ANALYSIS OF SIGNAL DETECTABILITY IN STATISTICALLY RECONSTRUCTED TOMOGRAPHIC IMAGES

by

Anastasia Yendiki

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical Engineering: Systems)
in The University of Michigan
2005

Doctoral Committee:

Professor Jeffrey A. Fessler, Chair
Professor Alfred O. Hero III
Assistant Professor Elizaveta Levina
Assistant Professor Thomas E. Nichols
Research Professor Neal H. Clinthorne

To my mother [43] and to the women's movement [94].

ACKNOWLEDGEMENTS

My dissertation owes its existence to the invaluable mentorship, support, and patience of Professor Jeff Fessler. It benefited immensely from the careful reading and thoughtful comments of Professors Clinthorne, Hero, Levina, and Nichols.

On the personal front, I am grateful to my parents for financing the costly operation that has been my education. I am also grateful to my dearest friends, Katerina Argyraki and Dimitris Makris, for being my partners in crime throughout this education. Going through this character-building experience, and the ones to come, at the same time as them (although at too distant locations) has been a great source of support and, often, entertainment. Finally, various individuals in Ann Arbor, who shall remain nameless, have shaped my experience here, making it at once easier and more difficult. I wouldn't be the person that I am without them and I will be grateful to them always.

TABLE OF CONTENTS

CHAPTER

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

ABSTRACT

ANALYSIS OF SIGNAL DETECTABILITY IN STATISTICALLY RECONSTRUCTED TOMOGRAPHIC IMAGES

by

Anastasia Yendiki

Chair: Jeffrey A. Fessler

Imaging in general, and emission tomography in particular, has become an important tool in many areas of medical diagnosis. Several common applications of emission tomography, such as the diagnosis of lung tumors or myocardial perfusion defects, involve the detection of a spatially localized target signal in an image reconstructed from noisy data. Such detection tasks are affected by various design parameters of the imaging system and reconstruction algorithm. This thesis is concerned with optimizing regularized image reconstruction methods for emission tomography with respect to the detectability of a spatially localized target signal in the reconstructed images.

We first consider the task of detecting a statistically varying signal of known location on a statistically varying background in a reconstructed tomographic image. We show that a broad family of linear observer models can achieve exactly optimal detection performance in this task if one chooses a suitable reconstruction method. This conclusion encompasses several well-known models from the literature, including those with a frequency-selective channel mechanism. Interestingly, the "optimal" linear reconstruction methods for many of these observer models are unregularized and in some cases quite unconventional. In the case of channelized models in particular, the observer's ability to prewhiten determines the extent to which its detection performance can benefit from regularization. That is, regularization is more important for channelized observers that have incomplete knowledge of the second-order statistics of the reconstructed images. Therefore, such models are more useful for the purposes of designing regularized reconstruction methods that optimize lesion detectability.

Subsequently, we investigate detection tasks where the location of the target signal is unknown to the observer. This location uncertainty complicates the mathematical analysis of observer performance significantly. We consider model observers whose decisions are based on the maximum value of a linear local test statistic over all possible signal locations. Several of our conclusions about the known-location task extend to this case. Previous approaches to this problem have used Monte Carlo simulations to evaluate the localization performance of maximum-statistic observers. We propose an alternative approach, where approximations of tail probabilities for the maximum of correlated Gaussian random fields facilitate analytical evaluation of detection performance. We illustrate how these approximations can be used to optimize the probability of detection (at low probabilities of false alarm) for the observers of interest. Using our analyses, one can rank and optimize image reconstruction methods with respect to unknown-location detectability without the need for time-consuming Monte Carlo simulations.

CHAPTER 1

Introduction

Objective evaluation of the quality of an image requires specifying the goal that the image will be used to achieve. Typical uses of medical images can be categorized into estimation, where the goal is to measure the value of some parameter of interest (*e.g.*, radioactivity uptake within an organ), and classification, where the goal is to decide among several possible states of the truth (*e.g.*, the presence or absence of a lesion). By evaluating images with respect to a task of interest, one can compare and rank the imaging systems or reconstruction algorithms that produce these images.

In several types of medical diagnosis, the goal is to detect a small abnormality, such as a lesion, in an image reconstructed from noisy tomographic data. Faced with choosing among a plethora of image reconstruction methods, one would like to find which method best facilitates this classification task. Furthermore, several reconstruction methods involve parameters that control the variance-resolution trade-off in the reconstructed images. Rather than resorting to trial-and-error, one would like to have a fast analytical method to determine optimal or at least reasonable choices for these parameters. Our objective is to optimize image reconstruction methods with respect to the performance of observers attempting to detect small features in the reconstructed images. We are particularly interested in optimizing the regularization parameters involved in penalized-likelihood methods.

We begin our treatment of this problem, in Chapter 2, with an overview of tomographic imaging systems and the basic principles of reconstructing images from data acquired by such systems. We also provide a mathematical formulation for the detection tasks considered in this work and the observers that perform such tasks.

In a typical clinical setting, these tasks are performed by human observers. However, not only are experiments with human observers time-consuming, but also the performance of human observers does not lend itself to optimization through analytical tools. Thus, our objective of optimizing reconstruction parameters analytically is better served by mathematical observer models.

Statistical decision theory can be used to derive an ideal observer model that takes advantage of all available statistical information to achieve optimal performance in a specific classification task. However, because the performance of this ideal observer is invariant to any nonsingular data transformation, it is not useful for evaluating and ranking image reconstruction methods. Optimization of reconstruction methods with respect to detectability must instead be performed for non-ideal observer models. Furthermore, researchers have found that human observers themselves do not perform ideally in simple detection tasks. As a result, various linear observer models that account for human suboptimality have been proposed in the literature and their performance in such tasks has been tested against that of humans with various degrees of success.

The performance of these suboptimal observer models is particularly amenable to analysis in the simple special case where the location of the signal to be detected is known a priori. As a result, the analysis of detection performance in tomographically reconstructed images has focused to a significant extent on such known-location tasks. These are also the tasks that we use, in Chapter 3, as our starting point for gaining intuition into the problem. We show that a broad family of linear observer models can exactly achieve optimal performance in the detection of a statistically varying signal of known location on a statistically varying image background without the need for regularization. Using frequency-domain approximations, we achieve fast evaluation of, and further insight into, the detection performance of model observers in these tasks.

Since in clinical practice the potential location of a lesion is typically not known in advance, efforts have been shifting lately to the analysis of the detection performance of observer models for tasks that involve location uncertainty. We investigate these models in Chapter 4. Due to the location uncertainty, the probability distribution of the decision variables associated with these models is difficult to obtain in closed form, thus complicating the analysis significantly. As a result, previous approaches to evaluating the performance of these observers have resorted to time-consuming simulations. We propose an alternative approach that can be used to evaluate unknown-location detection performance analytically. Our approach utilizes approximations for the tail distribution of the maximum of correlated Gaussian random fields. Such approximations have been applied to the problem of detecting activation in functional neuroimaging. We use them here to evaluate measures of unknown-location lesion detectability in tomographically reconstructed images.

In Chapter 5 we summarize the results of our work on the optimization of tomographic image reconstruction methods with respect to detection performance and we suggest some future goals.

The main contributions of this work are:

- An analytical investigation of the reconstruction methods that lead to optimal performance in known-location detection tasks for various commonly used observer models [41, 149, 151].
- An analytical investigation of the behavior of model observers when combined with

practical regularized reconstruction methods. In particular, our analysis illustrates the role of the observer's knowledge of second-order image statistics, as it relates to optimizing regularization with respect to detection performance.

- Derivation of frequency-domain approximations for detectability measures, using a continuous-to-discrete model for the system. This approach leads to a refinement of frequency-domain approximations and can be useful when the latter must be computed locally at every pixel in the image.
- Fast, analytical methods for ranking and optimizing image reconstruction with respect to observer performance in tasks that involve location uncertainty [148]. These methods obviate the need for time-consuming simulations that are required by other approaches to the unknown-location problem.

Finally, another part of our work focuses on system models for statistical tomographic reconstruction and their effect on quantification tasks, particularly for single-photon emission computed tomography (SPECT) systems. We present some results from this work in appendix A. Contributions in this area include:

- Design and implementation of system models for fully 3-D SPECT reconstruction, including a new family of SPECT system models that assume the imaged object to be composed of spherically symmetric volume elements.
- Investigation of the volume elements used by the system model and their effect on the quantification of radiotracer uptake in images reconstructed from simulated phantom data [147,150].
- Investigation of the method for depth-dependent response compensation used by the system model and its effect on the quantification of radiotracer uptake in images reconstructed from real patient data [72–75].

CHAPTER 2

Background

The topic of task-based design of imaging systems and image processing algorithms is of interest in any image formation application where a system or algorithm is needed to produce images that will be used in a prespecified objective task. However, the motivation for this investigation comes specifically from the problem of designing regularized reconstruction methods for emission tomography, and therefore we begin with a brief review of this problem. We first discuss the imaging systems used for emission tomography and the nature of the data produced by such systems. We then describe some of the reconstruction methods that one can use to form images from this data. Finally, we present a simple mathematical model for some of the classification tasks performed on such images, as well as mathematical models for the observers that perform the tasks.

2.1 Tomographic Imaging Systems

2.1.1 Overview

Tomography[1] is the imaging of transverse cross-sections of a three-dimensional object [87]. In *computed tomography* (CT), an imaging system acquires a set of projections through the three-dimensional object and an image reconstruction algorithm processes

---

[1] From the Greek $\tau\acute{o}\mu o\varsigma$, section, and $\gamma\rho\alpha\phi\acute{\eta}$, writing.

these projections to form images of the transverse cross-sections.

In *emission* CT, the imaged quantity is the spatial distribution of radioactivity throughout the object of interest. The source of this radioactivity is a substance injected into the object. For clinical use, the radioactive substance usually consists of a radionuclide bound to some molecule that can be absorbed by the part of the patient's physiology that is to be imaged. This substance is injected into the patient in a dosage small enough to avoid having unwanted biochemical effects and is called a *radiotracer*. The detector counts photons emitted along different directions by decaying atoms of the radionuclide.

Detectors in emission tomography systems typically consist of scintillating crystals that are coupled to arrays of photo-multiplier tubes (PMTs). Commonly used scintillating materials include thallium-activated sodium iodide (NaI(Tl)) and bismuth germanate (BGO). The γ-ray photons that are emitted by decaying atoms of the radiotracer within the patient's body hit the scintillating crystal on the detector. When the crystal absorbs a γ-ray photon, it releases a burst of light photons onto the photocathodes of the PMTs. In response to these light photons, the PMTs produce electrons in a manner that allows the position of the scintillation event to be estimated. The pulses that are generated at the anodes of the PMTs vary depending on how close the PMTs are to the location on the detector where the γ-ray photon was absorbed. The electronics of the detector count the pulses generated by these scintillation events and record their estimated locations. Thus the measurements form a sort of histogram of the number of photons absorbed at different locations on the surface of the detector.

Roughly speaking, the higher the radioactivity concentration at some region of the body, the more the photons that will hit the region on the detector that it faces, and the more the pulses that will be counted there. Thus the photon counts collected at the detector provide information about radiotracer uptake in different parts of the body. Radiotracer uptake can help distinguish between normal and diseased function, so emission tomography can be used to acquire functional information about the patient's body.

The development of emission tomography has provided safe and non-invasive diagnostic capabilities for a variety of medical conditions where early and accurate diagnosis is critical. Applications include detecting breast, lung or other tumors, diagnosing myocardial wall movement defects, and imaging brain function in epilepsy or stroke patients [20]. Tomographic images are used to monitor the treatment of specific patients, as well as to evaluate the effectiveness of new procedures and medications.

Emission tomography encompasses two families of imaging modalities, PET and SPECT. We proceed with discussing the basic characteristics of the two.

2.1.2 Positron Emission Tomography

In *positron emission tomography* (PET), positron-emitting radionuclides are used. When atoms of such a radionuclide decay, positrons are emitted. Each positron travels for a short distance before being annihilated by colliding with an electron. The annihilation results in the emission of two γ-ray photons in directions that form a 180^o angle. The detector forms a ring, with the patient lying in its interior. It comprises an array of scintillating crystals, either larger continuous crystals (typically made of NaI), or small discrete crystal blocks (typically made of BGO) [81]. A simplified diagram of a PET detector is shown in Fig. 2.1(a).

The coincidence of two photons arriving simultaneously at two different positions on the detector ring is an event of interest, recorded at the detector. Ideally, this coincidence signifies that a positron annihilation occurred somewhere along the *line of response* (LOR), *i.e.*, the line segment defined by the positions where the two photons hit the detector. The higher the radioactivity distribution within the body along a certain LOR, the more the events recorded for this LOR at the detector. Specifically, the count of events for each LOR can be viewed as a projection (integration) of the values of the radioactivity distribution along that line, as shown in Fig. 2.3. These projections can subsequently be used to reconstruct an image of the radionuclide distribution.

However, due to effects known as *photon attenuation* and *scattering*, departures from this ideal situation often occur. That is, not every positron annihilation results in a coincidence event recorded at the detector, and not every event recorded at the detector corresponds to a true annihilation along the respective LOR.

On the one hand, a significant number of photons are absorbed as they travel through their surrounding materials. As a result, the number of photons reaching the detector is only a fraction of the number of photons emitted from within the body. This effect is called attenuation and is an exponential function of the distance traveled and of the attenuation coefficient of the surrounding materials. If attenuation is not accounted for, LORs through highly attenuating materials, such as bone, will appear to generate less annihilation events than they actually do.

As a secondary result of attenuation, it is possible that one of the two photons produced by a positron annihilation is absorbed and only one stray photon reaches the detector. If two such stray photons from two different positron annihilations happen to reach the detector simultaneously, this event, called a *random coincidence*, will be erroneously recorded at the detector as if it were a true coincidence corresponding to a single annihilation.

