flux solutions is built for every source (location, direction, energy group) combination. This library was built before the inverse source algorithm was executed. The code uses the multigroup energy discretization. We used the fine-group shielding library from `SCALE`,¹¹ but the groups below 100 keV and above 1.332 MeV are ignored. The boundaries of these groups are listed in Table 5.1.

For the uncollided gamma simulations, it would be most accurate to use the

Table 5.1: The energies E and group index i for the 16 energy group structure used for this test problem. The structure is based on the fine-group gamma shielding cross section library from SCALE.¹¹ Energies, in keV, are the upper limit of the respective group. The groups above ^{60}Co 's highest peak (1.332 MeV) and below 100 keV are omitted.

i	E	i	E	i	E
1	1.44×10^3	7	7.00×10^2	13	3.00×10^2
2	1.33×10^3	8	6.00×10^2	14	2.60×10^2
3	1.20×10^3	9	5.12×10^2	15	2.00×10^2
4	1.00×10^3	10	5.10×10^2	16	1.50×10^2
5	9.00×10^2	11	4.50×10^2		1.00×10^2
6	8.00×10^2	12	4.00×10^2		

^a The last value is the lower boundary of the lowest-energy group.

pointwise cross sections at the respective energy of each line. However, all deterministic calculations involving the discrete-energy gammas were performed with group-collapsed cross sections. Since the groups are not wide, this introduces minimal error. Each of the three gamma energy lines appears in a distinct energy group.

Denovo uses a Cartesian spatial grid. The grid size for this problem was 52 by 53 by 54 cells. Although most cells were roughly cubes with 10 cm sides, the mesh was not uniform because of material interfaces. All of the results are plotted using this spatial grid. We use the step characteristic spatial discretization because it produces uniformly non-negative fluxes.

The adjoint source was specified using the first collision source feature of Denovo. This feature requires that a point source (or multiple point sources) be supplied. In this case, we had to execute each adjoint source one at a time, so we only specified one source per execution. For the 4π view, an isotropic point source is sufficient. However, for the 2π views, the point source must be anisotropic. In Denovo, this is accomplished by supplying a probability distribution function of the cosine of the source emission with respect to a specified vector, as described in section 4.5.3. For the discrete-ordinates calculation, the angular quadrature was a Gaussian-Legendre product set with $N = 8$.

The Denovo input file is included in the electronic archive attached to this thesis.

5.3 INVERSE MAPPING DETAILS

In this section we describe a few of the specific choices we made in applying the method developed in chapter 4.

5.3.1 Peak and continuum

As summarized in Equation 4.23, it is in general possible to have both peak and continuum sources and responses. However, only peak sources apply when dealing with gamma-emitting radionuclides. These discrete energy photons can scatter, producing a continuum response. Thus, in our circumstance, Equation 4.23 simplifies to

$$\begin{bmatrix} \mathbf{S}^{\text{peak}} \\ \mathbf{S}^{\text{pc}} \end{bmatrix} \vec{q}'^{\text{peak}} = \begin{bmatrix} \vec{r}^{\text{peak}} \\ \vec{r}^{\text{cont}} \end{bmatrix} , \quad (5.1)$$

since $\vec{q}'^{\text{cont}} = 0$. As stated above and demonstrated later, the continuum response in this case is not as significant as the peak response, so the simplified equation

$$\mathbf{S}^{\text{peak}} \vec{q}'^{\text{peak}} = \vec{r}^{\text{peak}} , \quad (5.2)$$

may be sufficient to compute the source map. We solve the test problem with synthetic data using only the peak responses and also solve it with the peak and continuum responses. The solutions in both cases turn out to be similar.

5.3.2 Prior distributions on $\vec{\alpha}$ and \vec{r}

The Bayesian method outlined in earlier chapters requires that we provide a prior distribution on both the model parameters $\vec{\alpha}$ and the responses \vec{r} . The distribution of the measured responses is more obvious because it corresponds to the measurement error. However, the distribution on the model parameters is more ambiguous. As we have already stated, the prior on the source intensities will be modeled as a log-normal distribution, which requires two parameters, the mean and standard deviation. We chose a small mean and a large standard deviation. The small mean reflects our belief that the source is likely not spread out over the entire region. (If there is reason to believe the source is widely distributed, then a larger prior mean would be recommended.) The large prior standard deviation shows that we have relatively little confidence in

the prior mean. In this case, we set the prior mean to 10^{-2} particles/second, which translates in logarithmic scale to $\alpha_0 = \log_e 10^{-2} \approx -4.6$. The log-normal prior standard deviation is set to 6, which gives a $\pm 1\sigma$ interval of $[2.5 \times 10^{-5}, 4.0]$ and a $\pm 3\sigma$ interval of $[1.5 \times 10^{-10}, 6.6 \times 10^5]$. Thus, the prior distribution spans a wide range of scales.

We assume that all the correlations between sources intensities in each cell and group or peak are zero. There are reasons that such correlations would exist. If two peaks are produced by the same nuclide, then a strong positive correlation should be assigned for the source in each cell for those two peaks. Although this is the case for the two ^{60}Co peaks, we treat them as uncorrelated in these test problems. If the source was believed to be spatially smooth, one could assign a positive correlation between neighboring cells.

The response standard deviations are based on estimates of the experimental error. For the MCNP synthetic measurements, we make 3 assumptions to compute the error:

1. The count time for all “measurements” is the same.
2. The total counts (but not the count rate) in each photopeak follows Poisson statistics.
3. We can tune the fictional “count time” to adjust the counting errors.

With those assumptions, we chose to set the error of the largest measurement to a specified relative error τ_{\min} . The Poisson standard deviation for a count rate r is

$$\sigma_r = \frac{\sqrt{c}}{t} = \frac{\sqrt{r \cdot t}}{t} = \sqrt{\frac{r}{t}} \quad ,$$

where c is the number of counts and t is the count time. Thus, the count time t to ensure that the largest measurement has error τ_{\min} is computed as

$$t = \frac{1}{\tau_{\min}^2 \max(\vec{r}_m)} \quad .$$

In the tests that follow, $\tau_{\min} = 0.01$. If the elements in the vector \vec{r}_m span many orders of magnitude, then this algorithm may produce relative errors for the smaller components that are greater than one. This is unphysical for Poisson variables. These elements are commonly the detector readings from directions pointing away from the source, and so we would expect them to be zero or very small. Here, we use the

criterion $\tau \geq 1$ to indicate very small entries. For those entries, the absolute standard deviation is set equal to a value similar in magnitude to the largest of the small values. This important step allows the fitting procedure to search within a wider window for the nearly-zero entries. Otherwise, the residual from these entries can overwhelm the chi-square functional, thereby diluting the impact of larger, more important data points.

The uncertainties for the experimentally-measured data are available via the propagated Poisson error, described further in section 5.4.3.

5.3.3 *Newton method implementation*

As stated in section 4.6, we use the bound constrained quasi-Newton optimization routine `fmin_tnc` to optimize $Q(\vec{\alpha})$. This routine takes as input

- the functional $Q(\vec{\alpha})$,
- the gradient of $Q(\vec{\alpha})$, $\vec{\nabla}_{\alpha} Q$,
- the initial guess $\vec{\alpha}_{\text{init}}$,
- lower and upper bounds for each component of $\vec{\alpha}$, and
- the maximum number of function evaluations.

We use the function's default convergence criterion of

$$f_{\text{tol}} = \sqrt{\epsilon} \quad , \quad x_{\text{tol}} = \sqrt{\epsilon} \quad , \quad (5.3)$$

where ϵ is the machine epsilon [94, §2.5.3], which is approximately $\sqrt{\epsilon} \approx 1.49 \times 10^{-8}$ for double-precision floating-point numbers. The stopping criteria are

$$|f(\vec{x}_n) - f(\vec{x}_{n-1})| \leq f_{\text{tol}} \quad , \quad \|\vec{x}_n - \vec{x}_{n-1}\| \leq x_{\text{tol}} \quad . \quad (5.4)$$

In practice, the routine typically stopped on the f test, before the x criterion was achieved. The maximum number of function evaluations was set to 2000, but this total was rarely reached.

5.3.4 *Restricting cells permitted to contain a source*

One way to narrow the search for the source is to exclude spatial cells which can not contain source. This process is described in chapter 4, especially section 4.5.

Equation 4.16 shows that this transformation can be accomplished simply by the multiplication of an identity matrix with missing columns. In the test problem, the physical rationale for removing such cells from the model parameter list $\vec{\alpha}$ is that we assume that the sources are not airborne. They must be located near or in the walls, floor, ceiling, or cabinets. Thus, we eliminate the spatial cells in the interior of the room for all peak energies. The excluded cells are evident in the source plots in the following sections, for example, Figure 5.4. The excluded cells appear as the white regions in z -planes 6 through 47.

The Denovo spatial mesh has $n_c = 52 \cdot 53 \cdot 54 = 148\,824$ cells. When the interior cells are excluded, there are $n'_c = 84\,039$ remaining cells where the source intensity and spectrum are sought. About 44% of the cells have thus been excluded.

5.4 RESULTS FOR CS-137 AND CO-60 SOURCE MAPPING

This section presents the results of the source mapping algorithm applied to the ^{137}Cs and ^{60}Co distributions. First we describe the results with the synthetic “measurements” generated by MCNP, and then we present those results using the experimental measurements.

5.4.1 Synthetic data: photopeak responses only

Here we present the simplest of the test problems, with the use of only the photopeak synthetic measurements. This corresponds to Equation 5.2. For this problem, there are 126 rows in this linear system (6 locations \times 7 directions \times 3 energy peaks). Using a uniform initial guess of $\vec{\alpha}_{\text{init}} = 10$, the Newton method optimized the Q functional. At the optimal point $\vec{\alpha}^*$, the posterior covariances were computed.

The measured responses are compared to the best-fit responses in Figure 5.2. The response computed via $S^{\text{peak}} \vec{q}^{\text{peak}}$ resulting from the true source distribution is also shown. As this figure shows, the best-fit and measurements match well for the larger fluxes, above roughly 10^{-2} . As the bottom subplot shows, the relative difference between the best-fit predictions and measurements is typically less than 10%. However, for the smaller fluxes, the agreement is much weaker. These smaller fluxes are observed when the detector is facing away from the source. The discrepancy is due to the different way in which the collimation is simulated in MCNP and Denovo.

Thus, it is not surprising that they are not identical. However, it is important that in either case the flux is small, close to zero. For these small responses, the values computed with the adjoints and the true source are much smaller than 10^{-8} . This is because the Denovo-estimated contribution for sources behind a collimated detector is zero, while the lead shields in MCNP allow a small contribution.

The mean of the posterior of the source distribution for each of the three gamma lines is shown in a volume rendering in Figure 5.3. Notice that for the ^{137}Cs peak, the source is predicted to be located in the southwest bottom corner, as it truly is. However, the predicted source is spread over a few separated cells. For the two ^{60}Co peaks, the algorithm predicts the source around the true location.

The distribution can be examined more thoroughly in Figures 5.4A to 5.4G, which provide in x - y slices the mean and the standard deviation of the posterior source distribution. In general, the predicted source distribution matches well, if not perfectly, the true distribution pictured in Figure 5.1.

The posterior uncertainty is notable largely because it is unchanged from the prior distribution. The measurements have not reduced the source uncertainty for the vast majority of cells. While it would clearly be preferable to have reduced the uncertainty in these low-source cells, it is not surprising that we have seen so little reduction given the underdetermined state of the equations. This large uncertainty also hints at other possible solutions. Since the cells adjacent to the predicted source cells have such a wide uncertainty, we have not ruled out the possibility of a source being located in one of the neighboring cells. Instead, the particular cell with the source should be taken as the center of a neighborhood where the source is likely to be found.

In the cells with significant source, the uncertainty is greatly reduced. Table 5.2 shows the estimate of the source at each energy, summing the contributions over space since most of the sources of the same energy are nearby. This value can be compared to the true source activity. In this case, the ^{137}Cs source magnitude is overestimated by a factor of roughly 90. The reported uncertainty does not reflect this discrepancy. However, after examining the location of the ^{137}Cs source prediction, we see that it is predicted to be near the outside of the wall. This means that the flux emerging from the wall on its way to the detector will be significantly reduced due to attenuation within the wall so that achieving the same detector response would require a much stronger source intensity. Thus, the overprediction is explainable and would have been eliminated had we restricted the sources to exist only adjacent to wall and floor cells

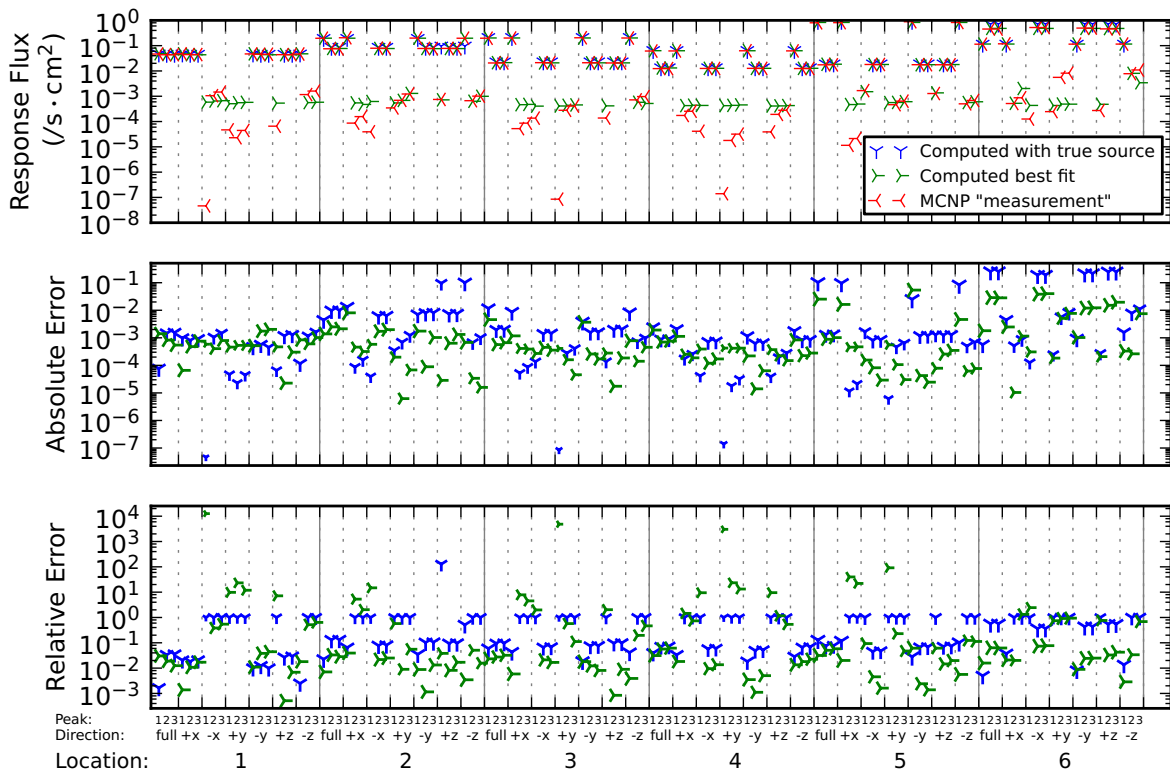


Figure 5.2: A comparison of the measured and computed responses for the source mapping using only photopeak synthetic responses. The continuum responses are ignored in this case. The synthetic measurements obtained with MCNP are in red. The green “computed best fit” responses are generated using the Denovo forward model and the best-estimate source prediction. Likewise, the “computed with true source” responses are generated with the Denovo forward model and the actual source configuration.

For measurements where no “computed with true source” data point is present, the value is off the bottom of the scale. The absolute error is the magnitude of the difference between the computed values and the measurement. The relative error is the ratio of the absolute error to the measurement value. The surface area of the markers in the bottom two subplots is proportional to the logarithm of each measurement. The computed responses with the largest relative error correspond to measurements with small magnitudes.






Figure 5.3: Volume rendering of the mean of the posterior distribution of the source using the photopeak synthetic responses. The source prediction has suppressed source regions in cells around the primary source cells. This results in dark areas near the sources visible in the insets.



(a) ^{137}Cs 661.7 keV emissions.

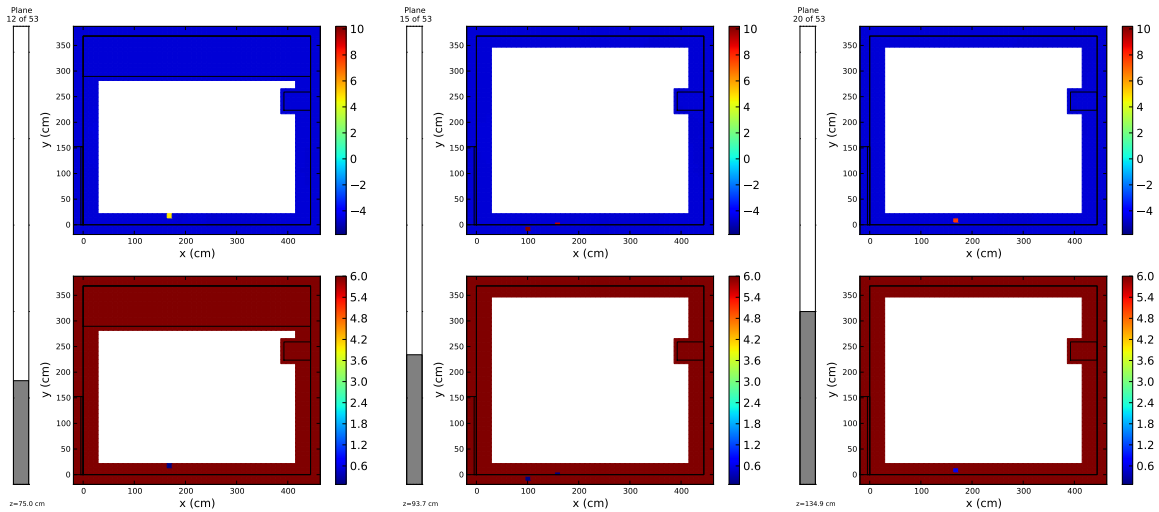
(b) ^{137}Cs 661.7 keV emissions.

Figure 5.4: Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using the photopeak synthetic responses. The natural logarithm of the source intensity is plotted, that is, $\log_e(q_{pijk}/q_0)$, where q_0 is a unit scalar needed to make the argument of the logarithm dimensionless. The units of q_{pijk} and q_0 are particles emitted per second.



(c) ^{60}Co 1173 keV emissions.

(d) ^{60}Co 1173 keV emissions.



(e) ^{60}Co 1332 keV emissions.

(f) ^{60}Co 1332 keV emissions.

(g) ^{60}Co 1332 keV emissions.

Figure 5.4: Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using the photopeak synthetic responses. (continued)

Table 5.2: Comparing the predicted source emissions against the true intensities. The intensities are in total emissions per gamma line per second. They are summed over all spatial cells. The true source intensities are computed with the activities from Table c.4 and the yields from Table c.2. Unless noted otherwise, the initial guess was set to ten, $\vec{\alpha}_{\text{init}} = 10$.

	Energy (keV)		
	661.7	1173.2	1332.5
True	$1.075 \pm 0.044 \times 10^5$	$2.19 \pm 0.14 \times 10^4$	$2.20 \pm 0.14 \times 10^4$
Peak	$8.9 \pm 0.7 \times 10^6$	$2.38 \pm 0.05 \times 10^4$	$3.90 \pm 0.27 \times 10^4$
Peak, $\vec{\alpha}_{\text{init}} = 1$	$1.09 \pm 0.10 \times 10^5$	$2.12 \pm 0.11 \times 10^4$	$2.12 \pm 0.49 \times 10^4$
Peak & continuum	$2.80 \pm 0.10 \times 10^4$	$1.21 \pm 0.06 \times 10^4$	$9.67 \pm 0.41 \times 10^3$
Experiment	$7.69 \pm 0.73 \times 10^5$	$4.21 \pm 0.52 \times 10^4$	$2.88 \pm 0.14 \times 10^4$

in the room’s interior. The two ^{60}Co source estimates are much closer to the true value differing by 9% and 77%, that constitute excellent and good agreement, respectively, in the context of the present application.

Varying initial guess As the high uncertainties in much of the posterior suggest, this problem can have many local minima in which the optimizer can get stuck. Indeed, using an initial guess equal to the prior mean, the optimizer found a local minimum that was far from the true solution. Through experimentation, we found that an initial source guess much larger than the prior mean allowed the Newton optimizer to come much closer to the functional’s value for the true source distribution.

To quantify the effect of different starting Newton iterates, we ran a sequence of optimizations with a range of values for the initial guess selected to cover about 9 decades in linear scale. Each initial guess was uniform in space and energy. Our experience of improved performance with a larger value of the prior mean cited above is confirmed with these results. First, Figure 5.5 shows the chi-square functional at the optimal $\vec{\alpha}^*$. A lower value, implying a closer fit, indicates a better solution to the IP. At lower values of the initial guess, the chi-square functional is almost 100 times larger than the smallest achieved chi-square. Since there are 126 responses, we should expect the chi-square value to be close to 126. The best achieved chi-square is within a




Figure 5.5: Change in chi-square and the number of function evaluations with respect to the intensity of the uniform initial guess for optimization.

factor of 5 to 10 from this optimal value, while the chi-square of the poorer converged distributions (resulting from initial guess smaller than about 1) is about 100 times too large. The optimizer seems to find the local minimum significantly faster, in fewer function calls, than it requires to find the true minimum.

Finally, Figure 5.8 presents the summed source intensity for each initial guess value (quantities similar to those reported in Table 5.2). The low-magnitude initial guesses tend to drastically underpredict the amount of source present. This is because these solutions typically predict the source in regions close to the detectors. Thus, not only is the source intensity incorrect, but the spatial distribution is not accurate as well. A sample prediction for the ^{137}Cs peak is given in Figure 5.6.

As the initial guess increases, the total source prediction becomes rather accurate.





Figure 5.6: The local minimum reached when using $\vec{a}_{\text{init}} = -6$. The slices in the x - y plane are the mean and standard deviation of the posterior distribution of the 661.7 keV emissions from ^{137}Cs using the uncollided, photopeak responses. The source is primarily predicted in cells close to detector locations.

The solutions produced by an initial guess of 1 to 6 seem to produce the most accurate solutions, where the sources are not located within the walls. The chi-square value in this range is around the best-achieved value of 500. This is the region with the most accurate predictions. The predicted source map with $\vec{a}_{\text{init}} = 1$ is shown in Figure 5.7, and the source intensities summed over space are given in a row in Table 5.2. The predicted source map matches almost exactly with the true source. The chi-square value is 505, and the chi-square per degree of freedom is 4.0. This prediction matches the true source better than the predictions with larger \vec{a}_{init} because the sources are not located within the concrete walls.

As the initial guess increases beyond 6, the source intensities are more often overpredicted. As seen in the example above for $\vec{a}_{\text{init}} = 10$, the reason for the overprediction is the placement of the source deep within the wall. Thus, this error is not as unreasonable as it first appears.



(a) ^{137}Cs 661.7 keV emissions. (b) ^{60}Co 1173 keV emissions. (c) ^{60}Co 1332 keV emissions.

Figure 5.7: Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using the photopeak synthetic responses with a favorable initial guess, $\vec{\alpha}_{\text{init}} = 1$.

5.4.2 Synthetic data: photopeak and continuum responses

In the previous subsection, the source mapping was performed using only photopeak responses, solving Equation 5.2. In this section, both the photopeak and continuum responses are considered via Equation 5.1. For the synthetic data, the continuum responses are simply the scattered fluxes binned according to the multigroup energy structure. They only contain the collided photons; they exclude the uncollided photons. The continuum responses are also computed using the F5 tally in MCNP, binned in the fine energy grid. These results are then summed into the SCALE energy groups. There are still three photopeak results $n_p = 3$, and there are 16 energy groups $n_g = 16$, so $n_{\vec{g}} = 19$. Using the same 6 locations and 7 detection directions per location, there are now a total of $6 \cdot 7 \cdot 19 = 798$ responses.

As for the previous results, we again used a uniform initial guess of $\vec{\alpha}_{\text{init}} = 10$. The Newton method required 1832 function calls and 184 Newton iterations to converge. At the optimal point $\vec{\alpha}^*$, the posterior covariances were computed.

The MCNP-measured responses are compared to the best-fit responses in Figure 5.9. In these figures, the first 126 response indices are for the peak responses, and the

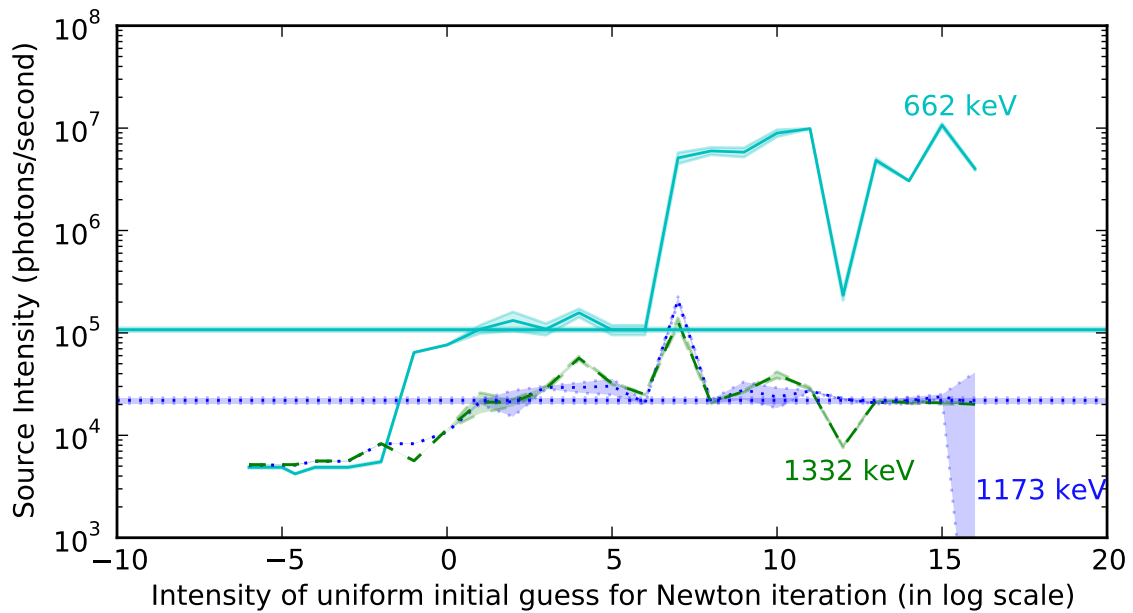


Figure 5.8: Change in source intensity summed over space with respect to the initial guess for optimization. The upper, cyan horizontal line corresponds to the true value of the 662 keV peak, while the lower, blue line corresponds to the intensity of the two ^{60}Co peaks. The shaded area around a line represents the 1σ confidence interval about the mean.

remainder are the continuum response. Figure 5.9A shows the photopeak responses, while Figure 5.9B shows the continuum responses. Also shown are the responses computed with the adjoints based on the true source distribution. Again the best-fit and measurements match well for the larger fluxes, above roughly 10^{-2} . The relative difference between the measurements and best-fit predictions is poorer than for the photopeak-only results (section 5.4.1). In that case, most of the large responses agreed to within 10%. Here, the relative difference for many of the large responses is between 10% and 100%.


The mean and standard deviation of the posterior distribution for the three energies are plotted in Figures 5.10A to 5.10L. The predicted source is similar to the true distribution and that predicted with only the photopeak responses. There are some false positives in areas without a real source, but the magnitude of the source in these cells is much lower than the value in the true source region (remember that the distribution is plotted on a logarithmic scale). This prediction is not noticeably better than for the photopeak-only prediction, and the uncertainties are no lower. We attribute this lack of change to the low amount of scattering in this problem. The uncollided flux at the detector is larger than the scattered flux, usually by a factor of 10–1000. For a medium with more highly scattering material and less air, this would presumably not be true.

The predicted source for the ^{60}Co is not a line, as the true spatial arrangement of the ^{60}Co sources were. However, the prediction does show some spatial spread of the source, in the z -planes surrounding the true height. Although the actual line source is not accurately reconstructed, the results provide some evidence that the source is not simply a single point.


The source at each energy summed over space is also provided in Table 5.2. The ^{137}Cs prediction is less than the true value by a factor of four. The predicted location for the source is about 25 centimeters closer to the nearby detectors (locations 5 and 3) than the true source, so it is not surprising that the source strength at this location is underestimated. The ^{60}Co peak intensities are within a factor of two of the true value.

5.4.3 *Experimental measurements*

The above results used synthetic “measurements” that were computed by MCNP. In this section, we present the results of the source mapping using the experimental




(a) Photopeak responses.




(b) Continuum responses.

Figure 5.9: Measured and computed responses for the source mapping with synthetic photopeak and continuum responses. For indices where no “computed with true source” data point is present, the value is off the bottom of the scale. See Figure 5.2 for an explanation of the absolute and relative error plots.



(a) ^{137}Cs 661.7 keV emissions.

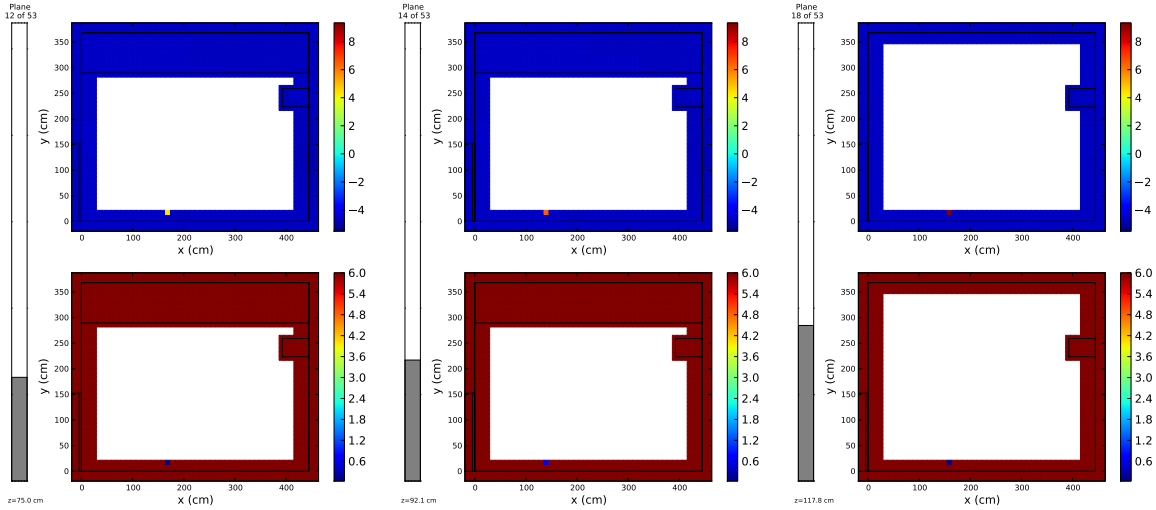
(b) ^{137}Cs 661.7 keV emissions.



(c) ^{137}Cs 661.7 keV emissions.

(d) ^{137}Cs 661.7 keV emissions.

Figure 5.10: Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using both the photopeak and continuum synthetic responses. The natural logarithm of the source intensity is plotted, that is, $\log_e(q_{pijk}/q_0)$, where q_0 is a unit scalar needed to make the argument of the logarithm dimensionless. The units of q_{pijk} and q_0 are particles emitted per second.




(e) ^{60}Co 1173 keV emissions. (f) ^{60}Co 1173 keV emissions. (g) ^{60}Co 1173 keV emissions.


Figure 5.10: Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using both the photopeak and continuum synthetic responses. (continued)

measurements. Appendix c contains a fuller description of the experiment. Since the source map prediction for the peak and peak-and-continuum results were very close, we will only use the photopeak responses in this test. This frees us from dealing with the full two-dimensional detector response function. However, we must still convert the scalar flux to detector counts, which is a function of the detector efficiency.

Detector efficiency The deterministic radiation transport code computes scalar fluxes, while the detector data is in count rates. To translate between the two, we must compute the intrinsic peak efficiency for the NaI detector at each of the three photon energies. The intrinsic peak efficiency is a measure of the full-energy scores in a detector versus the number of photons entering the detector at that energy (also known as the inward partial current). The peak intrinsic detector efficiency depends on the source-to-detector configuration, which is unknown to us as we seek to determine the source distribution itself. Thus, an average efficiency is employed. We computed this efficiency with the F8 tally in `MCNP` using a spatially-distributed source to capture the averaging process. The other step in converting from flux to counts is to change the scalar flux to inward partial current. This is a function of the mean chord length



(h) ^{60}Co 1332 keV emissions. (i) ^{60}Co 1332 keV emissions. (j) ^{60}Co 1332 keV emissions.



(k) ^{60}Co 1332 keV emissions. (l) ^{60}Co 1332 keV emissions.

Figure 5.10: Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using both the photopeak and continuum synthetic responses. (continued)

in the detector. A full description of this process is available in Appendix D, as well as a comparison of the experimental measurements with modeled values.

We express the response in Equation 5.2 as scalar fluxes, just as we did in the synthetic data cases above. This means that the measured count rates must be converted to scalar fluxes using the detector efficiency. The converted measured scalar fluxes are plotted in Figure 5.11.

Limited data The measured data set is smaller than the synthetic data set. Because of time constraints on the experiment, not all 6 locations \times 7 directions = 42 configurations were counted. The list of configurations experimentally counted is shown in Figure C.9, and the specific counts used for this experiment can be found at the bottom of Figure 5.11. As this second figure shows, there are several measured data points which are missing, even when the modeled flux was significant ($> 10^{-2}$ particles/cm²·s). The missing data is a consequence of the difficulty in finding peaks in noisy spectra. Since the source activities were low and the count times were limited, some of the peaks were not resolved. We would not expect to count a peak when the detector is facing away from the source, but all of the 4π detectors should be expected to register all three source peaks. Unfortunately, this is not the case here. However, this deficiency in the experimental data would not have materialized if the deployed sources were stronger or the counting times were extended beyond what was deemed reasonable in the context of this proof-of-principle study. Recognizing that this deficiency is not intrinsic to our new approach we circumvented the adverse effects of the limited number of experimental data points by supplementing them with a few synthetic points that we believe would have been experimentally obtained (to within experimental error) had we had access to stronger sources or a longer time frame. The process of supplementing the measured data with synthetic values is detailed below.

Because of the limitations in the data, attempts to map the source using only the measured data were unsuccessful. The predicted source distributions had many false positive source locations while missing the true locations. Given the several missing data points visible in Figure 5.11, the poor results are not surprising. However, we still wish to make some use of the data. Thus, we substituted 14 of the synthetic data points for missing photopeaks in the experimental data, out of 69 total experimental points. If the sources were more intense or massive, it would be less likely that one of the photopeaks was missed in the measured counts. Unfortunately, the




Figure 5.11: Measured and modeled scalar flux for the experiment in 2144 Burlington. If an experiment value is missing from the plot, its value is zero. The model in this case is MCNP.

radioisotope sources that were available to us for use in this experiment were weak. Even with long count times (several hours to a few days, see Figure c.5), the peaks were often barely visible, especially for configurations with a large source-to-detector separation, due to the $1/r^2$ geometric attenuation. The photopeaks from background radiation were often of the same magnitude as the peaks of interest, aggravating the challenge of locating the true peaks in the detector's response. To supplement the measured data, we replaced the missing data points with the synthetic data described in section 5.4.1. A photopeak data point was considered missing if it had a value of zero and the modeled value for that location-direction-peak combination was greater than 5×10^{-3} photons/cm²/s (see Figure 5.11). For the 662 keV energy, there were 2 missing data points, 6 missing for 1173 keV, and 6 missing for 1332 keV. Over all energies, 14 modeled values were above the cutoff. Hence, the number of synthetic points supplementing the experimental data adds up to 14 out of 69 total points. To make the synthetic measurements more realistic, they were randomly perturbed according to a Gaussian distribution with standard deviation equal to the estimated error of each synthetic measurement. In choosing the missing data points, we have used the information about the source location. Thus, in an end-use application, substituting the synthetic data is not a viable approach, but it is a reasonable alternative for the proof-of-principle purposes of this study.

The uncertainties in the non-zero experimental data are the values obtained through the propagation of the experimental Poisson error. For the synthetic supplemental data, we used the same method as for the purely synthetic case to estimate Poisson uncertainties for a fictitious counting time. For data points with a value of zero, the absolute uncertainty is assigned to be a small value, in this case 2×10^{-4} particles/cm²·s. The choice of this value controls the level of fitting on the zero responses by the optimization routine.

The computed and measured responses are given in Figure 5.12. The effect of the small absolute uncertainty assigned to zero data is evident in the suppression of the best-fit responses corresponding to measurements of value zero. These points, shown in green, are held to about 10^{-3} in this case.

The source distributions are shown for the three gamma lines in Figure 5.13. Although there are some small false positives, the general shape of the true source is predicted, perhaps not as well as for the previous, fully synthetic cases. This difference is largely due to the number of measurements available to achieve an





Figure 5.12: Measured and computed responses for the source mapping with experimental data. For indices in the top subplot where a data point is missing, the value is off the bottom of the scale. See Figure 5.2 for an explanation of the absolute and relative error plots. The Denovo-computed response using the true source distribution for the 662 keV peak at location 2 looking in the +z direction is about 100 times too large. This is because the ^{137}Cs source is located within the fuzzy region around $\cos \theta = \mu = 0$ modeled in Denovo (see section 4.5.3). This inaccuracy is also present in Figure 5.2 and Figure 5.9A, but it is obscured there by the larger number of responses.


accurate optimization of the predicted versus measured responses.

The source intensities in each peak are given in the last row of Table 5.2. The 662 keV peak intensity is overestimated by a factor of 7 because the source location is within the wall. The 1173 keV peak intensity is roughly double the true value. The 1332 keV intensity is the best estimate, being in error by 30%. Not coincidentally, the predicted spatial distribution for this peak is quite accurate.



(a) ^{137}Cs 661.7 keV emissions.

(b) ^{60}Co 1173 keV emissions.



(c) ^{60}Co 1332 keV emissions.

(d) ^{60}Co 1332 keV emissions.

Figure 5.13: Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using the photopeak responses of the experimental data. The natural logarithm of the source intensity is plotted, that is, $\log_e(q_{pijk}/q_0)$, where q_0 is a unit scalar needed to make the argument of the logarithm dimensionless. The units of q_{pijk} and q_0 are particles emitted per second.

5.5 CARPET SOURCE TEST

The source distribution for the previous test problem (see Figure 5.1) is highly localized. The ^{137}Cs sources are in a small box, and the linear arrangement of ^{60}Co sources approximates a point since one of the ^{60}Co sources is much more radioactive than the others. However, one of the advantages of the source mapping method is the flexibility to map more broadly distributed sources. To test this ability, we designed a source distribution with a planar source, which we term the carpet source because of its resemblance to a rectangular piece of carpet lying on the floor. The setting for the carpet source is still 2144 Burlington Engineering Labs. The sources are contained within

$$x \in [0, 350] \quad , \quad y \in [0, 250] \quad ,$$

and the ^{137}Cs source is at $z = 1$ and the ^{60}Co source is at $z = 2$ (all dimensions in cm). Within the carpet source region, the source strength varies linearly in space. The intensity of the ^{137}Cs source is described by

$$S''_{137\text{Cs}}(x, y) = \frac{x}{350} \quad \left[\frac{\text{particles emitted}}{\text{cm}^2 \text{ s}} \right] \quad ,$$

and the intensity of each of the ^{60}Co energy peaks is described by


$$S''_{60\text{Co}}(x, y) = \frac{y}{250} \quad \left[\frac{\text{particles emitted}}{\text{cm}^2 \text{ s}} \right] \quad .$$

This distribution is depicted in Figure 5.14.


For safety and security concerns, we do not have access to such a source. Thus, our only recourse is to simulate the carpet source. We generate a set of synthetic measurements in MCNP, just as in section 5.4.1. These scalar fluxes are shown as the 'MCNP "measurement"' in Figure 5.15.

Just as in the previous tests, we maximized the posterior distribution using the quasi-Newton function. It converged in 131 Newton iterations and 1276 function evaluations. The fitted response vector is plotted along with the synthetic measurements in Figure 5.15. The fit for the majority of the large fluxes is within 10% of the measured values, while even some of the smaller fluxes have a good fit.

The full source prediction for the three peaks is given in Figures 5.16A to 5.16K.



(a) ^{137}Cs 661.7 keV emissions.



(b) 1173 keV and ^{60}Co 1332 keV emissions.

Figure 5.14: The true distribution of the ^{137}Cs and ^{60}Co sources for the carpet source problem. In reality, the blue cells have no sources, but they are shown here with small source intensities to aid comparison with the results presented below.