On the other hand, some photons experience *Compton scattering* as they interact with free or bound electrons of the materials they travel through. In its interaction with an electron, the scattered photon loses part of its energy and its path is deflected by some angle inversely related to the amount of energy lost. Scintillating detectors can distinguish scattered photons based on their reduced energy but they have a limited energy resolution

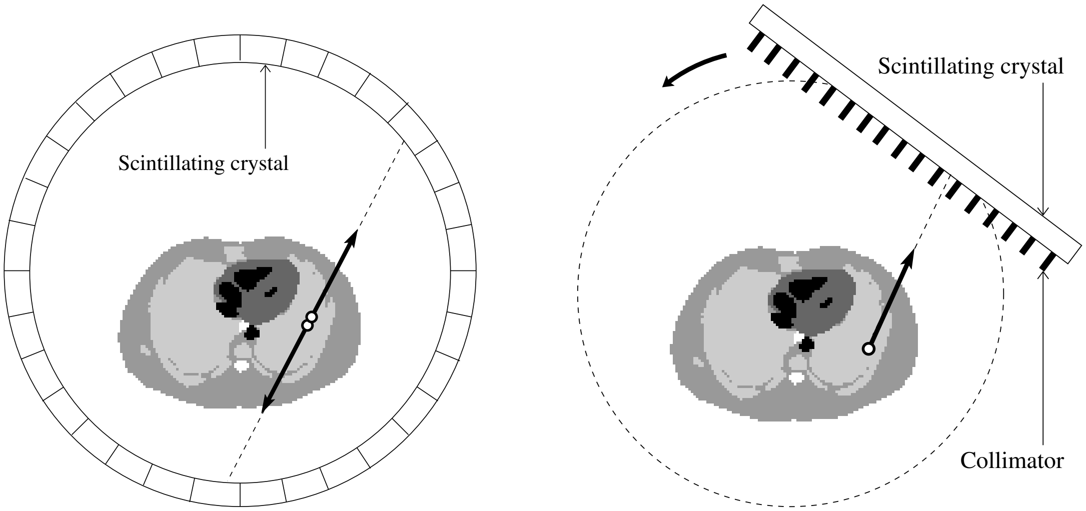

(a) PET (b) SPECT

Figure 2.1: Simplified diagrams of a PET and a SPECT system.

and therefore some scattered photons will still be counted. A *scattered coincidence*, although it involves two photons generated from the same positron annihilation, results in an erroneous LOR being recorded due to the deflection of the photon path.

2.1.3 Single-Photon Emission Computed Tomography

In *single-photon emission computed tomography* (SPECT), γ-emitting radionuclides are used. When an atom of such a radionuclide decays, a single (usually) γ-ray photon is emitted isotropically. The typical SPECT detector, known as an Anger camera, consists of a single large NaI crystal and rotates around the patient [153]. The orbit of the rotating detector can be arbitrary but in practice it usually is circular or elliptical. As the detector rotates at different angles along this orbit, it collects counts of γ-ray photons emitted from the patient's body at each angle. Similarly to PET, these counts are projections through the nuclide distribution inside the body and can be used to reconstruct an image of this distribution. A simplified diagram of a SPECT system is shown in Fig. 2.1(b).

Photons are emitted isotropically by decaying nuclei throughout the body, some in the direction of the detector. Due to attenuation, similarly to PET, not all photons travelling towards the SPECT detector will actually reach it. An important component of the SPECT detector is a lead collimator that is mounted on the scintillating crystal, facing the patient. Among the photons reaching the detector, those moving in appropriate directions will travel down the openings of the lead collimator. The collimator openings are usually parallel to each other, although geometries with convergent or divergent openings also exist.

The purpose of the collimator is to achieve spatial localization of the photon emissions, since in SPECT, with isotropically emitted single photons, it would otherwise be impossible to get LOR information for scintillation events. Each opening of the collimator is penetrated only by photons originating from a small range of angles around the direction of the collimator bore. The narrower the bore, the better the LOR resolution but also the smaller the number of photons that reach the scintillator and therefore the lower the detector sensitivity.

An important complicating characteristic of SPECT systems is their shift variance. The SPECT system response varies as a function of distance from the detector and as a result not all areas of the imaged distribution contribute to the measurements in the same way. Areas that are further away from the detector experience more blur than those that are closer. This depth-dependent blurring is illustrated in Fig. 2.2. A photon entering at a particular opening of the collimator could have originated in a greater range of locations within the body if it originated at distance d_2 than at distance $d_1 < d_2$. As a result, areas of the object that are closer to the detector “see” a finer resolution than those further away. The response of the system is thus characterized by a depth-dependent point-spread function (PSF) whose FWHM increases linearly with distance.

Photons hitting on the septa of the collimator are ideally absorbed. In practice, al-

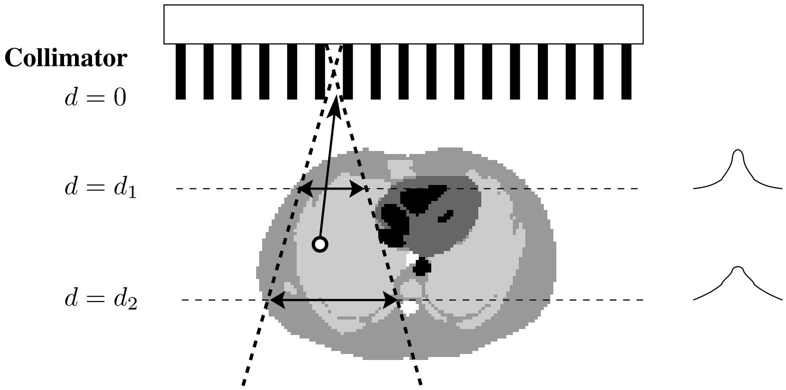

Figure 2.2: Depth-dependent blurring in SPECT.

though lead greatly attenuates γ-rays, some photons will still penetrate the septa and end up in adjacent openings, thus introducing another source of blur in the measurements. Furthermore, similarly to PET, some photons will end up in the wrong collimator opening due to scattering.

2.1.4 System Model

Let $\boldsymbol{f}$ be the object to be imaged and $\boldsymbol{p}$ a measurement of the object obtained through an imaging system. We model the imaging system as a linear operator $\boldsymbol{\mathcal{A}}$, which we will refer to as the *system operator*, mapping the object space into the measurement space[2]. We assume that the object can be written as a function $f : \mathbb{R}^2 \to \mathbb{R}$, and the measurement can be written as a function $p : \mathbb{R} \times [0, \pi] \to \mathbb{R}$. In its most general form, the system operator is defined by the superposition integral

$$\boldsymbol{p} = \boldsymbol{\mathcal{A}} \boldsymbol{f} \iff p(r, \varphi) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} a(r, \varphi; x, y)\, f(x, y)\, \mathrm{d}x\, \mathrm{d}y\,. \tag{2.1}$$

The operator $\boldsymbol{\mathcal{A}}$ captures only the deterministic properties of the system, so $\boldsymbol{p}$ would be the measurement obtained from $\boldsymbol{f}$ in the absence of random noise. The model in (2.1) is

---

[2] Emission tomography system models may be 3-D. However, we focus here on the 2-D case for simplicity and because the observer models that we will deal with in subsequent sections act on 2-D reconstructed images.

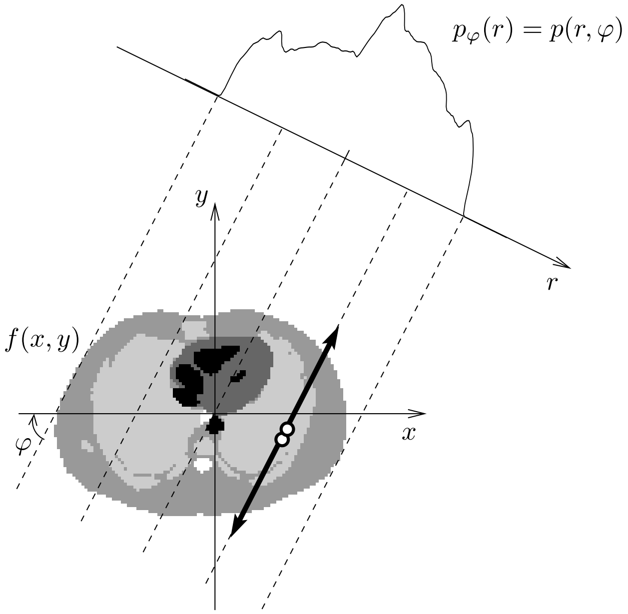

Figure 2.3: The Radon transform illustrated for a single projection angle.

broad enough to encompass various imaging systems, including the emission tomography modalities described in the previous sections. We now elaborate on the nature of $\boldsymbol{\mathcal{A}}$ for PET and SPECT systems.

In emission tomography, the object $\boldsymbol{f}$ is the unknown radioactivity distribution and the measurement $\boldsymbol{p}$ is a set of projections through the object at various angles. The ideal mapping from object space to projection space (in the absence of attenuation and blurring) is known as the *Radon transform*. It can be modeled as a linear operator $\boldsymbol{\mathcal{P}}$ defined by

$$\boldsymbol{p} = \boldsymbol{\mathcal{P}}\boldsymbol{f} \iff p(r,\varphi) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \delta(r - [x\cos\varphi + y\sin\varphi])\, f(x,y)\,\mathrm{d}x\,\mathrm{d}y, \tag{2.2}$$

where $\delta(\cdot)$ is the 1-D Dirac impulse. Here φ is the projection angle (the angle of the LOR in PET or the detector's rotation angle in SPECT) and r is the displacement in the perpendicular direction. The Radon transform is illustrated in Fig. 2.3.

The Radon transform maps a single point in object space to a sinusoidal line in projection space, which is why a set of tomographic data is also called a *sinogram*. Fig. 2.4 shows an example of tomographic data. When the true object is the chest phantom shown

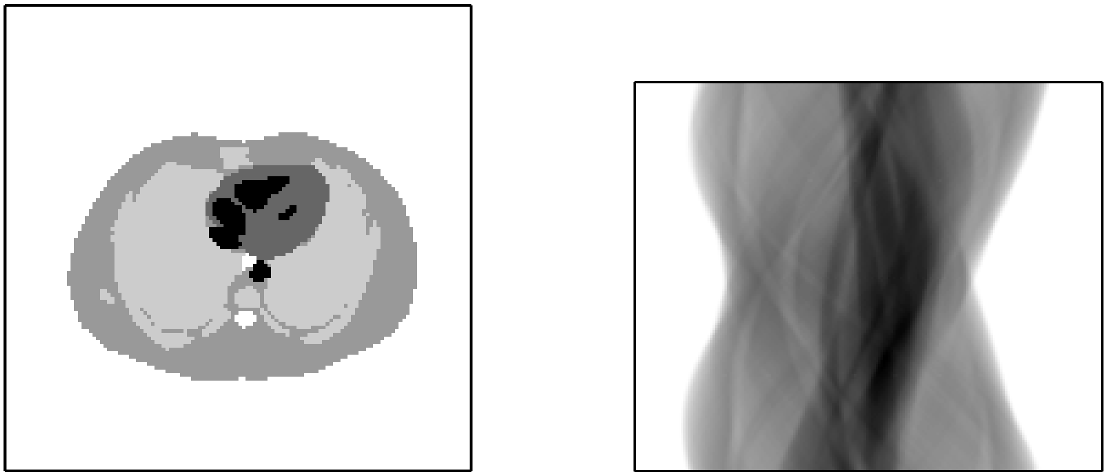

(a) True object $\boldsymbol{f}$

(b) Sinogram $\boldsymbol{p}$

Figure 2.4: Example of an object and its sinogram as produced by a PET system model. The model accounts for non-uniform attenuation in the object.

in Fig. 2.4(a), then the projections through this object, as they would be acquired by a PET system in the absence of imaging noise, would look like Fig. 2.4(b).

The ideal Radon transform in (2.2) involves line integrals of the object. However, real tomographic systems have a finite spatial resolution, which translates to blur introduced in these integrals. When this blur and other effects of real systems, such photon attenuation, are accounted for, the system model departs from (2.2). Let $b_\varphi(r; x, y)$ be the system blur function at position (x, y) within the object along a LOR of angle φ, given as a function of the radial displacement r. Let also $c_\varphi(x, y)$ be the attenuation factor experienced by photons emitted at position (x, y) and traveling towards the detector along the direction of φ. Then the tomographic system can be modeled as an operator $\boldsymbol{\mathcal{P}}_{b,c}$ such that

$$\begin{aligned} \boldsymbol{p} &= \boldsymbol{\mathcal{P}}_{b,c}\boldsymbol{f} \iff \\ p(r,\varphi) &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} b_\varphi(r - [x\cos\varphi + y\sin\varphi]; x, y)c_\varphi(x,y)\, f(x,y)\,\mathrm{d}x\,\mathrm{d}y\,. \end{aligned} \tag{2.3}$$

In PET, the system blur is sometimes assumed shift-invariant, *i.e.*, $b_\varphi(r; x, y) = b_\varphi(r)$

$\forall x, y$. Then (2.3) can be written as a convolution in the measurement domain. Furthermore, photon pairs that are emitted along the same LOR experience the same attenuation factor in PET, regardless of where on this LOR they originated, *i.e.*, $c_\varphi(x, y) = c_\varphi(x \cos\varphi + y \sin\varphi)$. Thus the attenuation factor in (2.3) can be factored out as a LOR-dependent scaling $c_\varphi(r)$. In this case the system operator $\boldsymbol{\mathcal{A}}$ for PET can be written as

$$\boldsymbol{\mathcal{A}}_{\text{PET}} = \boldsymbol{\mathcal{B}}_{\text{PET}} \boldsymbol{\mathcal{C}}_{\text{PET}} \boldsymbol{\mathcal{P}},$$

where $\boldsymbol{\mathcal{B}}_{\text{PET}}$ and $\boldsymbol{\mathcal{C}}_{\text{PET}}$ are measurement-domain operators denoting convolution with $b_\varphi(r)$ and multiplication with $c_\varphi(r)$ respectively. The above factorization does not hold when blur due to depth-of-interaction effects is taken into account, as in [79, 112]. Then the system blur becomes shift-variant and the more general formulation of (2.3) applies.

In SPECT, the attenuation factor experienced by a photon depends not only on the LOR (r, φ) along which it travels, but also on the position (x, y) where the photon originated. Therefore the effects of attenuation cannot be factored out of the Radon transform. Furthermore, the depth-dependent blur in SPECT systems is strongly dependent on the position (x, y). Thus for SPECT system models only the general formulation applies:

$$\boldsymbol{\mathcal{A}}_{\text{SPECT}} = \boldsymbol{\mathcal{P}}_{b,c}.$$

A tomographic system model can be more complex, accounting for the effects of scatter, detector efficiency, etc. Note that although $\boldsymbol{\mathcal{A}}$ is referred to as a “system” operator, here it depends on properties of both the tomographic imaging system (*e.g.*, geometry and detector efficiency) and the imaged volume (*e.g.*, attenuation and scatter).

Incorporating attenuation into the model of a PET or SPECT system requires knowledge of how γ-rays are attenuated along each LOR through the patient. It is becoming increasingly common to acquire this information by a transmission CT scan of the patient [71, 81]. In *transmission tomography*, an external x-ray source of known intensity is scanned around the patient. Measuring the intensity of the x-rays as they exit the patient's body from different angles can yield a map of the attenuation coefficient of the tissues throughout the body. This map is then used to calculate the attenuation factors $c(r, \varphi)$ corresponding to each LOR.

Incorporating scatter into the model of a tomographic system requires knowledge of the fraction of the measurement at each LOR that corresponds to scattered counts. There are several approaches to estimating this fraction and incorporating it into the model [71, 81]. The most commonly used method involves acquiring separate count measurements in nearby energy windows simultaneously with the main energy window. The ratio of scattered counts in the main energy window is then estimated as a fraction of the counts in the nearby energy windows. This is the basis for the popular triple energy window (TEW) scatter correction method [102]. Somewhat similarly, the effect of random coincidences in PET can be quantified by counting photons in a delayed time window, in addition to the main time window used to count coincidence events.

Incorporating a depth-dependent response into the model of a SPECT system requires knowledge of the PSFs of the system at different distances from the detector face. This information can be obtained by imaging a "point source" of radioactivity at different distances from the detector [15, 93, 127]. Ignoring septal penetration, the PSFs of a SPECT system with a parallel-hole collimator are usually modeled as Gaussian-shaped functions whose FWHM increases linearly with distance [21]. For septal penetration modeling, response functions with heavier tails are appropriate [73].

In all the above representations, the parameters x, y, r and φ were assumed to vary over a continuum of values. In practice, however, tomographic systems can acquire only a finite number of measurements and reconstruction algorithms can estimate only a finite number of image intensity parameters. Thus discrete representations are also useful.

In a discrete representation of the measurements, the detector collects n_{d} photon counts, which are lexicographically ordered in a vector $\boldsymbol{y} \in \mathbb{R}^{n_{\mathrm{d}}}$. In the 2-D case described above, the angle φ takes a discrete set of values φ_j, $j = 1, \ldots, n_\varphi$, corresponding to n_φ possible LOR angles in PET or n_φ rotation angles along the detector orbit in SPECT. At each angle, the detector collects photon counts at a finite number of radial positions, so r takes a discrete set of values r_i, $i = 1, \ldots, n_{\mathrm{b}}$, resulting in a total of $n_{\mathrm{d}} = n_{\mathrm{b}} n_\varphi$ recorded counts. The conceptual detector unit responsible for collecting one of these n_{d} counts is referred to as a *detector element* or *bin,* although it does not necessarily correspond to a physical detector element.

In a discrete representation of the object $\boldsymbol{f}$, a set of basis functions $v_j(x, y)$, $j = 1, \ldots, n_{\mathrm{p}}$, is chosen and the imaged distribution is assumed to be well approximated by a linear combination of these functions:

$$f(x, y) \approx \sum_{j=1}^{n_{\mathrm{p}}} f_j v_j(x, y). \tag{2.4}$$

The expansion coefficients f_j are then lexicographically ordered in a vector $\boldsymbol{f} \in \mathbb{R}^{n_{\mathrm{p}}}$. Typically the basis functions are spatially localized volume elements, arranged on a grid over the imaged area. The usual choice is cubic volume elements, but other shapes have been studied as well [82, 91, 138, 150].

For any of the system models discussed above, the problem at hand, in the absence of random noise, would be to find the object $\boldsymbol{f}$ given the measurement $\boldsymbol{p}$, where

$$\boldsymbol{p} = \boldsymbol{\mathcal{A}} \boldsymbol{f}. \tag{2.5}$$

However, even for this hypothetical noiseless measurement, it is not trivial to solve (2.5) for $\boldsymbol{f}$. For a measurement $\boldsymbol{p}$ in discrete space and an object $\boldsymbol{f}$ in continuous space, the problem is ill-posed. When adopting a discrete representation for both $\boldsymbol{p}$ and $\boldsymbol{f}$, the system operator becomes a matrix $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{n_{\mathrm{d}} \times n_{\mathrm{p}}}$. The ij element of this matrix, given by

$$a_{ij} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} a(r_i, \varphi_i; x, y) v_j(x, y) \, \mathrm{d}x \, \mathrm{d}y, \quad i = 1, \ldots, n_{\mathrm{d}}, \quad j = 1, \ldots, n_{\mathrm{p}},$$

represents the probability that a photon emitted at the jth volume element will be recorded at the ith detector element. Most of these probabilities are zero, resulting in a very sparse matrix $\boldsymbol{\mathcal{A}}$. The nature of this matrix in emission tomography is such that, even if it is full-rank, it usually is ill-conditioned, so inverting it is not trivial.

2.1.5 Measurement Model

The raw measurements acquired by a real tomographic imaging system are not related deterministically to the true radioactivity distribution, as (2.5) would suggest. The inherent randomness in radiotracer distribution, decay time, and photon detection result in randomness in the measurements, which can be treated as Poisson-distributed independent random variables [111, 152].

More specifically, if $\boldsymbol{Y}$ is a random vector representing the acquired tomographic data, then

$$\boldsymbol{Y} \sim \mathsf{Poisson}\{\boldsymbol{\mathcal{A}} \boldsymbol{f} + \boldsymbol{r}\}, \tag{2.6}$$

where $\mathsf{Poisson}\{\boldsymbol{\mu}\}$ denotes a multivariate Poisson probability distribution with mean vector $\boldsymbol{\mu}$. The vector $\boldsymbol{r} \in \mathbb{R}^{n_{\mathrm{d}}}$ contains the “background” counts (*e.g.*, random coincidences and/or scattered counts) collected at each detector element in lexicographic ordering. (Although it is possible to account for scatter within the system model $\boldsymbol{\mathcal{A}}$ [14, 103, 139], we assume here that $\boldsymbol{\mathcal{A}}$ models only unscattered events.)

Both $\boldsymbol{\mathcal{A}}$ and $\boldsymbol{r}$ are considered known, following the discussion in the previous section. The image reconstruction problem in emission tomography consists of estimating $\boldsymbol{f}$ given the data $\boldsymbol{y}$, where $\boldsymbol{y}$ is a realization of the random vector $\boldsymbol{Y}$ in (2.6).

2.2 Tomographic Image Reconstruction

2.2.1 Overview

As previously explained, the tomographic data $\boldsymbol{y}$ are in the form of projections through the true object $\boldsymbol{f}$ along different directions. This data cannot be used directly for the purpose of medical diagnosis. It would be extremely challenging for a human observer to reliably decide on the presence of a small anomaly and determine its exact position within the imaged area by looking at the projection views. It would also be very challenging to quantify radiotracer uptake within a region of interest in the object directly from the projection view data. This creates the need for a *reconstruction method, i.e.*, a mapping of the projection data $\boldsymbol{y}$ to an estimated image $\hat{\boldsymbol{f}}$.

The resolution of a particular imaging system is intrinsically limited by its various physical characteristics. Nevertheless, the resolution of the reconstructed images can be significantly poorer than this intrinsic resolution, as it is also affected by the reconstruction method that is applied to the data. This effect is undesirable, since important details that we seek to image, such as small lesions, may be missed if the resolution is not fine enough. On the other hand, the reduction of noise in the reconstructed images is also important, since noise artifacts may be mistaken for lesions. However, fine resolution and low noise levels are competing demands and we can only accomplish a trade-off between them.

For a broad review of image reconstruction methods for emission tomography, see, for example, [84]. In the following we list some common reconstruction methods, explaining the tuning parameters through which each of them controls the noise-resolution trade-off. We begin with backprojection, which, although not adequate as a reconstruction method by itself, plays a role in most practical reconstruction methods.

2.2.2 Backprojection

In the following, we use "′" to denote the adjoint of an operator or equivalently the complex transpose of a matrix. For the generic linear operator $\boldsymbol{\mathcal{A}}$ defined in (2.1), the adjoint operator $\boldsymbol{\mathcal{A}}'$ is defined by

$$\boldsymbol{g} = \boldsymbol{\mathcal{A}}'\boldsymbol{q} \iff g(x,y) = \int_0^\pi \int_{-\infty}^\infty a^*(r,\varphi;x,y)\,q(r,\varphi)\,\mathrm{d}r\,\mathrm{d}\varphi\,.$$

For the Radon transform operator in (2.2), the operation $\boldsymbol{\mathcal{P}}\boldsymbol{f}$ is a *(forward) projection* of the true object $\boldsymbol{f}$ and the adjoint operation,

$$\boldsymbol{g} = \boldsymbol{\mathcal{P}}'\boldsymbol{q} \iff g(x,y) = \int_0^\pi q(x\cos\varphi + y\sin\varphi, \varphi)\,\mathrm{d}\varphi, \tag{2.7}$$

is a *backprojection* of the sinogram $\boldsymbol{q}$. The latter amounts to replicating each projection onto image space along the direction of the respective projection angle and integrating the result over all angles.

We also define here a "diagonal" weighting operator $\boldsymbol{\mathcal{W}}$, performing sinogram-domain multiplication with some angle-dependent weights $w(\varphi)$, by

$$\boldsymbol{q} = \boldsymbol{\mathcal{W}}\boldsymbol{p} \iff q(r,\varphi) = w(\varphi)\,p(r,\varphi), \tag{2.8}$$

for some π-periodic angular-weighting function $w(\varphi)$. We assume that $w(\varphi)$ is real, in which case the operator is self-adjoint, *i.e.*, $\boldsymbol{\mathcal{W}}' = \boldsymbol{\mathcal{W}}$.

Although backprojection, as defined in (2.7), is a mapping from the sinogram domain to the image domain, it is not appropriate as a reconstruction method because it results in severe blurring. Using the definition of the Radon transform operator $\boldsymbol{\mathcal{P}}$ from (2.2), one can show that, for any weighting operator $\boldsymbol{\mathcal{W}}$ defined as in (2.8), we have [37]

$$\boldsymbol{\mathcal{P}}'\boldsymbol{\mathcal{W}}\boldsymbol{\mathcal{P}}\boldsymbol{f} = \frac{w(\varphi + \pi/2)}{|r|} ** \boldsymbol{f}, \tag{2.9}$$

where $r = \sqrt{x^2 + y^2}$ and the symbol "$**$" denotes 2-D convolution. That is, projecting an object onto the sinogram domain and subsequently backprojecting it onto the image domain, is equivalent to convolving the object with $1/|r|$, a particularly heavy-tailed PSF.

2.2.3 Analytical Reconstruction

Analytical reconstruction methods for estimating $\hat{\boldsymbol{f}}$ from $\boldsymbol{y}$ rely on the ideal mathematical model for the projection data, ignoring its statistics.

If one disregards the fact that real data are discrete and assumes that $\boldsymbol{\mathcal{A}}$ is a continuous-to-continuous operator as in (2.1), then it is possible to invert $\boldsymbol{\mathcal{A}}$ for some simple tomographic system models [99]. This is the premise of the popular *filtered backprojection* (FBP) method, which relies on the property (2.9) of the Radon transform operator. The $1/|r|$ blur in a backprojected image can be removed by applying a cone filter on the image or, equivalently due to the central-section theorem [87, p.117], by applying a ramp filter on the projections before backprojecting.

Since the ramp filter would greatly amplify the noise present in real data, practical implementations of FBP use an apodized ramp filter [26], essentially leaving some of the blur in the image to trade off some of the noise. There are many choices for the apodization window and its cut-off frequency is the parameter that controls the variance/resolution trade-off in FBP-reconstructed images.

Incorporating a detailed system model into FBP is not straightforward, especially for SPECT, since the ideal model used to derive (2.9) does not account for shift-variant effects. The exact inversion of the Radon transform has been extended to accommodate non-uniform attenuation [100, 101] but not depth-dependent detector response. Numerous variations of FBP exist that attempt to approximately model shift-variant effects. They range from the the simple but popular Chang attenuation correction [25] to more sophisticated methods, compensating for both uniform attenuation and depth-dependent PSFs [67, 123, 131].

Another family of analytical reconstruction methods, known as algebraic reconstruction techniques (ARTs), attempt to obtain $\boldsymbol{f}$ by solving the forward projection equation (2.5). Because of the size of the problem in typical tomographic applications, iterative algorithms are used to find the solution [24]. Several of the strategies for controlling the variance/resolution trade-off that are described in the following section, such as post-filtering or regularization, can be used with ARTs as well.

The main advantage of analytical reconstruction methods over the statistical methods described in the following section is that the former may be much faster. Also, their variance and resolution properties, especially in the case of FBP, are easier to analyze and control. However, analytical methods may suffer from higher bias due to their difficulty in incorporating accurate system models and/or higher variance due to their disregard for measurement statistics.

2.2.4 Statistical Reconstruction

In general, statistical reconstruction methods estimate $\boldsymbol{f}$ from $\boldsymbol{y}$ by maximizing an *objective function* $\Psi(\boldsymbol{f})$ of the form

$$\Psi(\boldsymbol{f}) = D(\boldsymbol{y}, \boldsymbol{\mathcal{A}}\boldsymbol{f}) - \beta R(\boldsymbol{f}) \tag{2.10}$$

$$\hat{\boldsymbol{f}} = \arg\max_{\boldsymbol{f} \geq \boldsymbol{0}} \Psi(\boldsymbol{f}), \tag{2.11}$$

where $D(\cdot,\cdot)$ and $R(\cdot)$ are referred to as the *data-fit* and *regularization* terms respectively and $\beta \geq 0$ is a *regularization parameter* that controls the relative weight of the two terms. The non-negativity constraint $\boldsymbol{f} \geq \boldsymbol{0}$ is used in emission tomography, where $\boldsymbol{f}$ consists of photon emission densities.

The specific choices for $D(\cdot,\cdot)$ and $R(\cdot)$ have a great impact on the quality of the reconstructed image $\hat{\boldsymbol{f}}$. For common choices, a closed-form expression for the maximizer in (2.11) either does not exist or is too expensive to compute directly due to the problem size, hence the need for iterative algorithms. Assuming that (2.10) has a unique maximizer and that the algorithm iterates until the estimate $\hat{\boldsymbol{f}}$ converges to this maximizer, the algorithm itself has no effect on the quality of $\hat{\boldsymbol{f}}$.