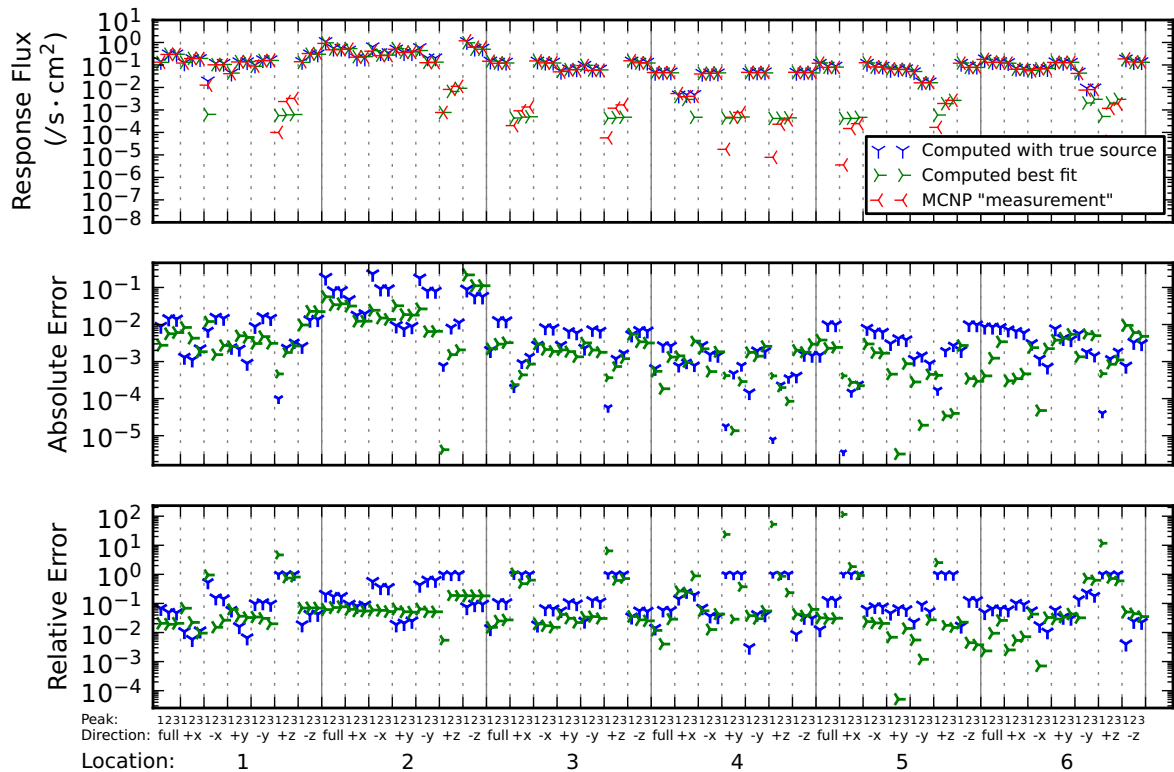



Figure 5.15: Measured and computed responses for the source mapping with the carpet source. For indices where no “computed with true source” data point is present, the value is off the bottom of the scale. See Figure 5.2 for an explanation of the absolute and relative error plots.

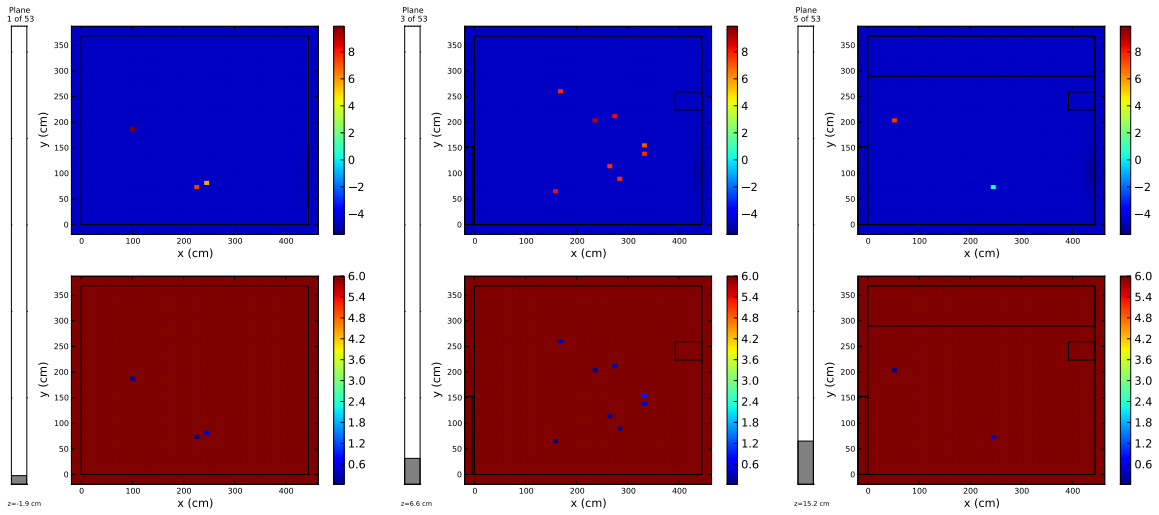
The ^{137}Cs source prediction is skewed to the $+x$ side of the rectangle, with no source being predicted for $x \leq 100$ cm. This bias is consistent with the true tilt of the source, but it fails to indicate that there is some source in the $x \leq 100$ cm region.

The ^{60}Co source predictions show a similar bias toward the $+y$ side of the source region. There is a source location or two that are falsely predicted under the cabinet. The most striking feature of the solution is the speckled pattern. The optimization routine apparently favors giving one cell a stronger source rather than distributing it smoothly over a few neighboring cells. While the speckled pattern does not accurately predict the smoothness of the true source, it does convey that it is spatially distributed over the floor.

The source intensity for each energy peak summed over space is given in Table 5.3.




(a) ^{137}Cs 661.7 keV emissions. (b) ^{137}Cs 661.7 keV emissions. (c) ^{137}Cs 661.7 keV emissions.

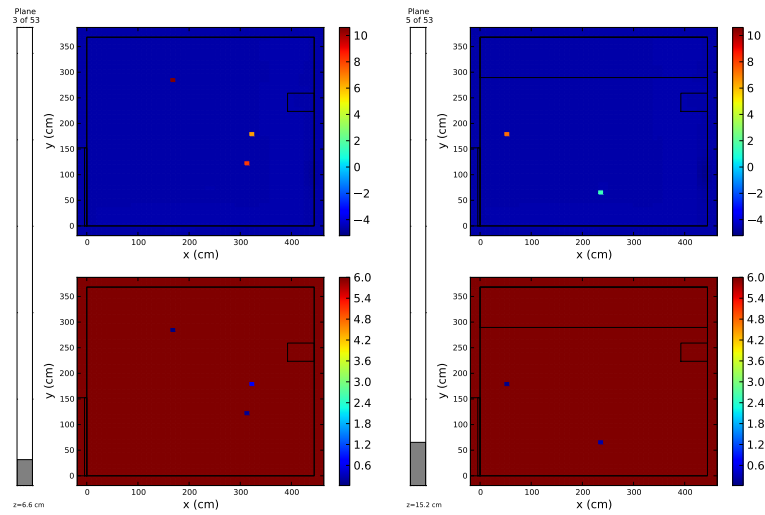


(d) ^{60}Co 1173 keV emissions. (e) ^{60}Co 1173 keV emissions. (f) ^{60}Co 1173 keV emissions.

Figure 5.16: Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using the photopeak responses for the floor-distributed source. The natural logarithm of the source intensity is plotted, that is, $\log_e(q_{pijk}/q_0)$, where q_0 is a unit scalar needed to make the argument of the logarithm dimensionless. The units of q_{pijk} and q_0 are particles emitted per second.



(g) ^{60}Co 1332 keV emissions. (h) ^{60}Co 1332 keV emissions. (i) ^{60}Co 1332 keV emissions.



(j) ^{60}Co 1332 keV emissions. (k) ^{60}Co 1332 keV emissions.

Figure 5.16: Slices in the x - y plane of the mean (upper plot) and standard deviation (lower plot) of the posterior distribution of the source using the photopeak responses for the floor-distributed source. (continued)

Table 5.3: Comparing the predicted source emissions per energy per second summed over space against the true values for the carpet test problem.

	Energy (keV)		
	661.7	1173.2	1332.5
True	$4.4 \pm 0.0 \times 10^4$	$4.4 \pm 0.0 \times 10^4$	$4.4 \pm 0.0 \times 10^4$
Peak	$6.3 \pm 0.2 \times 10^4$	$5.3 \pm 0.3 \times 10^4$	$8.6 \pm 0.2 \times 10^4$

All of the predicted intensities are within a factor of two of the true value. The overprediction is consistent with the source being predicted within the floor, whereas the true source was on top of the floor and so had only air attenuation before reaching the detector.

Here we discuss the results presented in the previous chapter and suggest extensions of this work.

6.1 CONCLUSIONS

Predicting source locations and intensities The locations of the source predictions when using the synthetic data were mostly accurate. There were few false alarms, in which a source is predicted in a location where there is no source. The source intensities for the false alarms were insignificant as compared to the intensity at the primary locations. The ^{137}Cs prediction usually looked like the point source that it should be, with the first test being an exception. The ^{60}Co line source was never accurately mapped, but this is not surprising given the large differences in intensity in the ^{60}Co sources. Although the predictions typically had one cell with the primary ^{60}Co source, there were also surrounding cells with a non-negligible ^{60}Co sources which gave the impression of a more distributed source.

As Table 5.2 demonstrated, the source magnitudes were generally on the correct order of magnitude, although they were not usually within the predicted uncertainty. In fact, the predicted uncertainty was consistently too small. In cases where the source prediction was deep within a wall, the source intensity was much larger than the true value. The larger intensity compensated for the extra material attenuation that was not present in the true configuration. One way to avoid this would be to eliminate the interior of the walls from the list of permitted source cell unknowns.

The elimination of most of the void cells was effective in limiting the source to desired locations. It had the added benefit of decreasing memory and computation demands of the algorithm.

Modest reduction in uncertainty The most striking conclusion when focusing on the standard deviation of the posterior is its near similarity with the prior (except where source cells were identified). The measured data had no effect on the uncertainty in the majority of cells. This implies that the resolution of the data is rather poor, at least in terms of this problem. The collimated detectors certainly increased the resolving power of the data, but more and better information is necessary to shrink the uncertainty on more cells. This extra data could be provided by a more sensitive directionally-dependent detector.

The high uncertainty also points to a multiplicity of solutions. For instance, in the ^{137}Cs source posterior, the cell with the primary source (which has a much smaller uncertainty) is surrounded on all sides by cells with the prior uncertainty. We interpret this to mean that the primary source cell could easily be shifted a cell or two without much change in the response. Thus, even if the primary cell has a small uncertainty, the surrounding cells with high uncertainty caution against putting too much faith into the predicted location and intensity.

Choosing the initial guess Although an obvious choice, using the mean of the prior distribution as the initial guess was an effective way to find a local minimum in the chi-square functional which was far from the true solution. Our experience indicates that selecting a uniform guess with a magnitude several orders of magnitude greater than the prior mean allows the Newton optimizer to find a minimum that is close to the true solution. With a favorable initial guess, the predicted source is accurate within the posterior uncertainty.

Peak and continuum responses For the problem of finding the ^{137}Cs and ^{60}Co sources in 2144 BEL, it appears that using only the photopeak responses produced by the uncollided flux is sufficient. The addition of the continuum response in the form of the scattered flux does not improve the source distribution prediction. Since the uncollided flux at the detector is significantly larger than the scattered flux, this result is reasonable. The scattered radiation would play a more important role in systems with shields or sources that were otherwise blocked from direct view of the detector.

Overall efficacy of the source mapping algorithm In the tests presented here, the source mapping algorithm produces predictions which, although far from perfect, capture the

most important spatial and energy features of the true source distribution. Due to the ill-posed, underdetermined nature of the problem, these order-of-magnitude estimates, for instance, of the total source intensities, are the best that could be expected with a small detector array. The location of the predicted sources was often within 20 cm of the true location; the worst estimate was in error by 80 cm. For a room with side lengths of 4 meters, accuracy in the tens of centimeters, that is, a few to 20%, is helpful for narrowing the search area. The optimal chi-square value was a useful metric in judging the proximity of the predicted distribution to the true solution. Finally, the uncertainties in the predicted source distribution indicated the existence of other nearby solutions, which could only be eliminated with higher-resolution data.

6.2 FUTURE WORK

A more complete detector response function To include the continuum response (instead of the photopeaks only) with the experimental data, the two-dimensional detector response would be required. We avoided the use of the full DRF by using only the peak intrinsic efficiency. The DRF can be generated with Monte Carlo methods, but empirical adjustments are still typically required, for instance, to describe the Gaussian broadening of the peaks.⁸⁴

As mentioned above, the DRF is dependent on the source-to-detector positioning. To capture this effect fully would require computing a separate DRF for each source cell, or at least modifying a base DRF for the effects of each unique orientation. This would be the most accurate DRF approach. However, one could average the DRF over various orientations, and use the resulting DRF to unfold the detector spectrum into the scalar flux at the detector. Then this flux could be fed to the source mapping algorithm. This should be more accurate than simply using the photopeaks, but it lacks some fidelity because the DRF must be somehow averaged.

The unfolding of the spectra to produce the scalar flux at the detector is a related inverse problem. The DRFs have been traditionally applied to convert from the flux to the response, which is the simpler direction. Detector covariance counting data could improve the results of the unfolding.

Finding a near-global minimum As already observed, there is a very real possibility of finding a local minimum when optimizing an underdetermined set of equations. In

this work, we relied on a heuristic rule which was effective in helping the Newton optimizer to find a minimum close to the true solution. However, this heuristic could be eliminated by the use or combination of other optimization methods. A brute force or Monte Carlo selection of initial guesses might be means to gain confidence in the Newton solution. Also, global stochastic optimizers might be useful, although they are typically much slower.

Proper response probability distribution The variance-based Bayesian method used in this work assumes Gaussian distributions on the model parameters and responses. The state variables were modeled with a log-normal distribution because they can not be negative. Likewise, the responses cannot be negative, and so a Gaussian distribution for these variables is not theoretically valid. In practice, the Gaussian distribution was serviceable, but a log-normal distribution or a non-variance-based method would be a sounder approach to adopt in future work.

Blind test The experimental validation in this work used non-blind data, in which the author both conducted the experiment and produced the source mapping. For a true validation of the methods, the data should be gathered by a third party to eliminate bias in the mapping procedure.

Other related inverse problems This method could be extended to related inverse problems discussed in section 1.2. If the radiation source affects the geometry or material cross sections, then the underlying problem is genuinely nonlinear and additional transport simulations would be necessary at each iteration of the optimization process. Adjoint flux solutions for the initial configuration would not be relevant for later iterations as the cross sections or geometrical details of the system change.

BIBLIOGRAPHY

- [1] A. Bachrach, R. He, and N. Roy. "Autonomous flight in unknown indoor environments". *International Journal of Micro Air Vehicles* 1.4 (2009), pp. 217–228.
- [2] Ahmed Badruzzaman. "Computational methods in nuclear geophysics". *Progress in Nuclear Energy* 25.2-3 (1991), pp. 265–290.
- [3] Allan F. M. Barton. *CRC handbook of polymer-liquid interaction parameters and solubility parameters*. CRC Press, 1990.
- [4] George I. Bell and Samuel Glasstone. *Nuclear reactor theory*. Malabar, FL: Robert E. Krieger Publishing, 1970.
- [5] M.J. Berger et al. "XCOM: Photon cross sections database". *NIST Standard Reference Database version 1.5* (2010).
- [6] Joseph Bertrand. *Calcul des probabilités*. Gauthier-Villars, 1889.
- [7] Åke Björck. *Numerical methods for least squares problems*. Philadelphia: Society for Industrial & Applied Mathematics, 1996.
- [8] K.C. Bledsoe, J.A. Favorite, and T. Aldemir. "Using the Schwinger inverse method for solutions of inverse transport problems in two-dimensional cylindrical geometries". *Annals of Nuclear Energy* 36.7 (2009), pp. 966–973.
- [9] Keith C. Bledsoe and Jeffrey A. Favorite. "A Comparison of the Differential Evolution and Levenberg-Marquardt Methods for Solving Inverse Transport Problems with Several Unknowns in Cylindrical Geometries". *Transactions of the American Nuclear Society* 101 (2009), pp. 411–413.
- [10] Keith C. Bledsoe and Jeffrey A. Favorite. "A Hybrid Differential Evolution/Levenberg-Marquardt Method for Solving Inverse Transport Problems". *Transactions of the American Nuclear Society* 102 (2010).
- [11] Steven M. Bowman. "SCALE 6: Comprehensive Nuclear Safety Analysis Code System". *Nuclear Technology* 174.2 (2011), pp. 126–148.
- [12] Dan G. Cacuci. *Sensitivity and Uncertainty Analysis: Theory*. CRC, 2003.
- [13] D.G. Cacuci and M. Ionescu-Bujor. "Best-estimate model calibration and prediction through experimental data assimilation – I: Mathematical framework". *Nuclear Science and Engineering* 165.1 (2010), pp. 18–44.

- [14] D.G. Cacuci and M. Ionescu-Bujor. "Sensitivity and Uncertainty Analysis, Data Assimilation, and Predictive Best-Estimate Model Calibration". In: *Handbook of Nuclear Engineering*. Ed. by D.G. Cacuci. Springer, 2010.
- [15] National Nuclear Data Center. "information extracted from the NuDat 2 database" (accessed 2011).
- [16] M.B. Chadwick, M. Herman, P. Obložinský, et al. "ENDF/B-VII.1 Nuclear Data for Science and Technology: Cross Sections, Covariances, Fission Product Yields and Decay Data". *Nuclear Data Sheets* 112.12 (2011). Special Issue on ENDF/B-VII.1 Library, pp. 2887–2996.
- [17] M.B. Chadwick, P. Obložinský, M. Herman, et al. "ENDF/B-VII.0: Next Generation Evaluated Nuclear Data Library for Nuclear Science and Technology". *Nuclear Data Sheets* 107.12 (2006). Ed. by P. Obložinský and M. Herman, pp. 2931–3118.
- [18] M.I. Chahine. "A general relaxation method for inverse solution of the full radiative transfer equation". *Journal of the Atmospheric Sciences* 29.4 (1972), pp. 741–747.
- [19] Richard Threlkeld Cox. "Probability, frequency and reasonable expectation". *American Journal of Physics* 14.1 (1946), pp. 1–13.
- [20] D.E. Cullen, J.H. Hubbell, and L. Kissel. EPDL97: *The evaluated photon data library, '97 version*. Tech. rep. UCRL-50400. Lawrence Livermore National Laboratory, 1997.
- [21] RH Dicke. "Scatter-hole cameras for x-rays and gamma rays". *The Astrophysical Journal* 153 (1968), p. L101.
- [22] P.A.M. Dirac. "Approximate rate of neutron multiplication for a solid of arbitrary shape and uniform density, Part I: General theory". In: *The Collected Works of P.A.M. Dirac, 1924-1948*. Ed. by R.H. Dalitz. The Collected Works of P.A.M. Dirac v. 1. Cambridge University Press, 1995, pp. 1115–1128.
- [23] YF Du et al. "Evaluation of a Compton scattering camera using 3-D position sensitive CdZnTe detectors". *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 457.1 (2001), pp. 203–211.

- [24] J.J. Duderstadt and L.J. Hamilton. *Nuclear reactor analysis*. Wiley, 1976.
- [25] *Eckert & Ziegler Reference & Calibration Sources: Product Information*. Tech. rep. Valencia, CA: Eckert & Ziegler Isotope Products, 2011.
- [26] M.B. Emmett and J.C. Wagner. "MONACO: A New 3-D Monte Carlo Shielding Code for SCALE". *Transactions of the American Nuclear Society* 91 (2004), pp. 701–703.
- [27] Thomas M. Evans et al. "Denovo: A New Three-Dimensional Parallel Discrete Ordinates Code in SCALE". *Nuclear Technology* 171.2 (2010), pp. 171–200.
- [28] J.A. Favorite, K.C. Bledsoe, and D.I. Ketcheson. "Surface and Volume Integrals of Uncollided Adjoint Fluxes and Forward-Adjoint Flux Products". *Nuclear science and engineering* 163.1 (2009), pp. 73–84.
- [29] J.A. Favorite and R. Sanchez. "An inverse method for radiation transport". *Radiation protection dosimetry* 116.1-4 (2005), p. 482.
- [30] Jeffrey A. Favorite. "Identification of an Unknown Material in a Radiation Shield Using the Schwinger Inverse Method". *Nuclear Science and Engineering* 152.1 (2006), pp. 106–117.
- [31] Jeffrey A. Favorite. "Identification of unknown densities in a source/shield system using the Schwinger inverse method". *Transactions of the American Nuclear Society* 93 (2005), pp. 588–589.
- [32] Jeffrey A. Favorite. "Using the Schwinger variational functional for the solution of inverse transport problems". *Nuclear Science and Engineering* 146.1 (2004), pp. 51–70.
- [33] Jeffrey A. Favorite and David I. Ketcheson. "Using the Levenberg-Marquardt Method for the Solution of Inverse Transport Problems". *Transactions of the American Nuclear Society* 95 (2006), p. 527.
- [34] Robin P. Gardner and Fusheng Li. "Use of the CEARXRF GUI-Based Monte Carlo-Library Least-Squares (MCLLS) Code for the Micro-Focused EDXRF analyzer". *X-Ray Spectrometry* 40.6 (2011), pp. 405–410.

- [35] R.P. Gardner and A. Sood. "A Monte Carlo simulation approach for generating NaI detector response functions (DRFs) that accounts for non-linearity and variable flat continua". *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 213 (2004), pp. 87–99.
- [36] Brian Gough. *GNU scientific library: reference manual*. Network Theory, 2009.
- [37] A. Griewank and A. Walther. *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*. Society for Industrial and Applied Mathematics, 2008.
- [38] Charles W. Groetsch. *Inverse problems: activities for undergraduates*. Washington, DC: The Mathematical Association of America, 1999.
- [39] C.W. Groetsch and D. Colton. *Inverse problems in the mathematical sciences*. Vol. 52. Vieweg Braunschweig, 1993.
- [40] Jacques Hadamard. "Sur les problèmes aux dérivées partielles et leur signification physique". *Princeton University Bulletin* 13.1 (1902), pp. 49–52.
- [41] P.C. Hansen. *Rank-deficient and discrete ill-posed problems: numerical aspects of linear inversion*. Society for Industrial & Applied Mathematics, 1998.
- [42] Per Christian Hansen. *Discrete Inverse Problems: insight and algorithms*. Fundamentals of Algorithms. Society for Industrial & Applied Mathematics, 2010.
- [43] W.K. Hastings. "Monte Carlo sampling methods using Markov chains and their applications". *Biometrika* 57.1 (1970), pp. 97–109.
- [44] T. He, R.P. Gardner, and K. Verghese. "The Monte Carlo–Library Least-Squares approach for energy-dispersive x-ray fluorescence analysis". *Applied Radiation and Isotopes* 44.10-11 (1993), pp. 1381–1388.
- [45] G.T. Herman. *Fundamentals of computerized tomography: image reconstruction from projections*. Springer Verlag, 2009.
- [46] JH Hubbell. "Review of photon interaction cross section data in the medical and biological context". *Physics in Medicine and Biology* 44 (1999), R1.
- [47] JH Hubbell et al. "Atomic form factors, incoherent scattering functions, and photon scattering cross sections". *Journal of Physical and Chemical Reference Data* 4 (1975), p. 471.

- [48] John H. Hubbell. *Bibliography of photon total cross section (attenuation coefficient) measurements 10 eV to 13.5 GeV*. Tech. rep. 5437. National Institute of Standards and Technology, 1994, pp. 1907–1993.
- [49] John H. Hubbell. “Review and history of photon cross section calculations”. *Physics in Medicine and Biology* 51 (2006), R245.
- [50] John H. Hubbell and William J. Veigele. *Comparison of theoretical and experimental photoeffect data 0.1 keV to 1.5 MeV*. Tech. rep. 901. US Department of Commerce, National Bureau of Standards, 1976.
- [51] Jérôme Idier, ed. *Bayesian approach to inverse problems*. Digital signal and image processing. Wiley, 2008.
- [52] I.C.F. Ipsen. *Numerical matrix analysis: Linear systems and least squares*. Philadelphia: Society for Industrial and Applied Mathematics, 2009.
- [53] F. James et al. “MINUIT reference manual”. CERN *Program Library Long Writeup D 506* (1994), p. 1998.
- [54] K.D. Jarman et al. “Bayesian Radiation Source Localization”. *Nuclear technology* 175.1 (2011), pp. 326–334.
- [55] Edwin T. Jaynes. “Information Theory and Statistical Mechanics”. *Physical Review* 106.4 (1957), pp. 620–630.
- [56] Edwin Thompson Jaynes. *Probability theory: the logic of science*. Cambridge University Press, 2003.
- [57] Harold Jeffreys. *Theory of Probability*. International series of monographs on physics. Clarendon Press, 1998.
- [58] Matthew A. Jessee. “Cross-Section Adjustment Techniques for BWR Adaptive Simulation”. PhD thesis. North Carolina State University, 2008.
- [59] Eric Jones, Travis Oliphant, Pearu Peterson, et al. *SciPy: Open source scientific tools for Python*. 2001–.
- [60] Benjamin Jowett. *The Republic of Plato*. Clarendon press, 1881.
- [61] C.T. Kelley. *Iterative methods for optimization*. Frontiers in applied mathematics. Society for Industrial & Applied Mathematics, 1999.

- [62] Dale Klein and Michael Corradini. “Fukushima Daiichi: ANS committee report” (2012).
- [63] G.F. Knoll. *Radiation Detection and Measurement*. John Wiley & Sons, 2010.
- [64] Erwin Kreyszig. *Introductory functional analysis with applications*. Wiley Classics Library. John Wiley & Sons, 1989.
- [65] WJM de Kruijf and JL Kloosterman. “On the average chord length in reactor physics”. *Annals of Nuclear Energy* 30.5 (2003), pp. 549–553.
- [66] H.H. Ku. “Notes on the use of propagation of error formulas”. *Journal of Research of the National Bureau of Standards. Section C: Engineering and Instrumentation* 70C (4 1966), p. 263.
- [67] J.R. Lamarsh. *Introduction to nuclear reactor theory*. Addison-Wesley, 1966.
- [68] E.W. Larsen. “Solution of multidimensional inverse transport problems”. *Journal of Mathematical Physics* 25.1 (1984), p. 131.
- [69] E.W. Larsen. “Solution of the Inverse Problem in Multigroup Transport Theory”. *Journal of Mathematical Physics* 22.1 (1981), pp. 158–160.
- [70] E.W. Larsen. “The inverse source problem in radiative transfer”. *Journal of Quantitative Spectroscopy and Radiative Transfer* 15.1 (1975), pp. 1–5.
- [71] Charles L. Lawson and Richard J. Hanson. *Solving least squares problems*. Philadelphia: Society for Industrial & Applied Mathematics, 1995.
- [72] Elmer E. Lewis and Warren F. Miller Jr. *Computational methods of neutron transport*. La Grange Park, IL: American Nuclear Society, 1993.
- [73] M.I.A. Lourakis. *levmar: Levenberg-Marquardt nonlinear least squares algorithms in C/C++*. <http://www.ics.forth.gr/~lourakis/levmar/>. accessed Nov. 2010. Dec. 2009.
- [74] Mark Lutz. *Programming Python*. O’Reilly Series. O’Reilly Media, 2010.
- [75] *MAESTRO-32 MCA Emulator for Microsoft Windows 2000 Professional and XP Professional Software User’s Manual*. Tech. rep. Ortec, 2006.
- [76] J. Mattingly and Eric S. Varley. “Synthesis of the Feynman-Y Neutron Multiplicity Metric Using Deterministic Transport”. *Transactions of the American Nuclear Society* 98 (2008), p. 572.

- [77] John Mattingly and Dean J. Mitchell. "A Framework for the Solution of Inverse Radiation Transport Problems". *IEEE Transactions on Nuclear Science* 57.6 (2010), pp. 3734–3743.
- [78] NJ McCormick. "Inverse radiative transfer problems: a review". *Nuclear Science and Engineering* 112.3 (1992).
- [79] NJ McCormick. "Unified approach to analytical solutions of three inverse transport problems". *Progress in Nuclear Energy* 34.4 (1999), pp. 425–430.
- [80] MCNP — *A General Monte Carlo N-Particle Transport Code, Version 5*. Los Alamos National Laboratory. 2003.
- [81] William Menke. *Geophysical data analysis: discrete inverse theory*. Academic Press, 1989.
- [82] N. Metropolis and S. Ulam. "The monte carlo method". *Journal of the American statistical association* 44.247 (1949), pp. 335–341.
- [83] N. Metropolis et al. "Equation of state calculations by fast computing machines". *The journal of chemical physics* 21 (1953), p. 1087.
- [84] W.A. Metwally, R.P. Gardner, and A. Sood. "Gaussian Broadening of MCNP Pulse Height Spectra". *Transactions of the American Nuclear Society* 91 (2004), pp. 789–790.
- [85] Carl D. Meyer. *Matrix analysis and applied linear algebra*. Philadelphia: Society for Industrial and Applied Mathematics, 2000.
- [86] Karen A. Miller and William S. Charlton. "An inverse transport approach to radiation source location for border security". In: *Annual Meeting of the European Safeguards Research and Development Association*. Aix-en-Provence, France, 2007.
- [87] Karen Ann Miller. "An inverse source location algorithm for radiation portal monitor applications". PhD thesis. Texas A&M University, 2010.
- [88] M. Morales et al. "Monte Carlo and least-squares methods applied in unfolding of X-ray spectra measured with cadmium telluride detectors". *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 580.1 (2007), pp. 270–273.
- [89] J.E. Morel. "A hybrid collocation-Galerkin- S_n method for solving the Boltzmann transport equation". *Nuclear Science and Engineering* 101.72 (1989).

- [90] S.G. Nash. "Newton-type minimization via the Lanczos method". *SIAM Journal on Numerical Analysis* (1984), pp. 770–788.
- [91] Stephen J. Norton. "A general nonlinear inverse transport algorithm using forward and adjoint flux computations". *IEEE Transactions on Nuclear Science* 44.2 (1997), pp. 153–162.
- [92] Stephen J. Norton. "Iterative inverse scattering algorithms: methods of computing Fréchet derivatives". *The Journal of the Acoustical Society of America* 106 (1999), p. 2653.
- [93] DE Peplow, RP Gardner, and K. Verghese. "Sodium iodide detector response functions using simplified Monte Carlo simulation and principal components". *Nucl. Geophys* 8.3 (1994), p. 243.
- [94] A. Quarteroni, R. Sacco, and F. Saleri. *Numerical mathematics*. Texts in applied mathematics. Springer, 2007.
- [95] T. Douglas Reilly. "Nondestructive assay of holdup". In: *Passive Nondestructive Assay of Nuclear Material: 2007 Addendum*. Ed. by Douglas Reilly. LA-UR-07-5149. 2007.
- [96] W.A. Rhoades and D.B. Simpson. *The TORT Three-Dimensional Discrete Ordinates Neutron/Photon Transport Code*. 3rd. ORNL/TM-13221. 1997.
- [97] Clive D. Rodgers. *Inverse Methods for Atmospheric Sounding: Theory and Practice*. Series on Atmospheric, Oceanic and Planetary Physics. World Scientific, 2000.
- [98] R. Sanchez and NJ McCormick. "General solutions to inverse transport problems". *Journal of Mathematical Physics* 22.4 (1981), p. 847.
- [99] R. Sanchez and NJ McCormick. "Inverse Source Problems in Linear Transport Theory". *Transactions of the American Nuclear Society* 95 (2006), pp. 523–524.
- [100] R. Sanchez and N.J. McCormick. "On the uniqueness of the inverse source problem for linear particle transport theory". *Transport Theory and Statistical Physics* 37.2-4 (2008), pp. 236–263.
- [101] J. Carlos Santamarina and Dante Fratta. *Discrete signals and inverse problems*. Wiley, 2005.
- [102] Vittorio Scipolo. "Scattered neutron tomography based on a neutron transport problem". MA thesis. Texas A&M University, 2004.

- [103] *Secure hash standard*. Tech. rep. National Institute of Standards and Technology, 2002.
- [104] Claude E Shannon. "A Mathematical Theory of Communication". *Bell Systems Technical Journal* 27 (1948), pp. 379–423, 623–656.
- [105] S. Shen, N. Michael, and V. Kumar. "3D Estimation and Control for Autonomous Flight with Constrained Computation". In: *ICRA*. Submitted. Shanghai, China, 2011.
- [106] J.Kenneth Shultis and Richard E. Faw. *Radiation shielding*. American Nuclear Society, 2000.
- [107] D.S. Sivia and J. Skilling. *Data analysis: a Bayesian tutorial*. Oxford University Press, 2006.
- [108] NG Sjöstrand. "What is the average chord length?" *Annals of Nuclear Energy* 29.13 (2002), pp. 1607–1608.
- [109] VA Solé et al. "A multiplatform code for the analysis of energy-dispersive X-ray fluorescence spectra". *Spectrochimica Acta Part B: Atomic Spectroscopy* 62.1 (2007), pp. 63–68.
- [110] Matthew Sternat and Jean C. Ragusa. "Cross section inference based on PDE-constrained optimization". In: *International Conference on Mathematics, Computational Methods & Reactor Physics (M&C 2009)*. Saratoga Springs, NY: American Nuclear Society, 2009.
- [111] James Stewart. *Calculus*. Belmont, CA: Thomson Learning, Inc., 2003.
- [112] Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*. Society for Industrial & Applied Mathematics, 2005.
- [113] Barry N. Taylor and Chris E. Kuyatt. *Guidelines for evaluating and expressing the uncertainty of NIST measurements results*. Tech. rep. 1297. Gaithersburg, MD: National Institute of Standards and Technology, 1994.
- [114] R. Trainham et al. "Covariance spectroscopy applied to nuclear radiation detection". In: *Proceedings of SPIE*. Vol. 8142. 2011, 81420Q.
- [115] Jag K Tuli. "Evaluated nuclear structure data file". *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 369.2 (1996), pp. 506–510.

- [116] M.P. Unterweger et al. *Radionuclide half-life measurements*. Tech. rep. 120. Gaithersburg, MD: National Institute of Standards and Technology, 2003.
- [117] Curtis R. Vogel. *Computational methods for inverse problems*. Society for Industrial & Applied Mathematics, 2002.
- [118] Christopher G. Wahl. “Imaging, detection, and identification algorithms for position-sensitive gamma-ray detectors”. PhD thesis. University of Michigan, 2011.
- [119] Eric W. Weisstein. *Divergence Theorem*. <http://mathworld.wolfram.com/DivergenceTheorem.html>. From MathWorld—A Wolfram Web Resource. 2011.
- [120] G. Milton Wing and John D. Zahrt. *A primer on integral equations of the first kind: the problem of deconvolution and unfolding*. Society for Industrial & Applied Mathematics, 1991.
- [121] Z. Wu and M.L. Adams. “Variable Change Technique Applied in Constrained Inverse Transport Applications”. *Transactions of the American Nuclear Society* 102 (2010), pp. 209–212.
- [122] H.C. Yi, R. Sanchez, and N.J. McCormick. “Bioluminescence estimation from ocean *in situ* irradiances”. *Applied Optics* 31.6 (1992), pp. 822–830.
- [123] Feng Zhang. “Events reconstruction in 3-D position sensitive cadmium zinc telluride gamma ray spectrometers”. PhD thesis. University of Michigan, 2005.
- [124] K.P. Ziock, M. Cunningham, and L. Fabris. “Two-sided coded-aperture imaging without a detector plane”. In: *Nuclear Science Symposium Conference Record (NSS’08)*. IEEE. 2008, pp. 634–641.

APPENDICES



As shown in section 3.1.1, in numerical transport calculations, we can compute a discrete approximation to the adjoint scalar flux $\vec{\phi}^+$ with the matrix equation

$$L^+ \vec{\phi}^+ = \vec{\sigma}_d \quad ,$$

where L^+ is a matrix approximating the adjoint transport operator and $\vec{\sigma}_d$ is a vector approximating the detector response. Then in the discrete approximation, the detector response from a source \vec{q} is

$$\mathcal{R} = \langle \vec{\phi}^+, \vec{q} \rangle = (\vec{\phi}^+)^T \vec{q} \quad . \quad (\text{A.1})$$

This is true only with the vacuum boundary conditions of Equation 3.13 and Equation 3.14. If the boundary conditions are not vacuum, this expression involves the bilinear concomitant, as in Equation 3.11. However, in all problems described here, these vacuum boundary conditions are imposed.

By definition, the detector response can be computed with the forward flux,

$$\mathcal{R} = \langle \phi, \sigma_d \rangle = (\vec{\phi})^T \vec{\sigma}_d \quad . \quad (\text{A.2})$$

We desire that the response computed via the adjoint or forward flux is the same, that is

$$\vec{\phi}^T \vec{\sigma}_d = (\vec{\phi}^+)^T \vec{q} \quad . \quad (\text{A.3})$$

However, due to discretization and truncation errors, these two inner products are generally not equal. The level of agreement of these two methods is examined in this appendix.

In the next section, a simple test problem is used to compare the accuracy of the adjoint-computed responses. In the second section, we examine the cause of numerical discrepancies based on the ordering of the adjoint and discretization operations.

A.1 COMPARING THE FORWARD- AND ADJOINT-COMPUTED RESPONSES

In this section we present a comparison using a test problem of the two methods of computing the detector response. For this application, we plan to use two production-level radiation transport codes to compute the forward and adjoint fluxes. The codes that we consider are TORT⁹⁶ and Denovo.²⁷ For initial tests, we desire that the relative difference between the forward-computed and adjoint-computed detector responses,

$$\epsilon = \frac{|(\vec{\phi}^+)^T \vec{q} - \vec{\phi}^T \vec{\sigma}_d|}{\vec{\phi}^T \vec{\sigma}_d} \quad (\text{A.4})$$

is as small as possible. The response computed using the forward numerical flux should match the response computed with the adjoint numerical flux to a certain number of digits. For real-world problems, $\epsilon = 10^{-2}$ or 10^{-3} may be sufficient, but higher precision may be helpful in numerical tests.

We conducted some numerical experiments with a two-group test problem in a heterogeneous cube with one source cell and another detector cell. The configuration is described in Table A.1. A unit source in both groups was placed in $x \in [1,2]$ and $y, z \in [0,1]$, and the detector was a unit response function in both groups in $x, y, z \in [4,5]$, all dimensions in centimeters. Both forward and adjoint fluxes were computed using a variety of spatial discretizations available in TORT and Denovo.

The forward and adjoint fluxes were then used to compute ϵ in Equation A.4. Table A.2 compares ϵ for three mesh refinements and three levels of flux inner-iterations convergence criteria ($10^{-3}, 10^{-6}, 10^{-9}$). In this table, two categories of spatial discretization methods emerge. In the first category, ϵ is within a factor of 100 of the convergence criterion. In the second category, ϵ is only a weak function of the convergence criteria, often only changing from a few percent to a few tenths of a percent over the range of convergence criteria. We call the first category *numerically-adjoint* spatial discretizations. Table A.3 categorizes the spatial discretizations that we considered in our numerical experiment.

Table A.1: Description of the problem for numerical-adjoint test. The spatial domain is a 5 cm cube, with $x, y, z \in [0, 5]$. Vacuum boundary conditions are applied, thus eliminating the need for the bilinear concomitant.

Macroscopic 2-Group Cross Sections (/cm)	Spatial Mesh	Angular Quadrature
$z \in [0, 3]:$		
$\sigma_t = \begin{bmatrix} 1 \\ 1 \end{bmatrix},$	$5 \times 5 \times 5$ (coarse)	LQ_{16}
$\sigma_s = \begin{bmatrix} 1/2 & 0 \\ 1/5 & 3/4 \end{bmatrix}$	$10 \times 10 \times 10$ (medium)	[72,
	$20 \times 20 \times 20$ (fine)	S4-2]
$z \in (3, 5]:$ void		

A.2 COMMUTATIVITY OF ADJOINT AND DISCRETE OPERATIONS

Although we can gain some insight through the results in the preceding section, it is also instructive to consider the methods that radiation transport codes use in discretizing the transport equation and applying the adjoint operation. Indeed, the order of applying the discretization and adjoint to the S_N equations is important. The discretized adjoint equation is different in general than the adjoint of the discrete forward system. As we will see, the common practice in S_N codes (at least TORT and Denovo) is to supply the solution of the discretized adjoint equation. The identity of Equation A.3 will be true (to the floating-point convergence tolerance) if the adjoint-of-the-discretized equation is used. However, if the discretized adjoint solution is used, then Equation A.3 will only be in close agreement when the truncation error in the computed fluxes are small and the inner iterations are fully converged. To make this clear, we develop a very simple discretized transport equation in the following subsections.

A.2.1 Methods

Before moving to the adjoint equation, we briefly state the relevant forward S_N equation and its discretized form. Next, the continuous adjoint equation is discretized. Finally we derive the adjoint of the discrete forward system. First, the energy and

Table A.2: Comparison of relative inner product difference ϵ of Equation A.4 for various TORT and Denovo spatial discretizations and flux convergence criteria.