The usual choice for the data-fit term $D(\cdot,\cdot)$ is the log of the conditional likelihood $p_{\boldsymbol{Y}|\boldsymbol{f}}(\boldsymbol{y}|\boldsymbol{f})$ of the data $\boldsymbol{y}$ given the true object $\boldsymbol{f}$, in which case (2.10) and (2.11) define a *penalized-likelihood* reconstruction method. For emission tomography, the Poisson statistical model in (2.6) yields the log-likelihood

$$D(\boldsymbol{y}, \boldsymbol{\mathcal{A}f}) = \sum_{i=1}^{n_{\mathrm{d}}} y_i \log([\boldsymbol{\mathcal{A}f}]_i + r_i) - ([\boldsymbol{\mathcal{A}f}]_i + r_i), \tag{2.12}$$

where terms that are independent of the true object $\boldsymbol{f}$ have been discarded. Other data-fit terms have been proposed for situations where the data statistics have been altered due to preprocessing. An example is the shifted-Poisson model for PET data that have been corrected by subtracting the randoms [146]. Alternatively, the assumption of Gaussian statistics leads to penalized weighted least-squares (PWLS) methods [33], which make use of a data-fit term that is quadratic in the object $\boldsymbol{f}$:

$$D(\boldsymbol{y}, \boldsymbol{\mathcal{A}f}) = -(\boldsymbol{y} - \boldsymbol{\mathcal{A}f})' \boldsymbol{W} (\boldsymbol{y} - \boldsymbol{\mathcal{A}f}), \tag{2.13}$$

for some measurement-domain weighting matrix $\boldsymbol{W}$.

Setting $\beta = 0$ in (2.10) results in an *unregularized* reconstruction method, which maximizes the data-fit term alone. Due to the noise in emission data and the conditioning of the problem, such methods yield unacceptably noisy images. Therefore some noise reduction strategy is needed and the literature offers numerous alternatives to choose from.

With most iterative algorithms, the estimate becomes increasingly noisy after each iteration [12]. Thus the simplest approach to noise reduction is to initialize the algorithm with a smooth image and to stop it at some early iteration, well before convergence [57, 133]. The number of iterations is then the parameter that controls the noise/resolution trade-off. However, stopping the algorithm before convergence makes the quality of the reconstructed image dependent on additional factors, namely the image used for initialization and the algorithm itself. Several iterative algorithms have been proposed and they often come with their own tuning parameters. For example, practical application of the popular *expectation maximization* (EM) algorithm [118] in tomographic reconstruction has been facilitated by the introduction of *ordered subsets* (OS) [63]. The OS idea consists of accelerating an algorithm by processing only a subset of projection data at each iteration. How these subsets are chosen affects the image that is reconstructed, if the algorithm is not allowed to converge.

Alternatively, one can allow the unregularized algorithm to converge and achieve noise reduction by post-filtering the noisy reconstructed image [16, 121] or even by filtering the estimate after each iteration [65, 120]. In these strategies, it is the smoothing filter that controls the variance/resolution trade-off. Other approaches control this trade-off through some smoothness constraints on the class of permissible true objects, such as the method of *sieves* [122, 134] or methods that substitute smooth volume elements for the traditional cubic voxels [82, 91].

Finally, noise reduction and convergence acceleration can be achieved by the introduction of the regularization term $R(\cdot)$ in (2.10). This term is typically a *roughness penalty*, favoring smooth images by penalizing the intensity differences between neighboring volume elements [46, 56, 78]. Then maximizing (2.10) with $\beta > 0$ corresponds to maximum a posteriori (MAP) estimation with a smoothness prior and the regularization parameter β can be used to adjust the variance/resolution trade-off in the reconstructed images. A smaller β in (2.10) gives more weight to the data-fit term, resulting in a noisier image,

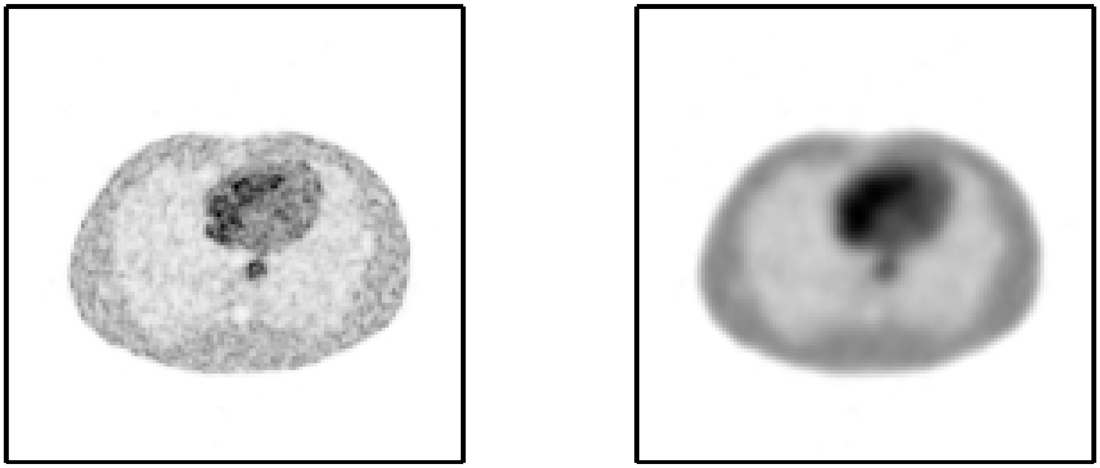

(a) Small β

(b) Large β

Figure 2.5: Examples of penalized-likelihood reconstruction from noisy data. The images have been produced with a small and a large value of the regularization parameter β.

while a larger β gives more weight to the penalty term, resulting in a smoother image. Examples of reconstructions of noisy data using a penalized-likelihood algorithm with a small and a large value of β are shown in Fig. 2.5(a) and Fig. 2.5(b) respectively.

Besides the value of β, the form of the penalty function also affects image quality significantly and various choices exist here as well. Assuming a discrete representation for the reconstructed image as in (2.4), a general form of a roughness penalty is

$$R(\boldsymbol{f}) = \sum_k w_k \, \psi([\boldsymbol{\mathcal{D}} \boldsymbol{f}]_k),$$

where the summation index k indicates the kth neighborhood of volume elements whose intensity differences are to be penalized. The cost function $\psi(t)$, also referred to as a *potential function*, specifies the penalty enforced for an intensity difference equal to t. The penalty weights w_k provide the flexibility to choose a different amount of regularization for each neighborhood, allowing local resolution control. The difference matrix $\boldsymbol{\mathcal{D}}$ is itself a matter of choice (*e.g.*, [80]) but it is usually chosen to produce first-order finite differences of the form

$$[\boldsymbol{\mathcal{D}} \boldsymbol{f}]_k = f_j - f_k,$$

for element j near element k. Typically, differences between each element and its nearest neighbors (or maybe its second-nearest neighbors as well) are included [56].

Various alternatives have been proposed for the cost function $\psi(t)$ [76, 78]. A usual choice is the quadratic

$$\psi(t) = t^2/2 \qquad \Longleftrightarrow \qquad R(\boldsymbol{f}) = \boldsymbol{f}' \boldsymbol{\mathcal{R}} \boldsymbol{f}, \tag{2.14}$$

where $\boldsymbol{\mathcal{R}} = \frac{1}{2} \boldsymbol{\mathcal{D}}' \boldsymbol{\mathcal{D}}$. It is very simple to optimize but it penalizes large intensity differences excessively, resulting in blurred object edges. For this reason, edge-preserving cost functions have been introduced, but these are more complicated to optimize and analyze. Moreover, some of them involve additional tuning parameters that control the degree of edge preservation. For example, the Huber penalty [62],

$$\psi(t) = \begin{cases} t^2/2, & |t| \le \delta \\ \delta|t| - t^2, & |t| > \delta, \end{cases}$$

is quadratic only for small values of t, with a parameter δ controlling how small.

In addition to the form of the data-fit and regularization functionals in (2.10), yet another design consideration that affects the quality of the reconstructed images is the accuracy of the system model $\boldsymbol{\mathcal{A}}$. As typical iterative algorithms apply both the forward projector $\boldsymbol{\mathcal{A}}$ and the backprojector $\boldsymbol{\mathcal{A}}'$ once per iteration, the increased computation required by a more complicated model needs to be weighed against its impact on image quality. For example, a simpler backprojector $\boldsymbol{\mathcal{B}} \neq \boldsymbol{\mathcal{A}}'$ is sometimes substituted for the adjoint in the interest of speed, but the errors due to this mismatch may accumulate, resulting in artifacts [155]. Appendix A compares how two different implementations of $\boldsymbol{\mathcal{A}}$ for 3-D SPECT affect reconstructed image quality.

As illustrated by this brief review, there is a plethora of methods for image reconstruction from tomographic data and each method involves various parameters that affect the quality of the reconstructed images. Thus one needs to establish a criterion for ranking these methods and choosing optimal values for their tuning parameters. Preferably this criterion should be related to some intuitive measure of image quality. The following section lists some approaches in the literature, with a focus on criteria that have been used to optimize the regularization parameters involved in penalized-likelihood reconstruction methods.

2.2.5 Methods for Regularization Design

Several general-purpose strategies exist for choosing regularization parameters in inverse problem solutions that involve maximization of an objective function similar to (2.10). One such strategy is the method of L-curves, which chooses the regularization parameter to achieve a balance between the magnitudes of the data-fit and regularization terms [17, 55]. This strategy has undesirable issues of non-convergence as the noise level goes to zero [135]. Other strategies choose the amount of regularization that minimizes some type of mean squared error (MSE) [66]. Yet others use a Bayesian framework, introducing a prior distribution on the regularization parameter [60, 115, 156].

General-purpose methods such as the ones listed above do not relate the choice of regularization to any intuitive measure of image quality that is of interest in medical imaging. As data-fit and regularization terms typically used in image reconstruction have unrelated units, balancing their magnitudes via the L-curve method is not reasonable, nor does it have any obvious, desirable effect on reconstructed image quality. How specific choices of MSE-related metrics or priors on the regularization parameter affect image quality is not obvious, either.

Several methods for regularization design that are more specifically tailored to image reconstruction exist in the literature. One approach is choosing the regularizer that leads to some desirable property, such as uniform and isotropic resolution [40, 95, 124] or optimal contrast-to-noise ratio [108], in the reconstructed images. Another approach is choosing the regularizer that leads to optimal performance in a task to be performed on the reconstructed images. Tasks of interest in a clinical setting can be split into the broad categories of estimation and classification. As far as estimation (quantification) performance is concerned, one could choose the regularizer so as to approach the uniform Cramér-Rao lower bound on the variance of biased estimators [59].

This thesis is concerned with optimizing image reconstruction in general, and the regularizer in penalized-likelihood methods in particular, with respect to classification performance. The following section describes this problem in more detail and presents related previous work.

2.3 Detection of Signals in Reconstructed Images

2.3.1 Overview

Classification is the task of assigning an object or event to one of a set of prespecified classes. This is performed by observing one or more *features* of the object or event and deciding which class it belongs to through the evaluation of those features [31]. Medical diagnosis tasks can often be described as classification problems, *e.g.*, classifying a patient as normal or diseased, classifying a condition as mild, moderate or acute, etc. Our focus here is the case where the features used for the purpose of such a classification are extracted from a set of tomographic data.

We refer to the entity performing the classification as the *observer*. In a clinical setting that would usually be a human, although it may also be a computer program or a combination of the two. However, our ultimate goal in studying observers here is a fast, automated means for adjusting the parameters of image reconstruction methods prior to the reconstruction so as to improve performance in classification tasks. Therefore, we will eventually turn our attention to mathematical observer models, which are simpler to study analytically. We will be interested especially in certain models that, according to the literature, have been found to perform in a manner correlated to human observer performance.

2.3.2 Classification Tasks

Let $\boldsymbol{f}$ be the true object being imaged (or its approximation in $\mathbb{R}^{n_\mathrm{p}}$). To express our uncertainty about the object, we allow it to be a realization of a random field. In emission tomography this uncertainty stems from the variability in patient physiology and radiotracer uptake. Assume that $\boldsymbol{f}$ belongs to exactly one of the n_C object classes $C_i, i = 0, \ldots, n_C - 1$ and H_i is the hypothesis that $\boldsymbol{f}$ belongs to class C_i. We consider an observer that has to determine which one of the hypotheses $H_i, i = 0, \ldots, n_C - 1$ is true.

In the applications that interest us, the true object $\boldsymbol{f}$ cannot be observed directly, so the decision must be based on the observed data $\boldsymbol{y}$ or, more generally, on a feature vector $\boldsymbol{v} = \boldsymbol{v}(\boldsymbol{y})$ extracted from the data. This task would be easy to perform if there were a deterministic, one-to-one mapping of the object $\boldsymbol{f}$ to the feature vector $\boldsymbol{v}$. However, the imaging process that produces $\boldsymbol{y}$ is characterized by loss of information (*e.g.*, from a continuous-space object to discrete-space data) and randomness. Furthermore, the process through which the feature vector $\boldsymbol{v}$ is extracted from the data $\boldsymbol{y}$ may introduce additional randomness and loss of information.

For a specific instance of the true object $\boldsymbol{f}$, the variability in the data $\boldsymbol{y}$ due to noise inherent in the imaging process is expressed by the conditional distribution $p(\boldsymbol{y}|\boldsymbol{f})$, which in emission tomography is Poisson as given in (2.6). For a specific instance of the data $\boldsymbol{y}$, the variability in the feature vector $\boldsymbol{v}$ due to uncertainty in the observer's feature extraction process is discussed later in this section.

An observer is characterized by a *decision rule, i.e.*, a mapping of the feature vector $\boldsymbol{v}$ to a hypothesis H_i. Observers that have previous knowledge about any aspects (deterministic or random) of the imaging and feature extraction process can perform the classification more effectively by incorporating such knowledge in their decision rule. The decision rules we focus on are deterministic, that is the observer makes the same decision every time it is provided with the same $\boldsymbol{v}$. Furthermore, the decision must specify exactly one of the H_i's to be true. We use D_i to denote the event that the observer decides that H_i is true. Classification then corresponds to a partition of the feature vector space into non-overlapping regions, each corresponding to one of the H_i's [31,44].