	No. of cells	Relative Inner Product Error		
		Convergence Criterion		
		10^{-3}	10^{-6}	10^{-9}
TORT	5^3	0.042	0.042	0.042
Optimum	10^3	0.033	0.033	0.033
theta-weighted	20^3	0.028	0.028	0.028
	5^3	0.11	0.11	0.11
Optimum	10^3	0.087	0.087	0.087
xyz nodal	20^3	0.023	0.023	0.023
	5^3	0.0273	0.0273	0.0273
Characteristic	10^3	0.0047	0.0049	0.0049
	20^3	0.0029	0.0024	0.0024
Denovo	5^3	0.14	8.4×10^{-5}	8.6×10^{-8}
Linear	10^3	0.019	4.3×10^{-5}	6.2×10^{-8}
discontinuous	20^3	0.0014	2.0×10^{-5}	1.7×10^{-8}
	5^3	0.078	2.5×10^{-5}	1.2×10^{-7}
Tri-linear	10^3	0.017	2.4×10^{-5}	0.0
discontinuous	20^3	0.0025	1.1×10^{-5}	4.2×10^{-9}
	5^3	0.0802	0.00013	2.5×10^{-8}
Step	10^3	0.020	2.8×10^{-5}	5.7×10^{-8}
characteristic	20^3	0.0023	1.8×10^{-5}	8.5×10^{-9}
	5^3	0.021	1.3×10^{-5}	2.5×10^{-8}
Weighted	10^3	0.022	6.6×10^{-5}	1.1×10^{-7}
diamond-difference	20^3	0.0014	2×10^{-5}	1.7×10^{-8}
	5^3	0.056	0.025	0.025
Theta-weighted	10^3	0.026	0.042	0.042
diamond-difference	20^3	0.0074	0.0061	0.0061
	5^3	0.17	0.23	0.23
Weighted	10^3	0.38	0.41	0.41
Diamond-difference with flux fixup	20^3	0.16	0.16	0.16

Table A.3: Numerically-adjoint spatial discretizations in TORT and Denovo in the test.

Numerically adjoint	Denovo weighted diamond difference Denovo linear discontinuous Denovo tri-linear discontinuous Denovo step characteristic
Not numerically adjoint	TORT mode 0 – optimum theta-weighted TORT mode 2 – optimum xyz nodal TORT mode 7 – characteristic Denovo weighted diamond difference, flux fixup Denovo theta-weighted diamond difference

angle variables are discussed.

Energy In the energy variable, the only change required for the adjoint is the reversal of E and E' in the scattering kernel. Whether this is done before or after the multigroup approximation is applied, the result is that the scattering matrix is transposed. Thus, for our current purpose, it is enough to consider the one-speed equation.

Angle In the discrete ordinates method, there are a set of M directions along which we solve the transport equation. Since the results of these equations are combined in a linear sum, we only consider one of the equations, corresponding to direction m .

Space To be concise, we treat here only one spatial variable x . Treating the full three spatial dimensions would add little to the final conclusions. The one-speed S_N equation for ordinate m is

$$\mu_m \frac{d\psi_m}{dx} + \sigma_t \psi_m(x) = q_m(x) \quad x \in [a, b] \quad ,$$

where q_m includes the scattering and fixed sources. From here on, the m index will be suppressed. We consider only vacuum boundary conditions $\psi(a) = 0$ for $\mu > 0$.

Discrete Forward Equation Integrating over a homogeneous cell i , the equation is

$$\frac{\mu}{\Delta_i} (\psi_{i+1/2} - \psi_{i-1/2}) + \sigma_{t,i} \psi_i = q_i ,$$

where Δ_i is the width of cell i , $\psi_{i\pm 1/2}$ are the right and left edge fluxes, respectively, and ψ_i is the cell-average flux. If we take as the closure for the cell-average flux a weighted sum

$$\psi_i = \frac{1 + \alpha_i}{2} \psi_{i+1/2} + \frac{1 - \alpha_i}{2} \psi_{i-1/2} ,$$

the forward discrete equation for one cell is

$$\frac{\mu}{\Delta_i} (\psi_{i+1/2} - \psi_{i-1/2}) + \frac{\sigma_{t,i}}{2} \left(\frac{1 + \alpha_i}{2} \psi_{i+1/2} + \frac{1 - \alpha_i}{2} \psi_{i-1/2} \right) = q_i . \quad (\text{A.5})$$

Discrete Adjoint Equation If the continuous one-speed S_N adjoint equation is

$$-\mu \frac{d\psi^+}{dx} + \sigma_t \psi^+(x) = q^+(x) ,$$

then analogously the discretized adjoint equation is

$$-\frac{\mu}{\Delta_i} (\psi_{i+1/2}^+ - \psi_{i-1/2}^+) + \frac{\sigma_{t,i}}{2} \left(\frac{1 + \alpha_i}{2} \psi_{i+1/2}^+ + \frac{1 - \alpha_i}{2} \psi_{i-1/2}^+ \right) = q_i^+ . \quad (\text{A.6})$$

Discrete Forward System If we divide the domain $[a, b]$ into I intervals, then we have I unknowns. Writing the system of I equations in terms of the cell-edge fluxes, we have

$$\begin{bmatrix} a_1 & & & & & & \\ b_1 & a_1 & & & & & \\ & \ddots & \ddots & & & & \\ & & & b_i & a_i & & \\ & & & & \ddots & \ddots & \\ & & & & & b_I & a_I \end{bmatrix} \begin{bmatrix} \psi_{1-1/2} \\ \psi_{1+1/2} \\ \psi_{2+1/2} \\ \vdots \\ \psi_{i+1/2} \\ \vdots \\ \psi_{I+1/2} \end{bmatrix} = \begin{bmatrix} 0 \\ q_1 \\ q_2 \\ \vdots \\ q_i \\ \vdots \\ q_I \end{bmatrix} . \quad (\text{A.7})$$




Figure A.1: Demonstrating the non-uniform mesh for the test problem. This is a cut of the x - y plane.

Equation A.9 would be equivalent. However, in the general case they are not.

Even though this is shown for 1D, the principle still applies to 3D. As long as the downwind coupling of two adjacent cells occurs, the system will be a block lower-triangular matrix. When transposed, the same issues will occur.

A.2.2 Numerical results

In general, Denovo solves for the adjoint by flipping the sign of each ordinate and transposing the scattering matrix. It does not further modify (with one exception) the spatially discretized equation (the 3D analog of Equation A.5). This procedure solves the discrete form of the continuous adjoint S_N equations. Thus, it should be possible to see a breakdown in the equivalence of the inner products $\langle \phi, \sigma_d \rangle$ and $\langle \phi^+, q \rangle$, at least in certain circumstances. If the forward discrete solution is a good approximation to the continuous forward solution, then the discretized adjoint solution should approximate the solution of the adjoint of the discrete system. However, when the discrete forward solution is inaccurate, Denovo's adjoint may be less trustworthy to reproduce $\langle \phi, \sigma_d \rangle = \langle \phi^+, q \rangle$.

We built a one-group problem to demonstrate this phenomenon. First, the mesh is non-uniform so that the adjoint-of-the-discrete is not the same as the discrete-of-the-adjoint. The Cartesian mesh is $18 \times 1 \times 1$, and the dimensions are 10 cm in x and 1 cm in y and z . The first and last cells are 1 cm in x . The middle 16 cells are periodic cells with period 2. The first of the two cells is 0.9 cm thick, and the second of the two is 0.1 cm thick (in the x direction). See Figure A.1 for an illustration. The boundary conditions are vacuum.

The cells are all of the same material, with total scattering cross section σ_t and scattering ratio $c = 1/4$. We vary the total macroscopic cross section σ_t to regions for which the transport solver will produce results with significant truncation error. This assumes that as the cells become increasingly optically thick, the local truncation error




Figure A.2: Accuracy of adjoint as discrete solution degrades.

of the solution will increase.

A unit source is in the far left cell, and a detector with unit response is in the far right cell. The measure of the adjoint error is

$$\epsilon = \left| \frac{\langle \vec{\phi}, \vec{\sigma}_d \rangle - \langle \vec{\phi}^+, \vec{q} \rangle}{\langle \vec{\phi}, \vec{\sigma}_d \rangle} \right| .$$

The iterative convergence tolerance is $\tau = 10^{-10}$, so $\epsilon \approx 10^{-10}$ is about as good as could be expected.

The relative error ϵ for the spatial discretizations available in Denovo is shown in Figure A.2. The behavior of the results fall into three groups:

Adjoint of discrete The weighted diamond difference (WDD) results produced $\epsilon \approx \tau$ for all σ_t . This is despite the extremely poor solutions (with large truncation error) at high σ_t , where the detector response was overestimated by much more than a factor of 10^{10} . This indicates that the adjoint implemented for WDD is indeed the adjoint of the discrete forward system, not the discretized adjoint equation. Both Denovo's source code and documentation verify that the WDD implementation uses different equations for the forward and adjoint spatial equations, in contrast to the other spatial discretizations.

Linear discretizations These methods (step characteristics, linear discontinuous, and

trilinear discontinuous) have $\epsilon \approx \tau$ at low σ_t , but at higher optical thicknesses $\epsilon \gg \tau$. This is the expected behavior based on the analysis in Appendix A.2.1.

Non-linear discretizations These methods (weighted diamond difference with negative flux fixup and theta weighted diamond difference) do show an improving ϵ as the cell optical thickness decreases. However, this error never reaches the iterative convergence tolerance τ . The non-linearities in these methods are likely to blame for the higher errors as compared to the linear methods' results.

Conclusions From the above analysis and numerical results, we conclude that Denovo solves the discretized adjoint equations, not the adjoint of the discretized forward S_N system. The lone exception to this is the WDD implementation, which solves the adjoint of the discrete forward system.

If the discrete solution of the forward and adjoint equations is accurate (that is, their local truncation error is small), then it is likely that ϵ will be small, at least for the linear discretizations. If the discrete solution is inaccurate, then we should not expect that ϵ be small, unless we use WDD.

As mentioned in section 4.2.3, one can include many kinds of uncertainties when using Bayesian inference. One source of uncertainty in the predictive model is physical properties; in the present case, this would include photon cross sections. In the end, we chose to exclude this data from the adaptable model parameters because the cross section uncertainties were much smaller than the source uncertainties. This appendix explains the reasoning behind the choice to exclude the cross sections from $\vec{\alpha}$.

For the deterministic transport calculations, we use the fine-group shielding cross sections from `SCALE`.¹¹ These, in turn, are generated from the Evaluated Nuclear Data File version seven (`ENDF/B-VII.0`).¹⁷ While this compilation of nuclear data includes estimated covariance matrices for neutron data (that is, when the incident particle is a neutron), there are no included uncertainties for the photon cross sections.

Some argue that since photon cross sections are theoretical, their uncertainties are negligible. While the uncertainties in neutron data can be much larger, photon cross sections can still have significant error. In the words of Hubbell, “theory is an interpolation of experiment for purposes of computing and compiling” photon cross sections.⁴⁹ To estimate the uncertainty in the cross section data, one would thus compute the spread of high quality experimental results about the evaluated data. There have been graphical comparisons of the evaluated data against measurements,^{47,50} but there is no database of experimental results to automate such a large task. Hubbell compiled a bibliography with a list of papers that contain total photon cross section measurements,⁴⁸ but one would be forced to copy the 22 000 data points from the 580 references.

Without going to these great lengths, there are limited estimates of the photon cross section error available. The `ENDF/B-VII.0` photon data is derived from the Evaluated Photon Data Library published in 1997 (`EPDL97`).²⁰ The `EPDL97` reference

Table B.1: Estimated uncertainties in photoelectric cross section in EPDL97 (assumed to be 1σ).

Energy range (keV)	Relative uncertainties (%)	
	Solid	Gas
0.5–1.0	10–20	5
1.0–5.0	5	5
5–100	2	2
100–10000	1–2	1–2

contains estimated uncertainties for the photoelectric effect cross section, partially copied in Table B.1. As Hubbell notes,⁴⁶

Since photoionization or atomic photoeffect is the dominant interaction at low photon energies, where the uncertainties are the largest, these very approximate and subjective percent uncertainties can be taken as a rough guide to the uncertainties of the total photon cross section σ_{tot} . . .

Since the bottom energy of the SCALE fine-group library is 10 keV, we could assume an uncertainty of 2% over the energy range for this problem (roughly 10 keV to 2 MeV). This small uncertainty is much less than the uncertainty on the source parameters.

SOURCE MAPPING EXPERIMENT



We conducted a series of radiation detector measurements to test the algorithms proposed in this work. A number of small radioactive gamma point sources were placed in different locations in a mostly empty room. Then we took multichannel counts using a sodium iodide (NaI) detector, moving the detector to different locations throughout the room. The goal of the experiment is to infer the locations and spectra of the sources based on the detector readings and the knowledge of the layout of the room, but without using the knowledge about the location or spectrum of the sources.

The location of the experiment was Burlington Engineering Labs (BEL) at NC State University, room 2144. The counts were performed from August 30, 2011 to September 15, 2011, a span of 17 days. This chapter is an account of the procedures, equipment, and data of the experiment. Section 5.4.3 demonstrates the use of this data to solve the IP using our new approach.

C.1 DETECTOR STATISTICAL ERROR

Before describing the experiment, we first review the concept of statistical error in radiation detection. The probability of detection in a particular energy channel can typically be described by the Poisson distribution [63, §3.IIB]. The Poisson distribution gives the probability of a discrete number of occurrences of a specified random event within a given time interval, assuming that each independent event has a constant and small probability. The Poisson distribution is a one-parameter distribution, with that one parameter being the expected value or mean \bar{x} . The probability distribution function for the Poisson distribution is

$$P(x) = \frac{\bar{x}^x}{x!} e^{-\bar{x}} ,$$

where $x \geq 0$ is the integer number of occurrences and $P(x)$ is the probability of observing x occurrences.

As Knoll observes, the Poisson distribution is normalized and its mean is \bar{x} :

$$\sum_{x=0}^n P(x) = 1 \quad , \quad \sum_{x=0}^n xP(x) = \bar{x} \quad ,$$

Furthermore, the variance is⁶³

$$\sigma^2 = \sum_{x=0}^n (x - \bar{x})^2 P(x) \quad ,$$

which results in the standard deviation being equal to the square root of the expected value, $\sigma = \sqrt{\bar{x}}$.

If the mean value of the Poisson distribution is not small, $\bar{x} \gtrsim 20$, then the distribution approaches a normal distribution,

$$P(x) = \frac{1}{\sqrt{2\pi\bar{x}}} \exp\left(-\frac{(x - \bar{x})^2}{2\bar{x}}\right) \quad ,$$

which still has mean \bar{x} and standard deviation $\sqrt{\bar{x}}$.

In the case of radiation detection, the Poisson and Gaussian model is valid for counts in an energy channel. With only one count, the best estimate for the mean \bar{x} is the observed counts x . Without further knowledge, we assume $\bar{x} \approx x$ and thus the estimated standard deviation of the count is $\sigma_x = \sqrt{\bar{x}}$. This rule applies for measurements that are counts, not counting rates, sums or differences of counts, averages of counts, or other derived values.⁶³

Linear propagation of error While the simple counts are assumed to be described by Poisson statistics, the quantities derived from these counts are not, and so the errors must be propagated through the various operations. As is the convention, the propagation of small errors through arbitrary operations can be well described by a linear approximation.⁶⁶ Using a first-order Taylor series approximation of the derived quantities \vec{r} from independent variables $\vec{\alpha}$, the covariance matrix C_r for \vec{r} is (for example, see §2.2 in [58])

$$C_r = S C_\alpha S^T \quad , \quad (\text{C.1})$$

where C_α is the covariance matrix of $\vec{\alpha}$ and S is the matrix of sensitivities. This is the same equation as Equation 2.43. Since a first-order linear expansion is used to derive this, it is only accurate for small errors and for functions which can be well approximated as linear in the local neighborhood.

Assuming that the quantities in vector $\vec{\alpha}$ follow a Poisson distribution, the covariance matrix C_α is known. Then the sensitivity coefficients in the matrix S are the only remaining unknown to compute C_r . These coefficients can be computed in several ways. In the main portion of this work, adjoint solutions are used to compute the sensitivity. For simple expressions, analytic forms of the derivatives can be found. In the propagation of uncertainty of these measurements, we used a similar analytic approach, but one that is computed automatically by the application of the chain rule. This technique is known as automatic differentiation.³⁷ We employed the Python uncertainties* package in performing these calculations.

C.2 EQUIPMENT

This section describes the sodium iodide multichannel detector system used in this experiment. In addition, it describes the radioactive sources and the configuration of the laboratory.

C.2.1 Detector system

The primary instrument for this experiment was the sodium iodide detector system. This system is property of the Department of Nuclear Engineering at NC State University. The detector is normally used for laboratory classes.

The detector itself is a sodium iodide (NaI) crystal coupled to a photomultiplier tube. The crystal has a diameter of 2 inches and a height of 2 inches, and so its shape is commonly referred to as 2" × 2". This is connected to a preamplifier and high-voltage supply. The detector signal is processed by a multichannel analyzer with 2048 channels, and the data is acquired via a USB-connected laptop. Table C.1 lists the individual components, and Figure C.1 describes how they were connected. Long cables to the detector, 10 feet or more, allowed the electronics to remain in one place

*Eric O. Lebigot, *Uncertainties: a Python package for calculations with uncertainties*.

throughout the experiment, with only the detector crystal and preamplifier being moved.

When we first assembled the detector system, we adjusted the amplifier gain so that the multichannel analyzer (MCA) captured approximately 0 to 3 MeV. Our reference for this adjustment was a 2615 keV gamma peak present in the background radiation, emitted by a daughter of ^{232}Th . Adjusting this peak to channel 1700 made the last channel, number 2048, approximately correspond to 3000 keV.

The Easy-MCA-2k MCA has a USB-interface. Ortec's Maestro, MCA acquisition software for Microsoft Windows[®], communicates with the MCA, enabling one to visualize and save the counted spectrum.⁷⁵ We used the ASCII Spe format to export each counted spectrum in a platform-independent file. These files are included in the electronic archive attached to this thesis.

One important distinction when using the Ortec hardware and Maestro detection software is the difference between live time and real time. Ortec has developed the Zero Dead Time (ZDT) mode to correct for dead time losses. Dead time is that time during which a detector system is busy processing an input pulse. Any additional pulses from further detected particles will be missed by the detector counting system during this dead time.⁶³ Thus, the detector reading will be an underestimate of its true value as the dead time increases. The ZDT is a correction of the real time to account for this underestimate. With an estimate of the real time t_r and the fractional dead time f_d , the live time t_l is [75, §4.2.7.6]

$$t_l = t_r(1 - f_d) \quad .$$

If we use the live time to compute the count rate, this yields the dead-time corrected value. The standard deviation in this count rate is \sqrt{N}/t_l .

c.2.2 *Radioactive sources*

The radioactive sources used in the experiment were low-activity gamma sources owned by the Department of Nuclear Engineering at NC State University. The two types of sources were ^{137}Cs and ^{60}Co . The important decay properties of these two radionuclides are listed in Table c.2. ^{137}Cs produces one strong gamma line at 662 keV and ^{60}Co produces two lines at 1173 and 1332 keV.

Table c.1: Detector components used in experiment. All components are branded Ortec except the NaI crystal.

Component	Model	Serial Number
NIM bin	4001C rev. 20	5347
High voltage power supply (HVSP)	556 rev. p	08280653
Amplifier (Amp)	572A rev. b	10238309
Photomultiplier and preamplifier (Preamp)	276	296
Multichannel analyzer (MCA)	Easy-MCA-2k rev. D	10131831
2"×2" NaI detector	905-3	
NaI crystal by Saint-Gobain	2M2/2	60004-00410-1




Figure c.1: A schematic of the detector components.

Table c.2: ^{137}Cs and ^{60}Co radioisotope properties. Gamma energies and branching fractions are from NNDc.¹⁵ Half-lives are from NIST.¹¹⁶

Isotope	Gamma Emissions		Half-Life	Loss in
	Energies (keV)	Intensities (%)	$t_{1/2}$ (days)	30 days (%) ^a
^{137}Cs	661.657 ± 0.003	85.10 ± 0.20	$11\,018.3 \pm 9.5$	0.19
^{60}Co	1173.228 ± 0.003	99.85 ± 0.03	1925.20 ± 0.25	1.1
	1332.492 ± 0.004	99.9826 ± 0.0006		

^a The relative amount of atoms which decay in $\Delta t = 30$ days, calculated as $1 - \exp(-\lambda \Delta t)$, with $\lambda = \log_e(2)/t_{1/2}$.

Each source was encapsulated in a small plastic disc with diameter 1 inch and thickness about $\frac{1}{4}$ inch. In Isotope Products' (a division of Eckert & Ziegler) classification, these are Type D sources [25, p. 46]. The active diameter (in which source material is present) in this case is about 5 mm. Where necessary, we assume that the source is deposited immediately below the top of the source disk. The source disks are made of plastic, which we assume is similar to the composition of polyethylene terephthalate [3, §27], a widely-used plastic with chemical form $(\text{C}_{10}\text{H}_8\text{O}_4)_n$. Using the measured values of the mass and volume of the discs, the average density of the plastic in this experiment was 1.12 g/cm^3 .

Several sources were used to increase the total activity and thus the detector count rate. Each source was nominally $1 \mu\text{Ci}$, as shown in Figure c.2 and tabulated in Table c.3. Unfortunately, only three of the sources were dated.

To determine the current activity of each source, we counted the sources one by one using an identical detector-source configuration. The source was placed closely to the detector to improve the absolute efficiency of the counting. By summing the counts in the photopeaks above the continuum, we computed the source's total uncollided contribution in the detector for that energy. The total counts N_i in a photopeak at energy E_i from a source with activity A is

$$N_i = A \epsilon_{\text{abs,peak}}(E_i) \gamma_i \quad . \quad (\text{C.2})$$

where γ_i is the average number of gamma particles emitted at energy E_i for one decay

Table c.3: The ^{137}Cs and ^{60}Co sources used in this experiment, with their specified calibration date and activity. Most of the sources are undated.

Id	Nuclide	Activity (μCi)	Calibration Date
11	^{137}Cs	1.01 ± 0.10^a	31 January 1983
12	^{137}Cs	1.00 ± 0.09	–
14	^{137}Cs	1.00 ± 0.09	–
15	^{137}Cs	1.00 ± 0.09	–
16	^{137}Cs	1.00 ± 0.09	–
21	^{60}Co	0.99 ± 0.10^a	6 March 1974
22	^{60}Co	1.00 ± 0.09	–
23	^{60}Co	1.00 ± 0.09	–
24	^{60}Co	1.00 ± 0.09	15 ^b March 2005
25	^{60}Co	1.00 ± 0.09	–

^a Professor Robin Gardner on November 10, 2011 estimated the uncertainties of the department-made sources as five to ten percent. The more pessimistic estimate is used here. The other uncertainties in this table are an estimate from a radioisotope supplier. We assume that these sources have similar uncertainty. Quoting from the supplier's catalog: "Sources are manufactured with contained activity values of $\pm 15\%$ of the requested activity value unless otherwise noted in the catalog." [25, p. 5] Assuming a uniform distribution of the true value in this interval, the estimated standard deviation or *standard uncertainty* is $15\%/\sqrt{3} \approx 9\%$ [113, §4.6].

^b The source is labeled March 2005 but does not specify a day. We assume for calculations that the day is the middle of the month.



(a) ^{137}Cs sources.



(b) ^{60}Co sources.

Figure c.2: The six ^{137}Cs sources and five ^{60}Co disc sources. Each source is labeled with the source id used throughout this experiment.

of the source, and $\epsilon_{\text{abs,peak}}$ is the absolute peak detector efficiency.

In this case, we have multiple sources that emit the same peak or peaks, but their activities are presumably different. For three of the sources, the activity A is known. We could look up γ_i , but it will be the same for peaks of the same energy. Likewise, we could estimate the absolute efficiency. However, neither of these steps are necessary, because these factors are the same for each source, assuming the source-detector configuration is identical. Instead, we apply Equation c.2 to one peak area counted from two different sources, one dated and the other undated. (The matter of calibrating the measured spectra and computing the peak areas will be deferred until a later section.) Taking the ratio of these two equations, we obtain

$$A^u = A^d \frac{N_i^u}{N_i^d} .$$

The u and d superscripts denote dated and undated sources, respectively. A^x is the activity of the radionuclide source x and N_i^x is the number of counts in photopeak i from source x . The $\epsilon_{\text{abs,peak}}(E_i)$ and γ_i factors cancel because they are the same for the same peak. For the ^{137}Cs source, one estimate of A^u is available, and this is the one we use. However, for the ^{60}Co source, there are two peaks and so Appendix c.2.2

can be applied twice. In addition, there are two dated ^{60}Co sources, so there are four activity estimates for each undated ^{60}Co source, two each for the two peaks and for the two dated sources. We have combined these with a weighted least-squares average.[†] Table c.4 gives the result of these calculations. The sources that have similar labels and were likely made at the same time ($\{12, 13, 23\}$, $\{14, 15, 16\}$, $\{22, 25\}$, as seen in Figure c.2) have computed $1\text{-}\mu\text{Ci}$ dates within a few years, lending some additional credibility to the computed activities.

We assume that each of the sources has constant activity over the 17 days of the measurements. This introduces only negligible errors, since the half-lives of the source nuclides are several years. As shown in the last column of Table c.2, both of the sources decay less than 1% in 17 days, a negligible change compared with the manufacturing uncertainty of 10%.

c.2.3 The room configuration

The experiment was conducted in laboratory room 2144 BEL. The room was mostly empty, with the exception of a few cabinets and counters. The walls are concrete block. A Cartesian coordinate system has its origin at the bottom northwest corner of the room. All components of the geometric model are right rectangular prisms, or bodies formed by the combination of these prisms. Figure c.3 shows the dimensions of the room and the major components. The room is pictured in Figure c.4. A complete description of the layout is available in the MCNP input file electronically attached to

[†]Assume a linear model $S\vec{\alpha} = \vec{r}$, where S is known. The vector \vec{r} is observed or measured as \vec{r}_m , with some level of certainty, described by the covariance matrix C_m . The purpose is to compute the optimal $\vec{\alpha}^*$, with corresponding covariance matrix C_{α}^* . In the case of using two estimated activities to obtain a better value for the true activity, \vec{r}_m is the two activity estimates. C_m is a diagonal matrix with the variances of the estimates on the diagonal. The unknown $\vec{\alpha}$ is a scalar in this case, and the matrix S is $[1, 1]^T$.

With a prior uncertainty on $\vec{\alpha}$, we could use Bayes' Theorem to obtain a posterior distribution on $\vec{\alpha}$. However, with no prior on $\vec{\alpha}$, the next best option is the maximum likelihood estimator [112, Example 3.4]:

$$\vec{\alpha}^* = C_{\alpha}^* S^T C_m^{-1} \vec{r}_m \quad , \quad C_{\alpha}^* = (S^T C_m^{-1} S)^{-1} \quad .$$

This could also be derived by letting $C_{\alpha} \rightarrow \infty$ in the Bayesian formulation. Since C_m is diagonal, these equations can be simplified to

$$\alpha^* = (\sigma_{\alpha}^*)^2 \sum_i \frac{\vec{r}_{m,i}}{\vec{\sigma}_{r_{m,i}}^2} \quad , \quad (\sigma_{\alpha}^*)^2 = \left(\sum_i \frac{1}{\vec{\sigma}_{r_{m,i}}^2} \right)^{-1} \quad .$$

This is commonly referred to as *weighted least squares*.⁷

Table C.4: Estimated sources activities based on computed activities of dated sources 11, 21, and 24 and on the measured peak area ratios. For each undated ^{137}Cs source, there is one peak and one dated ^{137}Cs reference source, so we have only one activity estimate. For each undated ^{60}Co source, there are two dated ^{60}Co reference sources and two peaks, yielding four activity estimates. The reported activity (and the derived 1 μCi date) are based on a weighted average of the four estimates.

Id	September 2011 Activity (μCi)	Estimated Date of 1 μCi Activity
11	0.52 \pm 0.05 ^a	
12	0.69 \pm 0.07	3 July 1995
13	0.67 \pm 0.07	6 July 1994
14	0.53 \pm 0.05	22 August 1983
15	0.52 \pm 0.05	14 November 1982
16	0.49 \pm 0.05	20 April 1980
21	0.0071 \pm 0.0007 ^a	
22	0.030 \pm 0.002	15 January 1985
23	0.078 \pm 0.005	13 April 1992
24	0.43 \pm 0.04 ^a	
25	0.052 \pm 0.003	9 March 1989

^a The activities for these sources were computed from the activities and dates found on their labels.

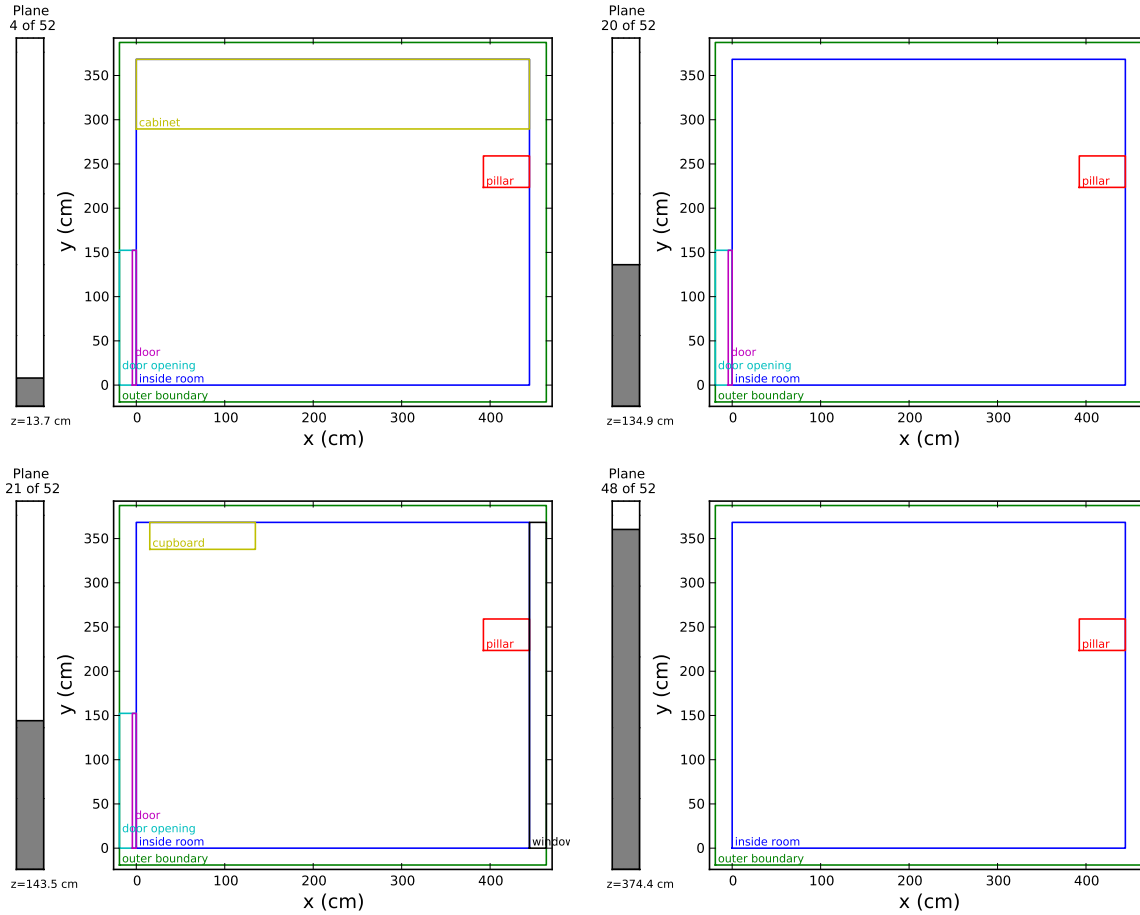


Figure c.3: Slices from the x - y plane of the 2144 BEL room layout.

this thesis.

For simplicity, we limit the number of distinct materials present in the room. The materials and their composition are listed in Table c.5. For finer detail, see the MCNP input file.


C.3 PROCEDURE

There were two sets of measurements during this experiment. The first set was necessary to determine the activities of the undated sources. The second was a series of counts from various detector locations to map the radioactive source.

Count times are shown in Figure c.5. The majority of counts were less than 2 hours.



(a) The windows on the south ($+x$) wall.



(b) The doors in the north ($-x$) wall.



(c) The counter and cabinets on the east ($+y$) wall.

Figure c.4: Pictures of the 2144 room.

Table C.5: The elemental composition of materials used in the 2144 model. The densities and compositions are from the Standard Composition manual of SCALE.¹¹ Integer values of composition are the atomic abundance per molecule, while real values indicate atomic weight percent.

Material	Density (g/cm ³)	Composition	Components
Void	0.0	-	Air in room
Oak Ridge concrete	2.2994	17.52 – carbon 3.448 – silicon 3.265 – magnesium 1.083 – aluminum 0.7784 – iron 0.6187 – hydrogen 0.1138 – potassium 0.0271 – sodium	Walls, floor and ceiling
Oak wood	0.7	10 – hydrogen 6 – carbon 5 – oxygen	Door
Carbon steel	7.8212	99. – iron 1. – carbon	Cabinets and cupboards
PVC	1.6	3 - hydrogen 2 – carbon 1 – chlorine	Ceiling tiles
Lead	11.344	100. – lead	Lead shields




Figure C.5: The count times for two counting sets. Each histogram bin is one hour wide.

The longer counts were typically for a night or weekend.

C.3.1 Source strength counts

Since many of the sources were undated, it was necessary to take a series of counts to estimate the activity of each source. As discussed in Appendix C.2.2, the activity of a particular radioactive source can be computed by comparing the ratio of the counted peak area with the peak area (at the same energy) counted from a reference source. In order to cancel the influence of detector efficiencies in this calculation, the source and detector configuration must be identical for the unknown and reference source. We used an identical configuration for all ^{137}Cs sources and another identical configuration for all ^{60}Co sources. One of these configurations is shown in Figure C.6. As the picture shows, the source-to-detector distance is relatively small to decrease the counting time needed to obtain good statistics.

For the ^{137}Cs sources, seven separate counts were taken – six for the six ^{137}Cs sources and one background count. For the ^{60}Co sources, five counts were needed for the five sources and one for background.




Figure c.6: The source and detector configuration for the source-strength counts.

c.3.2 Locating the sources counts

The second set of counts were to gather the detector data necessary to map the radiation sources. Counts were performed for two source configurations, in addition to background, for which all the sources were removed from the room. The two source configurations were

- ^{137}Cs only, with all six sources in a small box on the floor in the southwest corner of the room.
- ^{137}Cs and ^{60}Co simultaneously. The ^{137}Cs was located in the same place. The ^{60}Co sources were arranged in a line along the chalk tray, spaced roughly 10 cm apart. A line was chosen to provide a non-point distribution for the mapping algorithm.

The coordinates of the sources are given in Table c.6. A graphical presentation of the source locations and intensities are plotted in Figure 5.1. Pictures of the source locations are provided in Figure c.7. In retrospect, the first configuration could have been eliminated, since the peak energies of ^{137}Cs and ^{60}Co are sufficiently distant, hence unlikely to overlap peaks significantly.

While there were three source configurations (2 with sources and one background), the main change in this set compared to the source strength estimation counts was the movement of the detector. There were six locations, which are listed in Table c.7. We chose these locations to get reasonable spatial coverage of the room in all three



(a) ^{137}Cs source location in the southwest corner. The detector is at location 5.



(b) ^{60}Co source locations. The blackboard is on the $y = 0$ plane and the wood door is on the $x = 0$ plane. The detector is at location 6.

Figure c.7: The ^{137}Cs and ^{60}Co source locations.

Table c.6: Source locations for ^{60}Co . All dimensions are in cm. The y location of all sources was 9.5 cm, and 90.2 cm for z . All the ^{137}Cs sources (IDs 11-16) were located at (440, 5, 1) cm.

Nuclide	Source	x
^{60}Co	21	96.5
^{60}Co	22	108.0
^{60}Co	23	119.4
^{60}Co	24	130.2
^{60}Co	25	143.8

Table c.7: Detector locations. All dimensions are in cm.

Index	x	y	z
1	59.7	194.3	54.6
2	242.6	72.5	10.2
3	364.5	180.3	84.5
4	330.2	315.0	177.8
5	425.5	102.9	30.5
6	181.6	42.5	82.6

spatial directions. Locations five and six were added to be closer to the sources to obtain improved counting statistics.

At each location, a detector count with no collimation was taken for each source configuration. In addition, six collimated 2π counts were taken at a location: $\pm x$, $\pm y$, and $\pm z$. The collimation was accomplished as shown in Figure c.8. Lead blocks were used to shield gammas coming from the side or back of the detector. Custom machined blocks fit snugly around the detector crystal.

With six detector locations, seven directions at each location, and three source configurations, there are 126 possible combinations. However, many of these measurements were not taken due to time limitations. For the source mapping measurements, 73 counts were taken, but two of these were duplicates. All the possible combinations



(a) Side view.



(b) Front view.

Figure c.8: The 2π collimated detector. The lead bricks are $2 \times 4 \times 8$ inches in dimension. The blocks with a hollowed cylinder have an outside side-length of 4 inches and a height of 2 inches.

	Location					
	1	2	3	4	5	6
full	4c154b029	4b8185999	d6eba9291	e2b106ff8	838 b84	b1b12c912
+x	25de83e28	de8c8f515	7fb605d46			
-x	f37057ee5	228baa1a5	60875e5b2			
Direction +y	8a32724af	f01a591b1	d0e97f376			
-y	20a6e45ef	1468b862c	b7c fb7 ff5			
+z	c605a92a5	be03d1998	5988c556a			
-z	f6298d150	b0d6f3865	f6418327c			

Figure c.9: The IDs of the detector measurements that were taken. Blank squares indicate that no count was taken for that configuration. Each row is a collimation direction; each column is a location. The three contiguous blocks for a particular location and direction are the three source configurations: (1) ^{137}Cs only, (2) ^{137}Cs and ^{60}Co , and (3) background (with left-to-right ordering:

1	2	3
---	---	---

). The three characters are the first three digits of the ID hash of the measurement.

were conducted for locations 1, 2, and 3. For locations 4 through 6, only the isotropic combinations were counted. Figure c.9 lists the combinations that we counted.

C.4 PROCESSING THE DATA

We describe here the steps taken to preserve, categorize, and filter the detector data. Much of the data is available in the electronic archive attached to this thesis.

c.4.1 *Compiling raw database*

Using the Ortec Maestro data acquisition software, text files were exported with:

- the total counts per channel,
- the real and live length of count time, and
- the date and time of the count.

To each file, we manually added

- the detector location (1–6),
- direction of collimation or lack of collimation (4π or any of the six 2π views), and
- source configuration (^{137}Cs only, ^{137}Cs and ^{60}Co , or background)..

These files constitute the raw data. This data was saved in an HDF5[‡] file with appropriate metadata using PyTables.[§] A 40-character hexadecimal identification string was produced for each count using a SHA-1¹⁰³ hash of the count’s date and time stamp and the live count time.

c.4.2 *Compiling processed database*

After saving the spectra in the database of raw data, there were a few processing steps that were performed before the data could be used in the source mapping. These steps were:

1. peak identification,
2. energy calibration,
3. rebinning spectra to a standard energy structure,
4. background subtraction, and
5. determination of peak areas.

These tasks are outlined in the following paragraphs. Through each of these steps, the standard error of each of the derived quantities must be computed.

Peak identification Although there are many peak identification algorithms, we manually selected peaks in the spectrum by visualizing them in the graphical program PyMCA.¹⁰⁹ The channel of the peak center was recorded for those peaks from the background and source spectra. Depending on the relative strength of the source versus background, as well as the overlapping of close peaks, not all the peaks were identifiable in every spectra. The observed peak channels were then used to calibrate the detector. The four background peaks identified in most spectra are listed in Table c.8.

[‡]Hierarchical data format version 5.

[§]Francesc Alted, Ivan Vilata, and others, *PyTables: Hierarchical Datasets in Python*.

Table c.8: Gamma particle energies present in background. These were used in energy calibration. The peak energies are from the National Nuclear Data Center.¹⁵

Parent nuclide	Gamma energy (keV)
²¹⁴ Pb	351.9321 ± 0.0018
²¹⁴ Bi	609.320 ± 0.005
⁴⁰ K	1460.822 ± 0.006
²⁰⁸ Tl	2614.511 ± 0.010

Energy calibration The peak finding step produced a list of the peaks present in each spectrum and the channel about which each peak was centered. The energies of those peaks are known.¹⁵ Using these four to seven data points, we interpolated each channel bin to its corresponding energy. The amount of light produced versus the energy deposited in a sodium iodide crystal is almost linear, but there are some mild nonlinear effects.³⁵ To compensate for these scintillation nonlinearities, we used a quadratic function to interpolate the peak-and-energy data points:

$$E = g(N, \vec{a}) = a_2 N^2 + a_1 N + a_0 \quad ,$$

where the constant coefficients $\vec{a} = [a_0, a_1, a_2]^T$ are unknown, E is the energy, and N is the channel number. With the channel-energy pairs $\{(N_i, E_i)\}$, the coefficients \vec{a} can be determined from the least squares formulation

$$\chi^2 = \min_{\vec{a}} \sum_{i=1}^{n_p} \frac{(g(N_i, \vec{a}) - E_i)^2}{\sigma_{E_i}^2} \quad .$$

Since the number of peaks n_p is four or more, this least squares problem is overdetermined. Once the coefficients are computed, we converted the channel boundaries to energies.

Rebinning spectra Unfortunately, because each spectra is calibrated separately, each spectra has different energy bin edges. To simplify the computing of the net spectrum (gross minus background), every spectrum was rebinned or resampled onto

an identical energy grid. This process is described by Knoll in §18.IV.B.⁶³ To do the rebinning, an assumption must be made about the distribution of the counts within a channel. Knoll proposed fitting a polynomial through the nearby points. We took a simpler approach and assumed that the counts are distributed uniformly in each bin. This assumption guides the redistribution of an old bin's counts into any overlapping new bins. In the rebinning process, the total number of counts is preserved if the range of the new bins is equal to or larger than that of the old bins. The universal energy grid we chose went from 100 keV to 2900 keV, with 1806 edges and 1805 bins. The uncertainties package propagated the standard errors through this summation process.