Under these assumptions, the observer's decision rule can be represented as a comparison of a set of *test statistics* t_i, each of which is a function of the feature vector $\boldsymbol{v}$, $t_i = t_i(\boldsymbol{v}), i = 0, \ldots, n_C - 1$. The functionals $t_i(\cdot)$ are called *discriminant functions*. The decision rule is then:

$$\text{Decide } H_i \text{ if } t_i(\boldsymbol{v}) > t_j(\boldsymbol{v}) \ \forall j \neq i. \tag{2.15}$$

In the special case of a binary-hypothesis task ($n_C = 2$), this reduces to comparing a single test statistic $t = t(\boldsymbol{v})$ to some *threshold* τ, where τ is independent of the data $\boldsymbol{y}$:

$$\text{Decide } H_1 \text{ if } t(\boldsymbol{v}) > \tau, \text{ otherwise decide } H_0. \tag{2.16}$$

In the following we focus on *detection tasks*, where the observer must determine whether a certain detail of interest is present or not in the true object. Following convention, we use the term *signal* (or *target signal*) to refer to this detail of interest, which may or may not be present in $\boldsymbol{f}$, and not to the entire object $\boldsymbol{f}$. We use the term *background* to refer to what remains of the object in the absence of the signal. In emission tomography, the background would be a normal radionuclide distribution within the patient and a signal of interest could be, for instance, a lesion or a myocardial perfusion defect. Such signals of interest are identified in an emission scan as regions of either higher or lower radionuclide concentration than the normal region of the organ that is imaged. In this case, an additive model for the signal and background is an intuitive one [11, §8.4.5] and we adopt this model hereafter.

In addition to detecting the presence of the signal, it may also be of interest to assign it to one of a set of signal classes. In its most general form, the detection task is thus a decision between hypotheses of the form:

$$\begin{aligned} H_0 &: \quad \boldsymbol{f} = \boldsymbol{f}_b \\ H_i &: \quad \boldsymbol{f} = \boldsymbol{f}_b + \boldsymbol{f}_{s(i)}, \qquad\qquad i = 1, \ldots, n_C - 1. \end{aligned}$$

The usual formulation encountered in the literature assumes that the target signal is spatially localized and assigns a separate signal class to each of the $n_{\mathrm{L}} = n_C - 1$ different locations within the object where the signal may be centered. In this case, the task is signal detection and *localization.* In the most general formulation, the background $\boldsymbol{f}_b$ is treated as a random field. The signal $\boldsymbol{f}_{s(i)}$ at each different location may also be treated as a random field, to express uncertainty about its profile in addition to its location.

In the more realistic scenario where there are multiple candidate locations for the signal, the observer's decision rule has the general form (2.15). Although this formulation assumes a finite set of possible locations, it can be extended to consider a continuum of locations over $\mathbb{R}^2$ (e.g., [68]). Then the task consists of determining whether the signal is absent or present and, in the latter case, estimating its location $\boldsymbol{r} \in \mathbb{R}^2$. In the following we will focus on observers that consider finite sets of candidate locations, since the observer models that are usually studied in the literature fall under this category.

In the special case where there is only one possible signal location, or where identifying the signal location is not of interest, the task becomes a decision between the *signal absent* hypothesis H_0 and the *signal present* hypothesis H_1. Then the binary-hypothesis rule (2.16) applies.

Although in realistic diagnostic tasks the location of the signal of interest is rarely known a priori, known-location tasks are commonly used as a starting point in the literature, as they are simpler to analyze. Two further simplifications of the known-location task are considered often enough in the literature to have gotten their own names. These are the *signal-known-exactly* (SKE) tasks, where the signal profile $\boldsymbol{f}_s$ is deterministic, and the *background-known-exactly* (BKE) tasks, where the background $\boldsymbol{f}_b$ is deterministic.

Human observers have difficulty distinguishing details, such as small lesions, in a sinogram; reconstructed images are better suited for this purpose. To reflect common imaging practice, we focus hereafter on the case where the observer extracts its feature vector from a reconstructed image $\hat{\boldsymbol{f}} = \hat{\boldsymbol{f}}(\boldsymbol{y})$, rather than directly from the data $\boldsymbol{y}$. That is, $\boldsymbol{v} = \boldsymbol{v}(\hat{\boldsymbol{f}}(\boldsymbol{y}))$, where the mapping $\hat{\boldsymbol{f}}(\boldsymbol{y})$ is determined by the reconstruction method. Our goal is to find reconstruction methods that allow observers to achieve optimal performance in a specified detection task. We now discuss how this performance is quantified.

2.3.3 Figures of Merit

In binary-hypothesis tasks, one can evaluate the detection performance of an observer by tracing its *receiver operating characteristic* (ROC) curve. This is a plot of the *probability of detection* or *probability of a true positive, i.e.*, deciding that the signal is present when it is actually present:

$$P_{\mathrm{D}} \triangleq \mathsf{P}\{D_1 | H_1\}, \tag{2.17}$$

versus the *probability of false alarm* or *probability of a false positive, i.e.*, deciding that the signal is present when it is actually absent:

$$P_{\text{FA}} \triangleq \mathsf{P}\{D_1|H_0\}\,. \tag{2.18}$$

For a decision rule of the form (2.16), the trade-off between these two probabilities is plotted by varying the decision threshold τ. A common figure of merit for observers is the *area under the curve* (AUC)[3].

Another common figure of merit for binary-hypothesis tasks is the *signal-to-noise ratio* (SNR), defined as

$$\text{SNR} \triangleq \frac{\mathsf{E}[t|H_1] - \mathsf{E}[t|H_0]}{\sqrt{\frac{1}{2}\,\mathsf{Var}\{t|H_1\} + \frac{1}{2}\,\mathsf{Var}\{t|H_0\}}}, \tag{2.19}$$

where $\mathsf{E}[\cdot|H_i]$ and $\mathsf{Var}\{\cdot|H_i\}$ denote the expected value and variance respectively of some random variable under hypothesis H_i. In the case where the test statistic t is Gaussian-distributed under each hypothesis, the SNR is especially useful, since it is monotonically related to the AUC [11, p.819]:

$$\text{AUC} = \frac{1}{2}\left[1 + \text{erf}\left(\frac{\text{SNR}}{2}\right)\right], \tag{2.20}$$

where $\text{erf}(\cdot)$ is the error function. By inverting (2.20), one can also define the *detectability index* d_A:

$$d_A \triangleq 2\text{erf}^{-1}(2(\text{AUC}) - 1).$$

Obviously, $\text{SNR} = d_A$ when t is Gaussian-distributed under each hypothesis.

For binary-hypothesis tasks, tracing the ROC curve is enough to quantify the observer's performance. For detection tasks where the signal location is not known a priori, however,

---

[3] Alternatively, one may consider the *area under the error-exponent curve,* which is easier to analyze when the error exponent is available in closed form [54].

one may be interested not only in detection performance, but also in localization performance. Then the *probability of correct localization* becomes important. The concept of localization is often extended to include a *tolerance region,* so that the signal location determined by the observer is considered correct if it lies within a certain distance from the true signal location. In this case, the probability of correct localization is defined as

$$P_{\mathrm{CL}} \triangleq \mathsf{P}\{\cup_{\ell=1}^{n_{\mathrm{L}}} \cup_{j\in\mathcal{N}_\ell} (D_j \cap H_\ell) | \cup_{\ell=1}^{n_{\mathrm{L}}} H_\ell\} = \frac{\sum_{\ell=1}^{n_{\mathrm{L}}} \sum_{j\in\mathcal{N}_\ell} \mathsf{P}\{D_j|H_\ell\}\,\mathsf{P}\{H_\ell\}}{\sum_{\ell=1}^{n_{\mathrm{L}}} \mathsf{P}\{H_\ell\}}, \tag{2.21}$$

where $\mathcal{N}_\ell$ is the set of locations that lie within the tolerance region around location ℓ. When there is no tolerance, this set reduces to $\mathcal{N}_\ell = \{\ell\}$, $\ell = 1, \ldots, n_{\mathrm{L}}$.

Localization performance can be quantified by tracing the *localization ROC* (LROC) curve, a plot of the probability of correct localization in (2.21) versus the probability of false alarm, defined in the unknown-location case as

$$P_{\mathrm{FA}} \triangleq \mathsf{P}\{\cup_{\ell=1}^{n_{\mathrm{L}}} D_\ell | H_0\} = \sum_{\ell=1}^{n_{\mathrm{L}}} \mathsf{P}\{D_\ell|H_0\}\,. \tag{2.22}$$

If the tolerance region for correct localization is expanded to include all possible locations, *i.e.*, $\mathcal{N}_\ell = \{1, \ldots, n_{\mathrm{L}}\}$, $\ell = 1, \ldots, n_{\mathrm{L}}$, then the task becomes detection without localization and (2.21) reduces to the detection probability

$$P_{\mathrm{D}} \triangleq \mathsf{P}\{\cup_{\ell=1}^{n_{\mathrm{L}}} D_\ell | \cup_{\ell=1}^{n_{\mathrm{L}}} H_\ell\}\,. \tag{2.23}$$

The detection and false-alarm probabilities (2.23) and (2.22) respectively define a ROC curve for the unknown-location case, much like (2.17) and (2.18) do for the known-location case.

2.3.4 Image Reconstruction and Human Observer Performance

Humans are a natural starting point as observers of interest, since usually they are the ones performing clinical tasks in practice. The decision rule, that is, the explicit form of the discriminant $t(\cdot)$ that human observers use in signal detection tasks is not known. As a result, the SNR of human observers could not be calculated explicitly from (2.19), even if the statistics of $\boldsymbol{f}$ and $\boldsymbol{y}$ were known. The ROC curve must therefore be traced empirically to evaluate human performance [27, 86, 92].

To trace a ROC curve, the observer is presented with a series of realizations of the reconstructed image where the signal of interest may or may not be present. The observer is asked to rate her or his confidence in the signal's presence in each image. These confidence ratings represent different values of the detection threshold and thus different operating points on the ROC curve. The true-positive and false-positive probabilities at each of these points are estimated from repeated trials of the experiment and used to fit the observer's ROC curve. If the observer is also asked to determine the location of the signal in each image, the probabilities of correct localization can similarly be estimated and used to fit a LROC curve.

Various experimental studies have shown that the choice of reconstruction method can indeed affect the performance of human observers in detection/localization tasks. For example, Tourassi *et al.* investigated how human observers performed in detecting a lesion in images reconstructed from simulated SPECT data [128]. They found that EM reconstruction led to better performance than FBP reconstruction. LaCroix and Tsui found that EM with non-uniform attenuation correction led to improved detectability of myocardial perfusion defects in SPECT images, when compared to FBP without attenuation correction [77]. Wells *et al.* found that, compared to FBP, OSEM extended the range of lesion contrasts over which human observers perform accurate lesion detection and localization in thorax SPECT reconstructed images [141]. Gifford *et al.* found that compensating for the depth-dependent system response in image reconstruction for thorax SPECT significantly improved the localization performance of human observers [50]. They also showed how this performance varied per OSEM iteration. Later, Narayanan *et al.* found that compensating for non-uniform attenuation, scatter, and distance-dependent response in OSEM-reconstructed myocardial SPECT images led to better defect detection performance than compensating for only a subset of these factors [98].

Experiments like these offer valuable insight into the effect that various reconstruction parameters have on the performance of human observers. However, they are too time-consuming to perform for many values of a parameter of interest or to repeat every time that some aspect of the imaging process changes. Thus experiments with human observers are not practical as a tool to optimize reconstruction methods. This has led investigators to seek mathematical observer models whose performance in tasks of interest approximates the performance of humans and allows analytical treatment.

2.3.5 The Ideal Observer

As mentioned above, observers can achieve improved performance by incorporating previous knowledge about the imaging and feature extraction process in their decision rule. Human observers presumably acquire such knowledge through training, *i.e.*, looking at several images that have been generated by the same process. Mathematical observers attempt to model this knowledge in the form of various statistics of the feature vector included in their discriminant function. An important aspect of modeling human observers is determining the appropriate degree of knowledge about these statistics to include in the model's discriminant.

The *ideal observer* for performing a binary-hypothesis task given the data $\boldsymbol{y}$ has full knowledge of the statistics of $\boldsymbol{y}$. It has been shown [132, p.26] that the discriminant function of this observer is the well-known *likelihood ratio*:

$$t_{\circ}(\boldsymbol{y}) = \frac{p(\boldsymbol{y}|H_1)}{p(\boldsymbol{y}|H_0)}, \tag{2.24}$$

where $p(\cdot|H_i)$ denotes the probability distribution of a random vector under hypothesis H_i. The optimal threshold is equal to

$$\tau = \frac{(c_{10} - c_{00})p_1}{(c_{01} - c_{11})p_0}, \tag{2.25}$$

where p_i is the *a priori* probability of the hypothesis H_i, and c_{ij} is the Bayes cost associated with deciding on H_i when the truth is H_j. This observer is ideal in the sense that it minimizes an average Bayes cost (*e.g.*, average probability of error). Furthermore, by the Neyman-Pearson lemma, the discriminant (2.24) maximizes the true-positive probability for any given false-positive probability, thus maximizing the AUC.

Conducting the test defined by (2.24) and (2.25) requires knowledge of the statistics of both the object variability and the imaging noise, as well as the a priori probabilities of the hypotheses, referred to in medical diagnosis as *disease prevalence*. In practice, this much information is never known, so the ideal observer serves as a standard against which the detection performance of other observers can be compared.

The performance of the ideal observer is invariant to nonsingular transformations of the data $\boldsymbol{y}$ [44, §10.1]. However, using a feature vector $\boldsymbol{v} = \boldsymbol{v}(\boldsymbol{y})$ instead of the data $\boldsymbol{y}$ will lead to a decrease in the ideal detection performance if the mapping $\boldsymbol{v}(\boldsymbol{y})$ discards useful information in the data, *i.e.*, if $\boldsymbol{v}(\boldsymbol{y})$ is not a sufficient statistic for discriminating between H_0 and H_1. Thus the performance of the ideal discriminant for observation in the feature vector space,

$$t_\circ(\boldsymbol{y}) = \frac{p(\boldsymbol{v}(\boldsymbol{y})|H_1)}{p(\boldsymbol{v}(\boldsymbol{y})|H_0)}, \tag{2.26}$$

will generally be inferior to that of (2.24).