Background subtraction We are primarily interested in the net counts caused by the ¹³⁷Cs and ⁶⁰Co sources. Thus, the counts from background radiation should be removed. Since all the spectra have the same energy grid, the subtraction is straightforward, provided that we use the correct background. The corresponding background is evident in Figure c.9, where the background spectrum for each location-direction pair is given in the third subblock. Again, the package uncertainties propagated the error through this subtraction procedure. The net spectra can be rebinned to the multigroup energy grid for comparison with transport calculations.

Peak areas After the net spectra were obtained, we desired to compute the area under each identified photopeak because this represents the number of gamma photons at this peak energy that deposit all their energy in the detector. For each photopeak, we fit a Gaussian and exponential. The Gaussian is the model of the peak, and the exponential models the contribution from the continuum. Note that the exponential is a straight line on a linear-logarithmic plot. Together, this function

$$f(x, \vec{p}) = C \exp \left[-\frac{(x - \mu)^2}{2\sigma^2} \right] + ce^{-rx} \quad ,$$

has 5 unknowns $\vec{p} = [\mu, \sigma, C, r, c]^T$. The n_k channels \vec{N}_k are in the range of peak k , corresponding to counting rates \vec{d}_k . Then the model can be fit using a least squares

formulation with the chi-square metric

$$\chi^2 = \min_{\vec{p}} \sum_{i=1}^{n_k} \frac{\left(f(\vec{N}_{k_i}, \vec{p}) - d_i\right)^2}{\sigma_{d_i}^2} .$$

With this definition of χ^2 , the covariance of the final parameters can be estimated as $C_p = (J^T J)^{-1}$, where the Jacobian is defined as [36, §38.1]

$$J_{ij} = \frac{1}{\sigma_{d_i}} \frac{\partial f_i}{\partial p_j} . \quad (\text{C.3})$$

The Jacobian must be evaluated at \vec{p}^* , the values of the model parameters \vec{p} which minimizes χ^2 .

An example spectrum of the background is plotted in Figure c.10. The plot shows the peak fits as well. The chi-squared values for these particular peak fits are provided in Table c.9.

The peak area can be computed by one of two methods. First, the area A under the peak is equal to the integral of the Gaussian over its entire domain $(-\infty, +\infty)$, which is

$$A = C\sigma\sqrt{2\pi} .$$

The error in this estimate is computed by propagating the error on the C and σ least-squares parameters. If the least squares fit is poor, then this method will lead to a large estimated error in the peak area. An alternate method is to simply add the peak channels over the peak width and subtract the area under the linear continuum. This is the method used for this work, and it produced more consistent estimates of the peak area uncertainty.

After they are computed, the peak areas and accompanying uncertainties are saved in an HDF5 database file for the processed data, along with the net spectra. This is the file that is used in the source mapping procedure.




Figure c.10: A representative background spectrum taken with the NaI detector. The black bar — at each peak is the range over which the peak-fitting was performed. The red curves — are the Gaussian component of the fit, and the green lines — are the fit of the continuum.

Table c.9: The normalized χ^2 -value for the four peak fits in Figure c.10. The expected normal distribution around one for the normalized χ^2/n_{dof} is also provided. All of the fits except for the ^{40}K peak fit fall less than two standard deviations from the mean.

Parent nuclide	χ^2/n_{dof}	Expected normal distribution
^{214}Pb	1.002	1 ± 0.21
^{214}Bi	1.20	1 ± 0.15
^{40}K	1.95	1 ± 0.108
^{208}Tl	0.881	1 ± 0.109

A radiation detector gives data as a count. This is the number of successful detection events in the detector. However, to analyze the radiation field, it is necessary to relate the counts reported by the detector to the scalar flux ϕ or inward partial current J_- at the detector. The detector efficiency is a simplification of the full DRF described in section 4.4. While the DRF has two independent variables (in discretized form represented as a matrix), the detector efficiency has zero or one independent variable (represented as a scalar or vector).

D.1 EFFICIENCY DEFINITIONS

Detector efficiencies quantify the effectiveness of detectors in signaling an event as compared to the amount of radiation present. However, this concept is too general to be of much use, and we must be more precise with what signals and radiation we count. There are two different ways to specify these quantities, with both classifications being useful.

Absolute versus intrinsic efficiency The absolute efficiency $\epsilon_{\text{abs}}(E)$ is the fraction of radiation of energy E from the source which creates a signal in the detector. As Knoll defines it [63, §4-VI],

$$\epsilon_{\text{abs}}(E) = \frac{\text{number of pulses recorded}}{\text{number of particles with energy } E \text{ emitted by source}} .$$

The absolute efficiency depends on both detector properties as well as source-detector positioning. For instance, as the source-to-detector distance increases, the absolute efficiency decreases. In contrast, the intrinsic efficiency $\epsilon_{\text{int}}(E)$ depends mostly on the

detector properties. It is defined by Knoll as

$$\epsilon_{\text{int}}(E) = \frac{\text{number of pulses recorded}}{\text{number of particles with energy } E \text{ incident on detector}} .$$

It measures how effectively the detector is able to recognize radiation that enters it. Thus, a detector material with a high atomic number and a higher interaction cross section should have a greater ϵ_{int} . As Knoll points out, the intrinsic efficiency does depend slightly on the detector-source configuration, since this configuration affects the average path length of the radiation as it travels through the detector. Since interaction cross sections change with energy, ϵ_{int} is also a function of E . Intrinsic efficiencies are more frequently quoted because they have wider application.

Peak versus total efficiency In the above classification, we made no distinction between full and partial detection. Consider a monoenergetic radiation field incident on a detector. Some of the radiation particles could pass directly through the detector. For the particles that interact in the detector, they could eventually scatter out of the detector, carrying some of their energy away as primary or secondary radiation, or they could be absorbed, depositing all of their energy in the detector. If the particle escapes the detector with some of its energy, then the corresponding detector signal will be smaller. We would prefer for the detector to absorb all of the radiation particle energy so that the pulse signal would fully represent the initial incident energy. The full energy signals are simpler to extract from the detector spectrum than unfolding the overlapping continuum responses produced by partial-score events. The peak and total efficiencies describe this phenomenon. The peak efficiency is

$$\epsilon_{\text{peak}} = \frac{\text{number of full-energy pulses recorded}}{\text{number of radiation quanta}} ,$$

and the total efficiency is

$$\epsilon_{\text{total}} = \frac{\text{number of pulses recorded (of any energy)}}{\text{number of radiation quanta}} .$$

It is best to specify a particular efficiency as absolute or intrinsic and as peak or total. For gamma detectors, the efficiency is often given as intrinsic peak efficiency.

D.2 ESTIMATING PEAK INTRINSIC EFFICIENCIES

To convert the counts in a photopeak (corresponding to uncollided radiation) to scalar flux, we must apply the peak intrinsic efficiency ϵ_{ip} . Combining the definitions from above, the peak intrinsic efficiency is

$$\epsilon_{ip}(E) = \frac{\text{number of full-energy pulses (with energy } E)}{\text{number of particles with energy } E \text{ incident on detector}} \quad .$$

As mentioned above, this quantity is dependent on the source-detector spatial configuration. In this work, we have averaged this spatial dependence to avoid the need to explicitly account for it in the source mapping procedure. That is, if the true efficiency is $\epsilon_{ip}(\vec{x}, E)$, the average efficiency over the volume V is

$$\epsilon_{ip}(E) = \frac{1}{V} \int_V dV \epsilon_{ip}(\vec{x}, E) \quad .$$

This integral does not typically have a closed-form expression. Thus, we use Monte Carlo methods to compute it, specifically with `MCNP`.

For the experiment described in Appendix C, we used two detector configurations. The first was the bare NaI detector, used to collect a full view of the room. The second configuration was the collimated detector which collected the 2π views. We must compute ϵ_{ip} for both configurations, and for each of the three gamma peak energies. The F8 pulse-height tally of `MCNP` computes the number of full-energy pulses in the detector, using a narrow bin around the peak energy. The F1 current tally can supply the inward partial current, as long as we specify inward and outward cosine bins (the cosine bin edges set at $\{-1, 0, 1\}$). The inward partial current is equal to the total number of particles entering the detector per second.

For the bare detector configuration, we center the $2'' \times 2''$ NaI crystal (density 3.67 g/cm^3) at the origin of the `MCNP` simulation. The detector axis is parallel with one of the coordinate axes. We model a hypothetical isotropic source that is uniformly distributed within $x, y, z \in [0, 500] \text{ cm}$, but excluding the detector volume. The source is located only in one octant of the modeled spatial region to take advantage of the $\frac{1}{8}$ -symmetry. The characteristic dimension of the source region (5 m) is approximately the size of our laboratory. All materials in the simulation are vacuum except the detector. Three simulations are executed, one for each peak energy.

In terms of the tallies, the intrinsic peak efficiency for the bare detector is

$$\epsilon_{ip}(E) = \frac{\text{F8 tally of pulses of energy } E}{\text{F1 tally of inward current of particles of energy } E} \quad ,$$

where the tallies are over the detector volume and in a narrow energy bin around E .

The simulation for the collimated detector configuration is identical to that for the bare detector, except for the addition of the lead (density 11.344 g/cm³) collimator around the detector crystal. A front view is pictured in Figure D.1. Both the detector crystal and lead shield are 2 inches tall. In this case, the source is excluded from both the detector and shield regions. The intrinsic peak efficiency of the crystal itself is the same as for the bare detector

$$\tilde{\epsilon}_{ip}(E) = \frac{\text{F8 tally of pulses of energy } E}{\text{F1 tally of inward current of particles of energy } E} \quad .$$

However, we are interested in the intrinsic efficiency with respect to the crystal and shield configuration. Thus, we must account for the reduction in photons entering the crystal as a result of the collimator. Thus, the actual intrinsic peak efficiency of the collimated detector is

$$\epsilon_{ip}(E) = \frac{J_{-}^{\text{col}}}{J_{-}^{\text{bare}}} \tilde{\epsilon}_{ip}(E) \quad ,$$

where J_{-}^{col} and J_{-}^{bare} are inward partial currents for the collimated and bare detector crystals, respectively.

The computed efficiencies are given in Table D.1. For reference, the intrinsic total efficiencies ϵ_{it} are also provided.

D.3 PARTIAL CURRENT TO SCALAR FLUX

With the intrinsic peak efficiency, we can convert the detector counts in a photopeak A_p to the uncollided partial current J_{-} entering into the detector using the equation

$$A_p = \epsilon_{ip} J_{-} \quad .$$

However, the deterministic solution of the radiation field is in terms of the scalar flux, not the partial current. Thus, we must convert from partial current to scalar flux. If




Figure D.1: Lead collimation around NaI detector crystal.

we know the mean chord length^{65,108} $\langle l \rangle$ in the detector, then the total track lengths traversed in the detector by all incoming particles is $\langle l \rangle \cdot J_-$. Using the track-length definition of the scalar flux [106, §2.2.1], the scalar flux can be written as

$$\phi(E) = \frac{\langle l \rangle J_-(E)}{V} \quad , \quad (\text{D.1})$$

where V is the volume of the detector. This is evaluated using the definition of the mean chord length,

$$\langle l \rangle = \int_0^\infty l p(l) dl \quad . \quad (\text{D.2})$$

Here $p(l)$ is the normalized probability distribution function of the path lengths travelled through the volume.

While Equation D.1 provides the relation we need to convert between partial current and scalar flux, it does so with the introduction of a new unknown, the mean chord length $\langle l \rangle$. Dirac showed that for an isotropic angular flux, the mean chord length in a region with volume V and surface area S is²²

$$\langle l \rangle = \frac{4V}{S} \quad .$$

When computing the spatially-averaged ϵ_{ip} for the bare detector with the hypothetical

Table D.1: Peak and total intrinsic efficiencies for the bare and collimated detectors, as well as the mean chord length, as computed by MCNP. The one-sigma uncertainties are derived from the Monte Carlo statistical errors. For reference, the mean chord length predicted by the Dirac formula is 3.387 cm.

	Energies (keV)		
	662	1173	1332
Bare			
Total	0.549 ± 0.007	0.445 ± 0.006	0.424 ± 0.006
Peak	0.257 ± 0.004	0.140 ± 0.003	0.121 ± 0.003
$\langle l \rangle$ (cm)	3.36 ± 0.04	3.36 ± 0.04	3.36 ± 0.04
Collimated			
Total	0.228 ± 0.004	0.222 ± 0.004	0.221 ± 0.004
Peak	0.096 ± 0.002	0.060 ± 0.002	0.054 ± 0.002
$\langle l \rangle$ (cm)	3.31 ± 0.06	3.38 ± 0.06	3.39 ± 0.06

uniformly distributed isotropic source, the flux is nearly isotropic, and so this equation is valid. Note that the flux at a detector resulting from a point source (or a few point sources) is not isotropic. It is only with the spatial averaging of the source location that we can assume that the flux is isotropic.

However, for the shielded detector, the flux in the detector is not isotropic. As we did for the efficiency, we used Monte Carlo to evaluate the integral in Equation D.2. The simulations are similar to those to compute ϵ_{ip} , but here the detector volume is modeled as vacuum. (The collimator is still lead.) The NaI material is excluded because we are converting between the in-air scalar flux and the in-air partial current. The effect of the material is captured by the efficiency ϵ_{ip} .

The mean chord lengths are given in Table D.1. They are a weak function of energy because $p(l)$ changes with the penetrating power of the gammas through the lead shield.

D.4 FROM COUNTS TO FLUX IN THE 2144 BEL EXPERIMENT

We use this method to convert the count rate for a photopeak to the scalar flux in the 2144 BEL experiment described in Appendix c. This allows the comparison of the experimental data to the response computed by the deterministic or Monte Carlo transport calculations. This process also provides the key link between the count rate and scalar flux necessary for the source mapping algorithm.

In Figure D.2 and Figure 5.11, the comparison is made between the measured data and the model predictions assuming the source distribution is known. Figure D.2 compares the measured count rate in each photopeak with that computed by Monte Carlo or analytic means. (In this case, the transport calculation is simply the geometric attenuation because we are dealing with the uncollided flux.) The appropriate efficiency and mean chord length from Table D.1 are used in Equation D.1 and Equation D.2 to convert the predicted flux to the predicted count rate.

In Figure D.2, it is not surprising that the experiment registers zero counts for the directions pointing away from the source (roughly, the points where the model response is less than 10^{-3}). However, those isotropic (full) measurements and the directions pointing toward the source (roughly the points for which the model is greater than 10^{-2}) should not be zero. The fact that some are zero where they should be nonzero points to the difficulty in finding peaks in noisy spectra. This difficulty is exacerbated in this case because the radiation source peaks are often of a similar size as the background peaks. The negative consequences of these omissions are considered further in chapter 5.




Figure D.2: Measured and modeled photopeak responses. No ^{137}Cs and ^{60}Co measurements were taken from location 5. The experiment value is zero where there is no blue experiment point accompanying a green model (predicted) value.

RADIATION SOURCE RECONSTRUCTION WITH KNOWN GEOMETRY AND MATERIALS USING THE ADJOINT*



E.1 ABSTRACT

We present a method to estimate an unknown isotropic source distribution, in space and energy, using detector measurements when the geometry and material composition are known. The estimated source distribution minimizes the difference between the measured and computed responses of detectors located at a selected number of points within the domain. In typical methods, a forward flux calculation is performed for each source guess in an iterative process. In contrast, we use the adjoint flux to compute the responses. Potential applications of the proposed method include determining the distribution of radio-contaminants following a nuclear event, monitoring the flow of radioactive fluids in pipes to determine hold-up locations, and retroactive reconstruction of radiation fields using workers' detectors' readings. After presenting the method, we describe a numerical test problem to demonstrate the preliminary viability of the method. As expected, using the adjoint flux reduces the number of transport solves to be proportional to the number of detector measurements, in contrast to methods using the forward flux that require a typically larger number proportional to the number of spatial mesh cells.

E.2 METHODS

A common method to resolve inverse problems is to iterate on the unknown inputs using repeated forward calculations. The new input guess can be found using a

*This appendix is a reprint of a paper of the same title by Joshua M. Hykes and Yousry Y. Azmy appearing in the 2011 International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2011), Rio de Janeiro, RJ, Brazil, 2011. Two sections have been omitted due to their similarity with section 3.2 and Appendix A.1.

gradient-based or stochastic method. This is the approach taken in Refs. [10, 77]. If the forward model can be represented by a mapping $T_\alpha x = y(\alpha)$, and the computed responses by an inner product $r = \langle c, x \rangle$, where α is the unknown input, the typical algorithm is then:

- Choose initial guess α_0 .
- Loop for $i = 0, 1, \dots$ until r_i is close to the measured response:
 - Compute operator T_{α_i} and right hand side $y(\alpha_i)$.
 - Solve forward problem $T_{\alpha_i} x_i = y(\alpha_i)$.
 - Compute response $r_i = \langle c, x_i \rangle$.
 - Update inputs $\alpha_i \rightarrow \alpha_{i+1}$.

Since solving the forward problem can be a computationally-expensive task, it is undesirable that the task is inside the iteration loop. For inverse problems where the geometry and materials change significantly from iteration to iteration, this may be the only viable option. However, for problems in which the materials and geometry are known, there is another option.

It is well known that the response can be computed from the adjoint flux. In operator form, the Boltzmann transport equation is

$$L\psi = q \quad , \quad (\text{E.1})$$

where L is the streaming, collision, and scattering operator, ψ is the angular flux, and q is the external source. The inner product of this equation with the so-called adjoint flux ψ^+ results in

$$\langle \psi^+, L\psi \rangle = \langle L^+ \psi^+, \psi \rangle = \langle \psi^+, q \rangle \quad ,$$

where $\langle \circ, \circ \rangle$ denotes the inner product over independent variables. If we require that ψ^+ satisfy $L^+ \psi^+ = \sigma_d$, where σ_d is the detector response function, then $r = \langle \psi^+, q \rangle$.

Thus using the adjoint, we can modify the iteration algorithm to eliminate the expensive transport solve from inside the loop. Switching to the transport notation to emphasize the specificity of the algorithm, the optimization procedure becomes:

- Choose initial guess for source q_0 .
- Compute ψ^+ from $L^+ \psi^+ = \sigma_d$ for each detector reading.
- Loop for $i = 0, 1, \dots$ until r_i is close to the measured response:

Compute response $r_i = \langle \psi^+, q_i \rangle$.

Update source $q_i \rightarrow q_{i+1}$.

Note that an adjoint is necessary for each detector reading. If we have l detector locations and g energies (in groups or lines), then we need lg adjoint calculations. Also note that this algorithm is still general, since the method of updating the source is unspecified. In the following, we proceed to a concrete algorithm.

Imagine we have l detector locations and g energy groups (or spectral lines) per location, stored in the vector $\vec{r}_m \in \mathbb{R}_+^{lg}$ (m for measured). Here the nonnegative real line is written as $\mathbb{R}_+ = [0, \infty)$. The spatial domain is discretized into n cells (for instance, a cube mesh with $\sqrt[3]{n}$ cells per side). We compute a scalar adjoint flux $\vec{\phi}_i^+$, with a detector response as the adjoint source, for each detector location and group for $i \in \{1, \dots, lg\}$, such that

$$L^+ \vec{\phi}_i^+ = \vec{\sigma}_{d,i} \quad ,$$

where $\vec{\phi}_i^+, \vec{\sigma}_{d,i} \in \mathbb{R}_+^{ng}$ and $L^+ \in \mathbb{R}^{ng \times ng}$. Then we can compute the expected detector responses $\vec{r}_c \in \mathbb{R}_+^{lg}$, where the i -th detector response from a source guess $\vec{q} \in \mathbb{R}_+^{ng}$ is the well-known inner product (see, for example, §1-6 in [72])

$$[\vec{r}_c]_i = (\vec{\phi}_i^+)^T \vec{q} \quad \text{for } i = 1, \dots, lg \quad .$$

Here \vec{x}^T is the transpose of \vec{x} .

The transposed and stacked $\vec{\phi}_i^+$ vectors form the matrix

$$S = \begin{bmatrix} (\vec{\phi}_1^+)^T \\ (\vec{\phi}_2^+)^T \\ \vdots \\ (\vec{\phi}_{lg}^+)^T \end{bmatrix} \quad ,$$

so that $S \in \mathbb{R}^{lg \times ng}$. Then the computed responses \vec{r}_c are equal to the matrix-vector product

$$\vec{r}_c = S \vec{q} \quad .$$

Thus the cost functional $f : \mathbb{R}_+^{ng} \rightarrow \mathbb{R}_+$ that we seek to minimize can be written as

$$f(\vec{q}) = \|\vec{r}_m - \vec{r}_c\| = \|\vec{r}_m - S \vec{q}\| \quad .$$

Since negative sources are unphysical, the best source estimate is

$$\vec{q}^* = \arg \min_{\vec{q} \in \mathbb{R}_+^{ng}} \|\vec{r}_m - S\vec{q}\| \quad . \quad (\text{E.2})$$

If S is square and has full rank, one could find \vec{q}^* such that $f(\vec{q}^*) = 0$. If $l \neq n$ and we use the 2-norm, the minimization becomes a least squares problem. If $l < n$, the problem is underdetermined, and a potentially infinite number of solutions to the inverse problem exist. Although we had limited success with underdetermined problems with peaked (not widely distributed) sources, selecting from the class of admissible solutions begs for more information. Thus, we focus on determined and overdetermined systems here.

The constrained form of Equation E.2 is commonly called *nonnegative least squares* (NNLS).⁷ A number of algorithms exist to solve NNLS problems. For all numerical results below except Figure E.4E, we used the Fortran subroutine `nls.f` described in Ref. [71] and available on netlib.

E.2.1 Basis of the source reconstruction

It is likely that the number of spatial cells n will exceed the number of detector readings per group l . Without other modification, this would lead to an underdetermined system. This motivates us to reduce the number of degrees of freedom of the unknown source. We phrase this process in terms of the basis in which the source is expanded.

One could think of any number of basis functions for the source. Implicitly, the basis thus far has been the canonical basis, that is, ng basis vectors with each vector having one element equal to 1 and all other elements equal 0. We could form a canonical basis on a coarser spatial mesh, thus constraining the source to a constant over groups of finer mesh cells. Or one could coarsen some cells but not others, leading to an adaptive basis. If the source distribution is believed to be smooth over the domain, one could use a polynomial or Fourier basis. Obviously, the choice of basis should incorporate any prior knowledge of features of the source distribution.

For many of the non-canonical bases, there exists a linear map $P : \mathbb{R}^{pg} \rightarrow \mathbb{R}^{ng}$ so that

$$\vec{q} = P\vec{x} \quad . \quad (\text{E.3})$$

In this case, p is the spatial dimension of the smaller basis. We would choose $p \leq l$, so that the transformed problem

$$\min \|\vec{r}_m - A\vec{x}\| \quad \text{subject to } [\vec{q}]_i = [P\vec{x}]_i \geq 0 \text{ for } i = 1, \dots, ng \quad (\text{E.4})$$

is determined or overdetermined. Here $A = SP \in \mathbb{R}^{lg \times pg}$.

In the following numerical tests, we use an adaptive basis to ensure the problem is not underdetermined.

E.3 NUMERICAL RESULTS

E.3.1 Test problem

We constructed a two-group homogeneous system, described in Table E.1, to demonstrate our source reconstruction method. For simplicity, the source distribution varies in only the x dimension. In the finest mesh considered, the unknown source vector has $2 \cdot 3^4$ elements. Here we will only treat square linear systems, that is, the number of source unknowns equals the number of detector readings. In all the following, the detector response is unity in the group and spatial cell of interest, and zero elsewhere.

Full detector coverage The simplest way to solve this problem is to use $2 \cdot 3^4$ detector readings, one for each group per fine mesh cell, and solve for $2 \cdot 3^4$ unknown sources. This is shown in Figure E.1. Both the forward calculation and the adjoint calculations use the same 3^4 mesh, making this a blatant inverse crime. As we would expect, for this “illegal” system, it is possible to accurately identify the true source. While this is helpful to show the consistency of the method, it does not prove much about a real problem, except, perhaps, that the method still has a chance. The next series of tests moves closer to reality.

Adaptive source basis To reduce the number of detector measurements to a more realistic number, we switch from the canonical basis of the previous test to an adaptive basis using the formalism of §E.2.1. The concept of this basis is simple. Instead of jumping immediately to the finest mesh, this strategy begins at a very coarse mesh, where only a few detectors are needed. The source is then reconstructed on this coarse

Table E.1: Description of the source-reconstruction problem. All faces had vacuum boundary conditions.

(a) Spatial and material parameters.

Spatial Domain (cm)	Macroscopic 2-Group Cross Sections (/cm)	Spatial Mesh	Angular Quadrature
$x \in [0, 30]$ $y, z \in [0, 10]$	$\sigma_t = \begin{bmatrix} 1/10 \\ 1/2 \end{bmatrix}, \quad \sigma_s = \begin{bmatrix} 1/40 & 0 \\ 1/20 & 1/10 \end{bmatrix}$	uniform $3^4 \times$ 1×1	LQ_{16} [72, §4-2]

(b) Source parameters. The source in cell $\vec{c}[i]$ and group g is $\vec{q}[g, \vec{c}[i]] = \vec{s}_g[g] \cdot \vec{s}_i[i]$, with $g \in \{1, 2\}$ and $i \in \{1, \dots, 7\}$. All cells not in \vec{c} have zero source. See solid lines in Figure E.1.

$$\vec{s}_g = [1 \quad 1/5]^T$$

$$\vec{c} = [13 \quad 14 \quad 15 \quad 50 \quad 51 \quad 52 \quad 53]^T$$

$$\vec{s}_i = [1/2 \quad 1 \quad 3/4 \quad 1 \quad 1.1 \quad 1.2 \quad 1.3]^T$$




Figure E.1: Reconstructed source using 81 detector locations.

basis. The source basis (or mesh) is only refined in cells in which the solution exceeds a threshold δ . Then the appropriate number of detector measurements are added, and the source is recomputed on the refined mesh. This refinement process is repeated several times.

From a physical standpoint, this process represents an incremental deployment of detectors. Initially, only a small number of measurements are taken, since it should be possible to make a rough estimate of the sources based on these few readings. After the rough estimate is computed, more measurements can be taken to refine the estimate. Using this approach, we focus our attention on the important regions of the phase space. This should be more efficient than blanketing the entire space with detectors from the start.

To understand the procedure fully, one must distinguish three meshes:

- The *forward-computation* mesh is the finest uniform mesh. It is not refined.
- The *adjoint-computation* mesh is uniform at each refinement level, and it is refined uniformly.
- The *source-adaptation* mesh is non-uniform and adaptively refined. The smallest possible source-adaptation cell size in a given level is equal to the adjoint-computation cell size at that same level.

Figure E.2 shows the adaptive refinement process in action. In the beginning, the algorithm starts with a coarse adjoint-computation and source-adaptation mesh, in this case 3^1 . Six “detector readings” are computed, one per group in each coarse cell, taking the forward flux from the fine cell at the center of each coarse cell. Then $2 \cdot 3^1$ adjoint fluxes are computed, and the system of equations is solved. This stage is shown in Figure E.2A.

The next step is to move to a finer adjoint-computation mesh. Here we split each adjoint-computation cell in three, yielding 3^2 cells. If the source solution in a cell in the previous level was greater than the prescribed small criterion δ , set in this work to a few percent of the largest source magnitude, we split this source basis cell to match the current adjoint-computation cell size. Otherwise, this source-adaptation cell remains coarse. This typically reduces the number of unknowns, thus reducing the need for additional detector readings. Notice in Figure E.2B, there are $2 \cdot 3^2$ source unknowns. This is because each source solution in Figure E.2A was greater than δ . We include a detector reading (one per group) corresponding to the fine cell located at the


center of the coarse source-adaptation cell, meaning that the source cells and detectors are aligned.

This process repeats in Figures E.2C and E.2D. In Figure E.2C, the adaptive mesh requires only $2 \cdot 13$ detector readings instead of $2 \cdot 3^3$. For Figure E.2D, we need only $2 \cdot 19$ readings instead of $2 \cdot 3^4$, for the same fidelity as Figure E.1.


In this procedure, the detector locations were aligned with the source basis functions. This seems unlikely in applications, so in the next test we force source-detector misalignment.

Adaptive source basis with unaligned detectors Here we space the detectors evenly throughout the x domain, leaving off detectors at the right if necessary to match the number of source unknowns. In other respects, the procedure is identical to the previous test. Figure E.3 shows the process, with a plus mark indicating the location of a detector. The one significant difference in the solution versus Figure E.2 occurs for the third level, the right peak in Figure E.3C. This under-prediction is unsurprising, given that there is no detector in the errant source cell.


Computing “measured” detector readings with Monte Carlo To move closer to reality, the forward-computed measurement must come from a model other than the same deterministic model used for the adjoint, thus avoiding the inverse crime. In this case, we use a Monte Carlo solution for the forward flux, tallied in the same fine mesh as the deterministic forward solution. We used Monaco²⁶ from SCALE for this purpose, since supplying fictitious multigroup cross sections is straightforward in Monaco. Figure E.4 shows the results using the Monte Carlo responses. The first three levels show negligible differences from Figure E.3. However, the fine mesh solution in Figure E.4D is no longer accurate. This is the first sighting of ill-posed behavior, where a small change in the detector readings, about 10^{-2} , causes a large change in the obtained source distribution. The condition number of S is $4 \cdot 10^8$, which explains the large change in \vec{q} resulting from a small change in \vec{r}_m . If we apply truncated SVD regularization to the system, we can recover more of the solution (see Figure E.4E), but the computed source distribution no longer matches perfectly with the true distribution employed in the forward Monte Carlo computation. In the TSVD procedure, the singular value relative threshold was set equal to the discrepancy between the deterministic and MC forward fluxes.




(a) 3^1 adjoint-computation mesh, 3 detector locations.



(b) 3^2 adjoint-computation mesh, 9 detector locations.




(c) 3^3 adjoint-computation mesh, 13 detector locations.




(d) 3^4 adjoint-computation mesh, 19 detector locations.


Figure E.2: Reconstructed sources using adaptive procedure and “aligned” detectors. Each filled circle or x-mark denotes one source unknown. The detectors are located at the source unknowns.




(a) 3^1 adjoint-computation mesh, 3 detector locations.



(b) 3^2 adjoint-computation mesh, 9 detector locations.




(c) 3^3 adjoint-computation mesh, 13 detector locations.




(d) 3^4 adjoint-computation mesh, 19 detector locations.


Figure E.3: Reconstructed sources using adaptive procedure with evenly-spaced detectors. The plus marks indicate the x position of the detectors. The vertical positioning of the pluses on the graph is meaningless.




(a) 3^1 adjoint-computation mesh, 3 detector locations.



(b) 3^2 adjoint-computation mesh, 9 detector locations.




(c) 3^3 adjoint-computation mesh, 13 detector locations.



(d) 3^4 adjoint-computation mesh, 19 detector locations, not regularized.

Figure E.4: Reconstructed sources using adaptive procedure with Monte Carlo forward flux.



(e) 3^4 adjoint-computation mesh, 19 detector locations, regularized with TSVD.

Figure E.4: Reconstructed sources using adaptive procedure with Monte Carlo forward flux. (continued)

E.3.2 Comparison with an iterative method

A common method for the solution of this type of problem is the Levenberg-Marquardt (LM) algorithm. We implemented a solver using the LM method for the test problem with Monte Carlo-computed detector responses for a preliminary comparison of the adjoint method presented above to LM. Other than changing the inverse solver, the problem is identical to the adaptive source basis problem, whose solution was shown in Figure E.4. We used the `levmar` package⁷³ to solve for the unknown source using the functions for constrained LM. The number of transport solutions, forward solutions for LM and adjoint solutions for the adjoint method, is given in Table E.2, along with the number of iterations in the LM method. The results should be treated with caution because adjusting the convergence tolerances and other tuning would surely decrease the costs of the LM solver.

The LM algorithm requires derivatives of the detector responses with respect to the unknown parameters. In both the Favorite-Ketcheson paper and the Miller-Charlton paper, the authors use an adjoint to estimate these derivatives,^{33,86} while Mattingly and Mitchell use the simpler finite-difference approximation.⁷⁷ In this test, we use two approaches. First, we can compute the Jacobian with respect to \vec{q} since the problem is linear. For example, the Jacobian in Equation E.4 is the matrix $-A$. This is a simplification of the approach taken by the Favorite and Miller papers, since we only need the derivatives with respect to \vec{q} . In the tally of forward transport solves in Table E.2 for the analytic Jacobian, the count does not include the transport solves necessary for computing the Jacobian. Second, we let `levmar` compute the Jacobian

Table E.2: Comparison of number of transport solves needed for the adjoint method and the forward Levenberg-Marquardt iterative method. TS stands for the number of transport solves, and Its. is the number of iterations.

Adjoint Mesh	Adjoint	LM, analytic Jacobian		LM, finite-difference Jacobian	
	TS	Its.	TS	Its.	TS
3^1	6	5	6	2	27
3^2	18	28	398	40^b	1629
3^3	26	32	550	40^b	2348
3^4	38	22^a	485^a	40^b	3478

^a Iteration converges to errant solution.

^b Maximum number of LM iterations encountered.

using the finite-difference method. For both means of computing the Jacobian, if the number of transport solves is larger than the number of spatial cells times the number of groups, then it is more efficient in execution time to compute and store the inverse L^{-1} of the transport operator, but this may incur an unbearable memory cost.

As in other applications, the benefit of using the adjoint is a reduction in the number of transport solves from the order of the number of mesh cells to the order of the number of detector readings, that in our method is the same number as the distinct source values sought. In general, this is a favorable trade.

E.4 CONCLUSIONS

We have presented a method to efficiently reconstruct an unknown source distribution and spectrum using the adjoint flux if the geometry and material composition are known. The problem can be stated as a nonnegative least squares problem. Selecting an appropriate basis for the source is an important first step. In the examples presented, an adaptive basis method proves worthwhile for localized source distributions. The method is consistent for problems with an inverse crime, which shows its initial viability. For the problems with coarse meshes or the Monte Carlo-computed responses, the method proves suitable, although additional regularization is required. Using the adjoint flux becomes advantageous when the number of detectors is significantly smaller than the number of spatial cells in the transport calculation.



F.1 SETS, SPACES, AND SEQUENCES

Before discussing operators, we briefly describe the basics of set theory. Sets are collections of objects, for example: muffins, natural numbers, or continuous functions. Each object in a set is called an *element*. Sets are typically denoted by capital letters (X, Y, \dots) and their elements are usually listed within braces. Thus, the muffin set consisting of three types of muffins can be written

$$M = \{\text{blueberry muffin, chocolate chip muffin, lemon poppy seed muffin}\} \quad . \quad (\text{F.1})$$

The set of natural numbers is $\mathbb{N} = \{1, 2, \dots\}$, where the ellipsis indicates that the set continues indefinitely. Sets can also be specified with a condition, for example

$$X = \{x \mid x \text{ is a continuous function}\} \quad . \quad (\text{F.2})$$

The vertical bar $|$ is read “such that.” Sets contain both elements and subsets. If x is an element in X , we write $x \in X$. A set X contains a subset Y if every element of the subset is also in the set, that is, if $y \in Y$, then also $y \in X$. Then we write $Y \subset X$ or $X \supset Y$. Note that X contains itself, so $X \subset X$.

The *complement* of a subset A in a set X is

$$A^C = \{x \mid x \notin A \text{ and } x \in X\} \quad .$$

A subset is *open* if every point in the subset is contained in a ball that is also in the subset. A subset is *closed* if its complement is open. For example, the set consisting of all the points inside and on a sphere is not open, since points on the sphere can

not be contained in a ball entirely enclosed in the subset. An open ball, $B(x_0; r) \in X$, has a center at x_0 and a radius of r , and does not include the points at a distance r from x_0 . This radius is measured in terms of a metric (to be discussed below). An ϵ -neighborhood of x_0 is defined as $B(x_0; \epsilon)$. A *neighborhood* of x_0 is any subset of X containing an ϵ -neighborhood of x_0 .⁶⁴

Sets with an accompanying structure or tool are called *spaces*. Four common spaces are metric, vector, normed, and inner product spaces. Metric spaces are sets with a tool to measure distance between elements. (In the following sentences, x and y are elements of the respective space, and α is a scalar.) This tool is typically represented as $d(x, y)$. Vector spaces have two operations defined for elements: addition $x + y$ and scalar multiplication αx . Normed spaces have a tool, called a *norm* and written $\|x\|$, to measure an individual element's size. Inner product spaces are sets with a tool to measure the "alignment" of two elements. This tool is called an inner product $\langle x, y \rangle$, and it generalizes the angle between two vectors in Euclidean space. Both normed and inner product spaces are also vector spaces. These structures must satisfy certain axioms, which can be found in the first three chapters of the text by Kreyszig.⁶⁴

Sequences and their limits in spaces are important in analysis. A sequence x_1, x_2, \dots in a metric space X is said to converge to x if the limit

$$\lim_{n \rightarrow \infty} d(x_n, x) = 0 \quad .$$

When examining limits, it is often helpful to examine the distance between two elements in the tail of the sequence, which is formalized by Cauchy sequences. A *Cauchy sequence* in a metric space X is a sequence such that for every small tolerance $\epsilon > 0$ there is a threshold $N = N(\epsilon)$ such that⁶⁴

$$d(x_m, x_n) < \epsilon \quad \text{for every } m, n > N \quad .$$

If the limit of every Cauchy sequence in X converges to an element in X , the space is said to be *complete*. A complete normed vector space is called a Banach space, and a complete inner product space is called a Hilbert space. In this work, the inputs and outputs live in Hilbert spaces.




Figure F.1: Two-dimensional representation of a mapping T . Here $Tx_1 = y_1$ and $Tx_2 = y_2$.

F.2 OPERATORS

If the inputs and outputs of a system live in two spaces, we need a method to accept the input element and transform it to the output element. This is commonly called a *mapping*. Assuming X and Y are two sets, a mapping T of elements from X to elements in Y is written

$$T : \mathcal{D}(T) \rightarrow Y \quad , \quad (\text{F.3})$$

where $\mathcal{D}(T) = A \subset X$. $\mathcal{D}(T)$ is the *domain* of T and $\mathcal{R}(T)$ is the *range*. For T to be a mapping, T must associate every $x \in A$ with only one element $y \in Y$. The *image* of x with respect to T is $y = Tx$. We can also have an image of a set $M \subset A$, where the image of M is $T(M)$, and $T(\mathcal{D}(T)) = \mathcal{R}(T)$. A visualization of a mapping T is shown in Figure F.1. If X and Y are vector or normed spaces, then a mapping is called an *operator*.⁶⁴

There are three types of mappings: injective, surjective, and bijective mappings. An *injective* mapping takes each distinct point in the domain to a unique point in the range. Mathematically, for two points in the domain of the mapping T , $x_1, x_2 \in \mathcal{D}(T)$, if x_1 and x_2 are not equal, then the two images Tx_1 and Tx_2 are also not equal:

$$x_1 \neq x_2 \implies Tx_1 \neq Tx_2 \quad (\text{injective}).$$

Injective mappings are called *one-to-one* mappings.

A *surjective* mapping, also known as *onto* mapping, covers the entire space in which the image lives. If $T : \mathcal{D}(T) \rightarrow Y$, then T is surjective if the range $\mathcal{R}(T)$ is equal to Y :

$$\mathcal{R}(T) = Y \quad (\text{surjective}).$$

A mapping is *bijective* if it is both injective and surjective (or less formally, one-to-one and onto).

If T is bijective, then we can define the inverse $T^{-1} : Y \rightarrow \mathcal{D}(T)$, where $T^{-1}y_0 = x_0$ if $Tx_0 = y_0$. If T is injective but not surjective, we can still define an inverse, but we must limit its domain to the range of T , so that $T^{-1} : \mathcal{R}(T) \rightarrow \mathcal{D}(T)$.

In Tartaglia's cannonball example, is the map $T : X \rightarrow Y$ to compute flight distance r from the angle of elevation θ injective or surjective? If $X = Y = \mathbb{R}$, the real line, then T is neither. Since $\sin 0 = \sin \pi = 0$, T is not one-to-one. It cannot be onto because the range $\mathcal{R}(T) = [-v_0^2/g, v_0^2/g] \neq \mathbb{R}$. However, if $X = [0, \pi/2]$ and $Y = [0, v_0^2/g]$, the mapping T is surjective, but it is still not one-to-one, since two angles can give the same range.

F.2.1 Adjoint operators

The adjoint $T^+ : H_2 \rightarrow H_1$ of an operator $T : H_1 \rightarrow H_2$ for Hilbert spaces H_1 and H_2 is defined such that

$$\langle Tx, y \rangle_{H_2} = \langle x, T^+y \rangle_{H_1} \quad , \quad (\text{F.4})$$

for all $x \in H_1$ and $y \in H_2$.⁶⁴ The inner products in H_1 and H_2 are denoted $\langle \circ, \circ \rangle_{H_1}$ and $\langle \circ, \circ \rangle_{H_2}$, respectively.