As previously mentioned, observers that use the tomographic data $\boldsymbol{y}$ as their feature vector would not be appropriate as models of humans. The observer models described in the following sections extract their feature vector from the reconstructed image. Some of them use the reconstructed image itself as the feature vector, *i.e.*, $\boldsymbol{v} = \hat{\boldsymbol{f}}(\boldsymbol{y})$, and it may be safe to assume that a sensible reconstruction method would not discard useful information in the data. Other models, however, include a non-invertible feature extraction step in an attempt to capture the suboptimality of human observers.

2.3.6 Linear Observers

In contrast to the performance of an ideal observer that cannot be improved by any data transformation, the performance of human observers can benefit from appropriate reconstruction methods, as discussed in section 2.3.4. The ideal discriminants for all but the simplest detection tasks are nonlinear, but those discriminants have been known to overestimate the performance of human observers. It has been argued that linear discriminants are more appropriate to model the suboptimality of human observers in known-location tasks. Thus, although the specific mechanisms through which each model attempts to capture this suboptimality may vary, most of the popular observer models that have been proposed for this purpose are linear [10,13,22,23,32,85,96].

A *linear discriminant* can be defined as the inner product of some *template* $\boldsymbol{w}$, which is a real vector in the same space as $\hat{\boldsymbol{f}}$, with $\hat{\boldsymbol{f}}$:

$$t(\boldsymbol{y}) = \boldsymbol{w}'\hat{\boldsymbol{f}}(\boldsymbol{y}). \tag{2.27}$$

Strategies have been proposed for estimating the template $\boldsymbol{w}$ directly from human observer data [3,4]. Furthermore, explicit mathematical forms, involving the statistics of the reconstructed images, have been proposed for the template $\boldsymbol{w}$ and their performance has been compared to that of humans. In the following, we list the linear observer models that are most commonly encountered in the objective image quality literature, e.g., [7–9].

The Hotelling Observer

The *Hotelling observer* (HO) [61] assumes knowledge of the first- and second-order statistics of the reconstructed image $\hat{\boldsymbol{f}}$ under both the signal absent and signal present hypotheses. This results in the *ideal linear discriminant*, that is the one that achieves maximum SNR:

$$\boldsymbol{w}_{\mathrm{HO}} \triangleq \boldsymbol{\mathcal{K}}_{\hat{\boldsymbol{f}}}^{\dagger}\Big(\mathsf{E}\Big[\hat{\boldsymbol{f}}|H_1\Big] - \mathsf{E}\Big[\hat{\boldsymbol{f}}|H_0\Big]\Big), \tag{2.28}$$

where

$$\boldsymbol{\mathcal{K}}_{\hat{\boldsymbol{f}}} \triangleq \tfrac{1}{2}\,\mathsf{Cov}\Big\{\hat{\boldsymbol{f}}|H_1\Big\} + \tfrac{1}{2}\,\mathsf{Cov}\Big\{\hat{\boldsymbol{f}}|H_0\Big\}$$

is the unconditional covariance matrix of $\hat{\boldsymbol{f}}$, known as the *intra-class scatter matrix* in the pattern classification literature, the superscript "†" denotes a pseudo-inverse, and $\mathsf{E}[\cdot|H_i]$ and $\mathsf{Cov}\{\cdot|H_i\}$ denote the mean vector and covariance matrix respectively of some random vector under hypothesis H_i.

The HO is the ideal observer in SKE tasks where the data have a Gaussian distribution under both hypotheses and it is the ideal linear observer otherwise. No linear transformation of the data can improve the SNR of this ideal linear observer, just like no transformation of the data can improve the performance of the ideal observer. The performance of the Hotelling observer was found to be well-correlated with human observer performance in a simple detection task in [42].

The Prewhitening Observer

The *prewhitening* (PW) observer assumes knowledge of the second-order statistics of the reconstructed image $\hat{\boldsymbol{f}}$ under the signal absent, but not under the signal present hypothesis. The corresponding template is

$$\boldsymbol{w}_{\mathrm{PW}} \triangleq \mathsf{Cov}\Big\{\hat{\boldsymbol{f}}|H_0\Big\}^{\dagger}\Big(\mathsf{E}\Big[\hat{\boldsymbol{f}}|H_1\Big] - \mathsf{E}\Big[\hat{\boldsymbol{f}}|H_0\Big]\Big). \tag{2.29}$$

In SKE tasks the covariance of the reconstructed image $\hat{\boldsymbol{f}}$ is equal under both hypotheses. In such tasks, if the reconstruction method is linear, then the optimal linear discriminant (HO) reduces to the PW discriminant.

The Non-Prewhitening Observer

The *non-prewhitening* (NPW) observer assumes knowledge only of the first-order statistics of the reconstructed image $\hat{\boldsymbol{f}}$ under both the signal absent and signal present hypotheses. The corresponding template is

$$\boldsymbol{w}_{\mathrm{NPW}} \triangleq \mathsf{E}\left[\hat{\boldsymbol{f}}|H_1\right] - \mathsf{E}\left[\hat{\boldsymbol{f}}|H_0\right]. \tag{2.30}$$

In SKE tasks, if the reconstruction method is linear and the background and imaging noise present in the reconstructed image $\hat{\boldsymbol{f}}$ are additive, white, and independent, then the optimal linear discriminant (HO) reduces to the NPW discriminant. However, rarely is the noise in reconstructed images white in practice.

The Region-of-Interest Observer

The *region-of-interest* (ROI) observer assumes knowledge only of the first-order statistics of the target signal $\boldsymbol{f}_s$. The corresponding template is that of a simple *matched filter*:

$$\boldsymbol{w}_{\mathrm{ROI}} \triangleq \mathsf{E}[\boldsymbol{f}_s|H_1]. \tag{2.31}$$

If the reconstruction method is linear and chosen so that the reconstructed image $\hat{\boldsymbol{f}}$ is an unbiased estimate of the true object $\boldsymbol{f}$, then the NPW discriminant reduces to the ROI discriminant. However, due to the smoothing mechanisms of practical reconstruction methods, rarely is $\hat{\boldsymbol{f}}$ unbiased.

2.3.7 Channelized Linear Observers

Channelized observer models include a set of frequency-selective channels in an attempt either to construct an efficient basis for the approximation of the ideal linear observer [45], or to model the frequency selectivity that is believed to characterize human visual perception [96]. We are interested primarily in the latter type of channel, since sub-optimal observers are the ones used for the purpose of image reconstruction optimization. However, the analysis in the following chapters applies to either channel flavor.

Conceptually, channelized observers first pass the reconstructed image $\hat{\boldsymbol{f}}$ through a set of M bandpass filters. In known location tasks, these observers form a new feature vector $\hat{\boldsymbol{c}} \in \mathbb{C}^M$ from the values of the filter outputs at the location of the target signal center, potentially with the addition of some noise:

$$\hat{\boldsymbol{c}}(\boldsymbol{y}) = \boldsymbol{\mathcal{C}}' \hat{\boldsymbol{f}}(\boldsymbol{y}) + \boldsymbol{\varepsilon}_{\text{int}}, \tag{2.32}$$

where $\boldsymbol{\mathcal{C}} = [\boldsymbol{\mathcal{C}}_1, \ldots, \boldsymbol{\mathcal{C}}_M]$ consists of M operators. The mth of these operators applies the impulse response of the mth bandpass filter and samples the output at the center of the target signal. Typically this filtering step is not invertible and it greatly reduces the dimensionality of the detection problem (*e.g.*, $M = 4$ in [52]). The *internal noise* vector $\boldsymbol{\varepsilon}_{\text{int}}$ models inherent uncertainty in the observer's decisions. A generic channelized linear observer forms its test statistic t_{ch} by applying a template $\boldsymbol{w} \in \mathbb{C}^M$ to the output of the filter bank:

$$t_{\text{ch}}(\boldsymbol{y}) = \boldsymbol{w}' \hat{\boldsymbol{c}}(\boldsymbol{y}). \tag{2.33}$$

The Channelized Hotelling Observer

The *channelized Hotelling observer* (CHO), introduced by Myers and Barrett [96], has been found to be particularly successful in predicting human observer performance in known-location tasks. It applies the optimal linear discriminant with respect to the output $\hat{\boldsymbol{c}}$ of the M-channel filter bank. This corresponds to the template

$$\boldsymbol{w}_{\text{CHO}} \triangleq \boldsymbol{\Pi}_{\hat{\boldsymbol{c}}}^{\dagger} \Big(\mathsf{E}[\hat{\boldsymbol{c}} | H_1] - \mathsf{E}[\hat{\boldsymbol{c}} | H_0] \Big), \tag{2.34}$$

where

$$\boldsymbol{\Pi}_{\hat{\boldsymbol{c}}} \triangleq \tfrac{1}{2}\,\mathsf{Cov}\{\hat{\boldsymbol{c}}|H_1\} + \tfrac{1}{2}\,\mathsf{Cov}\{\hat{\boldsymbol{c}}|H_0\}$$

is the unconditional covariance matrix of $\hat{\boldsymbol{c}}$.

The Channelized Non-Prewhitening Observer

The *channelized non-prewhitening* (CNPW) observer applies the NPW template with respect to the output $\hat{\boldsymbol{c}}$ of the filter bank:

$$\boldsymbol{w}_{\mathrm{CNPW}} \triangleq \mathsf{E}[\hat{\boldsymbol{c}}|H_1] - \mathsf{E}[\hat{\boldsymbol{c}}|H_0]\,. \tag{2.35}$$

A special case is the introduction of a single “eye filter”, proposed to model the frequency response of the human visual system, to the NPW observer [85]. This model has been superseded in the literature by models using multiple bandpass filters.

Channel Profiles

For tasks involving the detection of circularly symmetric signals, typically the channel responses are also modeled as circularly symmetric. Channels with various passband profiles, including square or difference-of-Gaussians, have been proposed in the literature to model the frequency selectivity of human observers. As an example, Fig. 2.6 shows profiles of the frequency responses of three channel sets used with the CHO model in [1]. These are the square (SQR), sparse difference-of-Gaussians (S-DOG), and dense difference-of-Gaussians (D-DOG) channel sets. Gabor functions have also been proposed as channel profiles [29, 88]. For tasks involving the detection of signals that are not circularly symmetric, channels with spatial orientations, such as the ones proposed in [64], may be more appropriate.

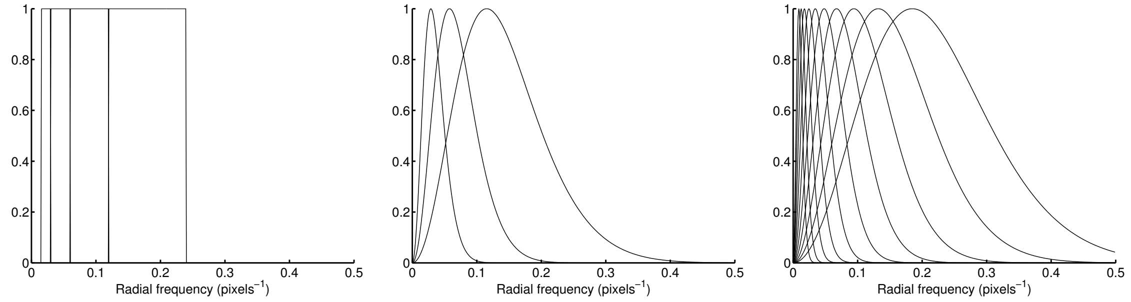

(a) SQR channels (b) S-DOG channels (c) D-DOG channels

Figure 2.6: Profiles of channel frequency responses for three examples of channel sets. The square (SQR), sparse difference-of-Gaussians (S-DOG), and dense difference-of-Gaussians (D-DOG) channels shown here have been investigated with the CHO model in [1].

2.3.8 Model Observers for Unknown-Location Tasks

For tasks where there are multiple candidate locations for the target signal, it can be shown [68] that the ideal discriminant is the following *modified generalized likelihood ratio:*

$$t_{\max_\circ}(\boldsymbol{v}(\boldsymbol{y})) \triangleq \max_{\ell=1,\ldots,n_{\mathrm{L}}} \frac{\sum_{j\in\mathcal{N}_\ell} p(\boldsymbol{v}(\boldsymbol{y})|H_j)p_j}{p(\boldsymbol{v}(\boldsymbol{y})|H_0)}. \tag{2.36}$$

The ideal strategy for choosing the most suspicious location $\ell_\circ$ and deciding whether the signal is present or absent is

$$\begin{aligned} \ell_\circ(\boldsymbol{v}(\boldsymbol{y})) &= \arg\max_{\ell=1,\ldots,n_{\mathrm{L}}} \frac{\sum_{j\in\mathcal{N}_\ell} p(\boldsymbol{v}(\boldsymbol{y})|H_j)p_j}{p(\boldsymbol{v}(\boldsymbol{y})|H_0)} \\ \text{Decide } H_{\ell_\circ(\boldsymbol{v}(\boldsymbol{y}))} &\quad \text{if} \quad t_{\max_\circ}(\boldsymbol{v}(\boldsymbol{y})) > \tau, \text{ otherwise decide } H_0, \end{aligned}$$

for some data-independent threshold τ. This strategy is ideal in the sense that it maximizes the probability of correct localization for any given false-positive probability, thus maximizing the area under the LROC curve. It also minimizes the average Bayes cost, for some constraints on the costs allocated to different types of erroneous decisions. In [68], this result is also generalized for a continuum of locations, where the sum in (2.36) is replaced by an integral.

If the a priori probabilities of all candidate signal locations, p_ℓ, $\ell = 1, \ldots, n_{\mathrm{L}}$, are equal and there is no localization tolerance, then (2.36) reduces to the usual generalized likelihood ratio known from detection theory (with the location ℓ as a nuisance parameter). In the binary-hypothesis case where $n_{\mathrm{L}} = 1$, the ideal discriminant in (2.36) becomes equivalent (to within a scaling factor) to the likelihood ratio in (2.26).

In general, the observer models used in the literature for unknown-location tasks compute some scalar *local test statistic* $t_\ell = t_\ell(\boldsymbol{v}(\boldsymbol{y}))$ for each of the candidate locations and compare the *maximum test statistic,*

$$t_{\max} \triangleq \max_{\ell=1,\ldots,n_{\mathrm{L}}} t_\ell, \tag{2.37}$$

to some threshold. However, unlike the ideal local test statistic in (2.36), the test statistics that have been used to model the suboptimality of human observers are linear, *i.e.*,

$$t_\ell = \boldsymbol{w}_\ell' \boldsymbol{v}(\boldsymbol{y}), \tag{2.38}$$

for some *local template* $\boldsymbol{w}_\ell$, which is a real vector in the same space as the feature vector $\boldsymbol{v}$. The local test statistics are typically based on the channelized linear observer models described in previous paragraphs. The *maximum CHO* (MaCHO) uses a local template $\boldsymbol{w}_\ell$ that follows (2.34), whereas the *maximum CNPW* (MaCNPW) observer uses a local template $\boldsymbol{w}_\ell$ that follows (2.35).

2.3.9 Correlation with Human Observers

There is an extensive literature comparing the performance of the observer models above to that of human observers in some simple detection tasks, especially tasks where the location of the target signal is known a priori. We review here some representative findings.

Known Signal Location

Rolland and Barrett found good correlation between the performance of the CHO and human observers for the task of detecting signals of known location in images with statistically varying "lumpy backgrounds" [113]. Wollenweber *et al.* compared human observers to CHO models with radially symmetric or oriented channels in the detection of myocardial perfusion defects in SPECT images. They found good correlation between human and CHO performance [142]. Gifford *et al.* found good correlation in how the CHO and human observers performed in SKE/BKE lesion detection in simulated SPECT images as the scatter compensation strategy in the reconstruction algorithm was varied [48]. Abbey and Barrett compared the effect of various parameters of the imaging process (including the amount of post-smoothing) on the detection performance of several observer models [1]. They found the CHO to best capture the effect of these parameters on human observer performance. Sankaran *et al.* also found good correlation in the effect of the cut-off frequency of a post-filter on the performance of the CHO and human observers detecting myocardial perfusion defects in SPECT images [114].

Unknown Signal Location

Research on observer models for tasks where the location of the signal is unknown is less extensive. Gifford *et al.* compared the performance of some observer models with a maximum test statistic to the performance of human observers in a task involving location variability in simulated thorax SPECT images [51]. They found the MaCNPW model to better capture human performance than the MaCHO model as the level of scatter in the images increased. Based on these results, the authors suggested that the prewhitening ability of humans may degrade when location uncertainty is introduced in the detection task. The authors also found the MaCNPW model to be reasonably correlated with humans in a study involving PET images in various display formats [47]. The MaCHO, however, was not investigated in that study.

Three-Dimensional Observer Models

The observer templates described above are applied to single slices of the reconstructed object. Wells *et al.* found that results on human observer performance in detection tasks that involve single image slices may not generalize to tasks where multiple slices are available [140]. There have been efforts to construct observer models that base their decisions on multiple slices of the reconstructed object. However, research on the correlation of these models to humans is still in preliminary stages. Gifford *et al.* recently investigated a multislice CHO and multislice CNPW model in cardiac SPECT [49]. They found that both modeled human performance on OSEM-reconstructed images equally well, but the multislice CHO modeled human performance on FBP-reconstructed images better than the multislice CNPW.

2.3.10 Image Reconstruction and Model Observer Performance

As previously mentioned, the performance of human observers does not lend itself to analytical optimization. Mathematical observer models that are well-correlated with humans are thus an attractive alternative for gaining insight on how image reconstruction affects classification performance. There have been increasing efforts in recent years to analyze the effect of reconstruction methods on the performance of some of the model observers described in previous paragraphs.

One approach to optimizing reconstruction parameters with respect to model observer performance is to perform simulations and evaluate performance in a manner analogous to experiments with human observers. For example, Myers and Hanson used the detection performance of the NPW observer in an SKE task, computed through simulations, to compare ART and maximum-entropy reconstruction and optimize their parameters [97]. Wagner *et al.* suggested more generally that knowledge of the task that is to be performed on the reconstructed images is a form of prior knowledge that should be included in Bayesian image reconstruction [137]. They similarly examined optimizing a regularized maximum-entropy method.

To avoid time-consuming simulations, efforts have shifted towards deriving analytical expressions for the detection performance of observer models when combined with some popular reconstruction methods. These expressions can subsequently be used to optimize the reconstruction parameters. Qi and Huesman used such expressions to show that MAP outperforms FBP in terms of SKE/BKE detection for the NPW and PW observer models [106]. The analysis of known-location tasks has been extended to include channelized observers [145] and statistically varying backgrounds and signal profiles [104].

Deriving analytical expressions for the performance of model observers for unknown-location tasks is more challenging. Swensson analyzed this performance under some special condition of statistical independence [126]. Khurd and Gindi used Swensson's results to analyze the LROC performance of the MaCHO for images reconstructed with penalized-likelihood methods for the case where the local test statistics at different candidate locations are independent [69]. However, the assumption of statistical independence generally does not hold when the candidate locations are near each other. For this more general situation, analytical approximations of the moments of the local test statistics have been combined with simulations to trace the LROC [70, 107].

In the following chapters we will illustrate which reconstruction methods optimize the performance of various model observers for known-location tasks, drawing on our presentation in [41, 149, 151]. We will also derive analytical approximations for the performance of model observers for unknown-location tasks, expanding on our presentation in [148].

CHAPTER 3

Known Signal Location

The assumption that there is only one possible location where the target signal may appear is fairly removed from realistic clinical tasks. However, due to their simplicity, known-location tasks have received considerable attention in the literature and we will use them as our starting point as well. We begin by establishing our notation for the objects, measurements, and reconstruction methods that we will be considering. We then analyze the detection performance of the observer models for known-location tasks that were described in chapter 2, with the objective of finding reconstruction methods that allow each of these observers to achieve optimal performance in such tasks. We show that, for several observer models, the optimal reconstruction method need not include regularization. However, in the case of channelized observers, which are the ones of greatest practical interest, the importance of regularization increases as the observer's prewhitening capabilities deteriorate.

3.1 The Detection Task

3.1.1 Overview

As discussed in section 2.3.2, we adopt an additive model for the background $\boldsymbol{f}_b$ and the target signal $\boldsymbol{f}_s$. The known-location detection task is a decision between the following pair of hypotheses:

$$\begin{aligned} H_0 &: \boldsymbol{f} = \boldsymbol{f}_b && \text{(signal absent)} \\ H_1 &: \boldsymbol{f} = \boldsymbol{f}_b + \boldsymbol{f}_s && \text{(signal present).} \end{aligned} \tag{3.1}$$

The background $\boldsymbol{f}_b$ and the signal $\boldsymbol{f}_s$ are random and statistically independent. We denote their expectations by $\bar{\boldsymbol{f}}_b$ and $\bar{\boldsymbol{f}}_s$ respectively. We denote their covariances by $\boldsymbol{\mathcal{K}}_b$ and $\boldsymbol{\mathcal{K}}_s$ respectively. We assume that the signal is concentrated in space around a given location. Thus the only source of randomness in the signal is its profile and not its location.

In the general case where both background and signal are random, $\boldsymbol{\mathcal{K}}_b$ and $\boldsymbol{\mathcal{K}}_s$ are positive definite. In a SKE task we have $\boldsymbol{\mathcal{K}}_s = \mathbf{0}$ and thus a deterministic signal $\boldsymbol{f}_s = \bar{\boldsymbol{f}}_s$. In a BKE task we have $\boldsymbol{\mathcal{K}}_b = \mathbf{0}$ and thus a deterministic background $\boldsymbol{f}_b = \bar{\boldsymbol{f}}_b$. In all cases, we assume that $\bar{\boldsymbol{f}}_b$, $\bar{\boldsymbol{f}}_s$, $\boldsymbol{\mathcal{K}}_b$ and $\boldsymbol{\mathcal{K}}_s$ are known.

3.1.2 Measurement Model

In the following, all moments are averages over both $\boldsymbol{y}$ and $\boldsymbol{f}$, except when they are accompanied by the subscript $\boldsymbol{f}$, in which case they are averages only over $\boldsymbol{f}$, or when they are conditional averages over $\boldsymbol{y}$ for a given instance of $\boldsymbol{f}$, in which case the usual $\cdot|\boldsymbol{f}$ notation is used. We assume knowledge of the following statistics of the measurement $\boldsymbol{y}$:

$$\bar{\boldsymbol{y}} \triangleq \mathsf{E}[\boldsymbol{y}|H_1] - \mathsf{E}[\boldsymbol{y}|H_0] \tag{3.2}$$

$$\boldsymbol{\Pi}_i \triangleq \mathsf{Cov}\{\boldsymbol{y}|H_i\}, \;\; i = 0, 1. \tag{3.3}$$

In emission tomography, where the measurement $\boldsymbol{y}$ is a noisy sinogram, its elements are independent and Poisson-distributed conditional on the object $\boldsymbol{f}$ following (2.6). Then the conditional mean and covariance of $\boldsymbol{y}$ are, respectively,

$$\begin{aligned} \mathsf{E}[\boldsymbol{y}|\boldsymbol{f}] &= \boldsymbol{\mathcal{A}}\boldsymbol{f} + \boldsymbol{r} \\ \mathsf{Cov}\{\boldsymbol{y}|\boldsymbol{f}\} &= \mathsf{diag}\{\boldsymbol{\mathcal{A}}\boldsymbol{f} + \boldsymbol{r}\}, \end{aligned}$$

where we use the notation $\mathsf{diag}\{\boldsymbol{a}\}$ for a diagonal matrix with diagonal elements equal to the elements of vector $\boldsymbol{a}$. Both the linear system operator $\boldsymbol{\mathcal{A}}$ and the randoms vector $\boldsymbol{r}$ are assumed to be deterministic and known. The moments of the sinogram $\boldsymbol{y}$ under hypothesis $H_i, i = 0, 1$ are then given by

$$\begin{aligned} \mathsf{E}[\boldsymbol{y}|H_i] &= \mathsf{E}_{\boldsymbol{f}}\big[\mathsf{E}[\boldsymbol{y}|\boldsymbol{f}]\,\big|H_i\big] = \boldsymbol{\mathcal{A}}\,\mathsf{E}_{\boldsymbol{f}}[\boldsymbol{f}|H_i] + \boldsymbol{r} && (3.4) \\ \mathsf{Cov}\{\boldsymbol{y}|H_i\} &= \mathsf{E}_{\boldsymbol{f}}\big[\mathsf{Cov}\{\boldsymbol{y}|\boldsymbol{f}\}\,\big|H_i\big] + \mathsf{Cov}_{\boldsymbol{f}}\big\{\mathsf{E}[\boldsymbol{y}|\boldsymbol{f}]\,\big|H_i\big\} \\ &= \mathsf{diag}\{\boldsymbol{\mathcal{A}}\,\mathsf{E}_{\boldsymbol{f}}[\boldsymbol{f}|H_i] + \boldsymbol{r}\} + \boldsymbol{\mathcal{A}}\,\mathsf{Cov}_{\boldsymbol{f}}\{\boldsymbol{f}|H_i\}\,\boldsymbol{\mathcal{A}}'. && (3.5) \end{aligned}$$

Under the assumption that $\boldsymbol{f}_b$ and $\boldsymbol{f}_s$ are independent, using (3.4) and (3.5) for each of the two hypotheses in (3.1) and substituting in (3.2) and (3.3) yields, for emission tomography,

$$\begin{aligned} \bar{\boldsymbol{y}} &= \boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_s && (3.6) \\ \boldsymbol{\Pi}_0 &= \mathsf{diag}\big\{\boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_b + \boldsymbol{r}\big\} + \boldsymbol{\mathcal{A}}\boldsymbol{\mathcal{K}}_b\boldsymbol{\mathcal{A}}' && (3.7) \\ \boldsymbol{\Pi}_1 &= \boldsymbol{\Pi}_0 + \mathsf{diag}\big\{\boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_s\big\} + \boldsymbol{\mathcal{A}}\boldsymbol{\mathcal{K}}_s\boldsymbol{\mathcal{A}}'. && (3.8) \end{aligned}$$

3.1.3 Linear Reconstruction Methods

We focus here on linear reconstruction methods, as several common tomographic reconstruction techniques either are or can be approximated as linear. We denote a generic linear reconstructor by an operator $\boldsymbol{\mathcal{Z}}$. The reconstructed image is then given by

$$\hat{\boldsymbol{f}}(\boldsymbol{y}) = \boldsymbol{\mathcal{Z}}\boldsymbol{y}. \qquad (3.9)$$

We may view the reconstruction $\hat{\boldsymbol{f}}$ either as a vector in a Hilbert space, in which case $\boldsymbol{\mathcal{Z}}$ is a general linear mapping from $\mathbb{R}^{n_\mathrm{d}}$ to that Hilbert space, or as a discrete representation in $\mathbb{R}^{n_\mathrm{p}}$, in which case $\boldsymbol{\mathcal{Z}}$ is a matrix in $\mathbb{R}^{n_\mathrm{p} \times n_\mathrm{d}}$. Combining the linear reconstruction method in (3.9) with the data moments from (3.2) and (3.3) yields the following expressions for the moments of the reconstructed image:

$$\mathsf{E}\left[\hat{\boldsymbol{f}}|H_1\right] - \mathsf{E}\left[\hat{\boldsymbol{f}}|H_0\right] = \boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}} \tag{3.10}$$

$$\mathsf{Cov}\left\{\hat{\boldsymbol{f}}|H_i\right\} = \boldsymbol{\mathcal{Z}}\boldsymbol{\Pi}_i\boldsymbol{\mathcal{Z}}', \quad i = 0, 1. \tag{3.11}$$

An example of a well-known family of reconstruction methods that can be approximated as linear is the penalized-likelihood family, which was discussed in section 2.2.4. In general, linearity is a good approximation for reconstructors of that family except maybe when the non-negativity constraint is enforced [35]. Here we assume that the target signal appears on a background that is sufficiently high to render the non-negativity constraint inactive around the signal location, so these reconstructors can be considered linear near the signal.

Consider, for instance, the PWLS reconstructor, which uses the quadratic data-fit term in (2.13). When the regularization term is also quadratic, as in (2.14), we have a *quadratically PWLS* (QPWLS) reconstructor. We will assume hereafter that the regularization parameter β has been absorbed within the regularization term $R(\boldsymbol{f})$, *i.e.*, the estimate is obtained by maximizing

$$\hat{\boldsymbol{f}}(\boldsymbol{y}) = \arg\max_{\boldsymbol{f}\geq\boldsymbol{0}}[D(\boldsymbol{y}, \boldsymbol{\mathcal{A}}\boldsymbol{f}) - R(\boldsymbol{f})]. \tag{3.12}$$

If the non-negativity constraint is ignored, one can show that (3.12) with data-fit term (2.13) and regularization term (2.14) is equivalent to taking $\hat{\boldsymbol{f}}(\boldsymbol{y}) = (\boldsymbol{\mathcal{A}}'\boldsymbol{W}\boldsymbol{\mathcal{A}} + \boldsymbol{\mathcal{R}})^{-1}\boldsymbol{\mathcal{A}}'\boldsymbol{W}\boldsymbol{y}$. Thus QPWLS corresponds to the linear reconstructor

$$\boldsymbol{\mathcal{Z}} = (\boldsymbol{\mathcal{A}}'\boldsymbol{W}\boldsymbol{\mathcal{A}} + \boldsymbol{\mathcal{R}})^{-1}\boldsymbol{\mathcal{A}}'\boldsymbol{W}. \tag{3.13}$$

For $\boldsymbol{\mathcal{R}} = \boldsymbol{0}$, we get the unregularized WLS reconstructor, which simply maximizes the data-fit term (2.13).

We consider observers that decide between hypotheses H_0 and H_1 based on a feature vector $\boldsymbol{v}$ extracted from the reconstructed image $\hat{\boldsymbol{f}}$ and a decision rule that compares some scalar test statistic

$$t = t(\boldsymbol{v}(\hat{\boldsymbol{f}})) = t(\boldsymbol{v}(\boldsymbol{\mathcal{Z}}\boldsymbol{y})) \tag{3.14}$$

to a threshold as in (2.16). The specific form of the discriminant function depends on the observer model considered. Our goal is to optimize the reconstructor $\boldsymbol{\mathcal{Z}}$ with respect to the performance of various observer models of interest in the detection of $\boldsymbol{f}_s$. In the remainder of this chapter we revisit the observer models for known-location tasks that we listed in section 2.3 and examine the choice of $\boldsymbol{\mathcal{Z}}$ that maximizes detection performance for each of these models.

3.2 Ideal Observer

The ideal discriminant is given by the likelihood ratio in (2.26) or, equivalently, by any monotonic function, such as the logarithm, of the likelihood ratio. In general, the ideal discriminant is nonlinear in the feature vector $\boldsymbol{v}$. For example, if the feature vector is assumed to be Gaussian-distributed (which is a reasonable approximation in some cases, as discussed in the following section), the log-likelihood ratio is

$$\begin{aligned} t_\circ(\boldsymbol{v}) &\triangleq \log \frac{p(\boldsymbol{v}|H_1)}{p(\boldsymbol{v}|H_0)} \\ &= \frac{1}{2}\boldsymbol{v}'\Big(\mathsf{Cov}\{\boldsymbol{v}|H_0\}^\dagger - \mathsf{Cov}\{\boldsymbol{v}|H_1\}^\dagger\Big)\boldsymbol{v} \\ &+ \Big(\mathsf{Cov}\{\boldsymbol{v}|H_1\}^\dagger\,\mathsf{E}[\boldsymbol{v}|H_1] - \mathsf{Cov}\{\boldsymbol{v}|H_0\}^\dagger\,\mathsf{E}[\boldsymbol{v}|H_0]\Big)'\boldsymbol{v}, \end{aligned} \tag{3.15}$$

where terms that are independent of the feature vector $\boldsymbol{v}$ have been discarded. The ideal discriminant in (3.15) is quadratic in $\boldsymbol{v}$ and thus not Gaussian-distributed itself. (The exception to this is the SKE case, when the observation covariance is equal under both hypotheses and the quadratic term in (3.15) vanishes.) Therefore, one would have to optimize the AUC of the ideal test statistic directly rather than work with the SNR.

However, as discussed in section 2.3.5, the performance of the ideal observer is invariant to any nonsingular data transformation. For this reason, the ideal observer is generally not used to evaluate image reconstruction methods, and linear observer models are instead used in the literature for this purpose.

3.3 Linear Observers

3.3.1 Overview

We first consider a generic linear observer that uses as its feature vector an image $\hat{\boldsymbol{f}}$ produced by a linear reconstruction method. Combining the linear discriminant in (2.27) with the reconstruction moments from (3.10) and (3.11) yields the following expressions for the moments of $t = t(\boldsymbol{y})$:

$$\mathsf{E}[t|H_1] - \mathsf{E}[t|H_0] = \boldsymbol{w}'\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}} \tag{3.16}$$

$$\mathsf{Var}\{t|H_i\} = \boldsymbol{w}'\boldsymbol{\mathcal{Z}}\boldsymbol{\Pi}_i\boldsymbol{\mathcal{Z}}'\boldsymbol{w}, \quad i = 0, 1. \tag{3.17}$$

Hereafter, we assume that the test statistic t is Gaussian-distributed and focus on maximizing the SNR, in which case the AUC is also maximized. For a linear observer $\boldsymbol{w}$ and a linear reconstructor $\boldsymbol{\mathcal{Z}}$, the test statistic t is a weighted sum of the elements of the measurement vector $\boldsymbol{y}$, so usually t can be approximated as Gaussian-distributed by the central limit theorem.

Furthermore, the probability distribution of an image $\hat{\boldsymbol{f}}$ reconstructed from Poisson data $\boldsymbol{y}$ through a penalized-likelihood method of the form (3.12) can be approximated as Gaussian [109]. This is an additional argument for considering t to be Gaussian-distributed when it is the product of a linear observer $\boldsymbol{w}$ applied on an image $\hat{\boldsymbol{f}}$ reconstructed by a penalized-likelihood algorithm.

Substituting the moment expressions (3.16) and (3.17) into (2.19) yields the SNR of the generic linear observer in (2.27):

$$\text{SNR}^2_{\text{lin}} = \frac{(\boldsymbol{w}'\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}})^2}{\boldsymbol{w}'\boldsymbol{\mathcal{Z}}\check{\boldsymbol{\Pi}}\boldsymbol{\mathcal{Z}}'\boldsymbol{w}} = \frac{\boldsymbol{w}'\boldsymbol{\mathcal{Z}}(\bar{\boldsymbol{y}}\,\bar{\boldsymbol{y}}')\boldsymbol{\mathcal{Z}}'\boldsymbol{w}}{\boldsymbol{w}'\boldsymbol{\mathcal{Z}}\check{\boldsymbol{\Pi}}\boldsymbol{\mathcal{Z}}'\boldsymbol{w}}, \tag{3.18}$$

where

$$\check{\boldsymbol{\Pi}} \triangleq \frac{1}{2}\boldsymbol{\Pi}_1 + \frac{1}{2}\boldsymbol{\Pi}_0 \tag{3.19}$$

is the unconditional covariance of the data. The left-hand side of the SNR expression (3.18) has the form of a generalized Rayleigh quotient. This form is maximized with respect to $\boldsymbol{\mathcal{Z}}'\boldsymbol{w}$ when (*e.g.*, see [31, p.120])

$$\boldsymbol{\mathcal{Z}}'\boldsymbol{w} \propto \check{\boldsymbol{\Pi}}^{-1}\,\bar{\boldsymbol{y}}\,. \tag{3.20}$$

When the observer template $\boldsymbol{w}$ and the reconstructor $\boldsymbol{\mathcal{Z}}$ satisfy the condition (3.20), the observer at hand achieves the maximum SNR that is attainable with any combination of a linear observer and linear reconstruction method.

Substituting the optimality condition (3.20) into the SNR expression (3.18) yields:

$$\text{SNR}^2_{\text{lin}} \le \bar{\boldsymbol{y}}'\,\check{\boldsymbol{\Pi}}^{-1}\,\bar{\boldsymbol{y}} \triangleq \text{SNR}^2_{\text{lin}_\circ}. \tag{3.21}$$

The upper bound on the right-hand side of (3.21) is the maximum SNR attainable by any linear observer and linear reconstructor in the detection task at hand.

For emission tomography, substituting (3.7) and (3.8) into (3.19) yields

$$\check{\boldsymbol{\Pi}} = \boldsymbol{\Pi} + \boldsymbol{\mathcal{A}}\boldsymbol{\mathcal{K}}_f\boldsymbol{\mathcal{A}}', \tag{3.22}$$

where

$$\boldsymbol{\Pi} \triangleq \mathsf{diag}\big\{\boldsymbol{\mathcal{A}}(\bar{\boldsymbol{f}}_b + \tfrac{1}{2}\bar{\boldsymbol{f}}_s) + \boldsymbol{r}\big\}$$

is the imaging noise covariance for the unconditional mean object and

$$\boldsymbol{\mathcal{K}}_f \triangleq \boldsymbol{\mathcal{K}}_b + \tfrac{1}{2}\boldsymbol{\mathcal{K}}_s$$

is the unconditional covariance of the object. Since $\mathbf{\Pi}$, $\boldsymbol{\mathcal{K}}_b$ and $\boldsymbol{\mathcal{K}}_s$ are positive definite, so is $\check{\mathbf{\Pi}}$ in (3.22), thus the ratio in (3.18) is well-defined provided $\boldsymbol{\mathcal{Z}}'\boldsymbol{w}$ is nonzero.

Substituting (3.6) and (3.22) into (3.21) yields the optimal SNR for the detection task at hand in emission tomography:

$$\mathrm{SNR}^2_{\mathrm{lin}_\circ} = \bar{\boldsymbol{f}}'_s \check{\boldsymbol{\mathcal{F}}} \bar{\boldsymbol{f}}_s, \tag{3.23}$$

where we define

$$\check{\boldsymbol{\mathcal{F}}} \triangleq \boldsymbol{\mathcal{A}}'\check{\mathbf{\Pi}}^{-1}\boldsymbol{\mathcal{A}} = \boldsymbol{\mathcal{A}}'(\mathbf{\Pi} + \boldsymbol{\mathcal{A}}\boldsymbol{\mathcal{K}}_f\boldsymbol{\mathcal{A}}')^{-1}\boldsymbol{\mathcal{A}} \tag{3.24}$$

$$= (\boldsymbol{\mathcal{I}} + \boldsymbol{\mathcal{F}}\boldsymbol{\mathcal{K}}_f)^{-1}\boldsymbol{\mathcal{F}} = \boldsymbol{\mathcal{F}}(\boldsymbol{\mathcal{I}} + \boldsymbol{\mathcal{K}}_f\boldsymbol{\mathcal{F}})^{-1}, \tag{3.25}$$

where $\boldsymbol{\mathcal{I}}$ is the identity operator, and $\boldsymbol{\mathcal{F}} \triangleq \boldsymbol{\mathcal{A}}'\mathbf{\Pi}^{-1}\boldsymbol{\mathcal{A}}$. From (3.24) we derive (3.25) using the “push-through” identity $\boldsymbol{\mathcal{A}}(\boldsymbol{\mathcal{I}}+\boldsymbol{\mathcal{B}}\boldsymbol{\mathcal{A}})^{-1} = (\boldsymbol{\mathcal{I}}+\boldsymbol{\mathcal{A}}\boldsymbol{\mathcal{B}})^{-1}\boldsymbol{\mathcal{A}}$ [58]. The operators $\check{\boldsymbol{\mathcal{F}}}$ and $\boldsymbol{\mathcal{F}}$ have a form analogous to that of the *Fisher information* matrix [132, p.80] for estimation problems that involve a linear system and additive Gaussian noise. In the SKE/BKE case, we have $\check{\boldsymbol{\mathcal{F}}} = \boldsymbol{\mathcal{F}}$.

A simple combination that satisfies the optimality condition (3.20) is $\boldsymbol{\mathcal{Z}} = \boldsymbol{\mathcal{I}}$ (which is not a reconstruction method) and $\boldsymbol{w} = \check{\mathbf{\Pi}}^{-1}\,\bar{\boldsymbol{y}}$, which corresponds to the Hotelling observer (2.28) for detection in the raw measurement domain rather than the reconstruction domain. However, even when we restrict attention to observers that are applied to reconstructed images, usually there are still many ways to satisfy (3.20), as the analysis that follows indicates.

3.3.2 Fisher Observers and Reconstructors

As shown in more detail below, several of the mathematical observers that have been proposed in the literature can achieve the optimal SNR when paired with simple reconstructors that correspond to some power of the Fisher information operator $\check{\boldsymbol{\mathcal{F}}}$ applied to a backprojection of the data. For lack of a better term we refer to this family of reconstruction methods as *Fisher reconstructors* and we allow them to include a regularization component:

$$\boldsymbol{\mathcal{Z}}_q \triangleq \boldsymbol{\mathcal{H}}^{(q)}\boldsymbol{\mathcal{A}}'\check{\boldsymbol{\Pi}}^{-1} = (\check{\boldsymbol{\mathcal{F}}} + \boldsymbol{\mathcal{R}})^{(q)}\boldsymbol{\mathcal{A}}'\check{\boldsymbol{\Pi}}^{-1} \tag{3.26}$$

for some $q \in \mathbb{R}$ and $\boldsymbol{\mathcal{H}} \triangleq \check{\boldsymbol{\mathcal{F}}} + \boldsymbol{\mathcal{R}}$ for some regularization operator $\boldsymbol{\mathcal{R}}$. We use the notation

$$\boldsymbol{\mathcal{H}}^{(q)} \triangleq \begin{cases} \boldsymbol{\mathcal{H}}^{q}, & q \geq 0 \\ (\boldsymbol{\mathcal{H}}^{\dagger})^{-q}, & q < 0. \end{cases}$$

where "†" denotes a pseudo-inverse[1]. Typically the regularizer $\boldsymbol{\mathcal{R}}$ is chosen so that $\boldsymbol{\mathcal{H}}$ is nonsingular. The reason we allow $\boldsymbol{\mathcal{H}}$ to be singular in the analysis that follows is to accommodate the unregularized case $\boldsymbol{\mathcal{R}} = \mathbf{0}$, since $\check{\boldsymbol{\mathcal{F}}}$ itself may be singular.

The estimated image produced by the Fisher reconstructor (3.26) is

$$\hat{\boldsymbol{f}} = \boldsymbol{\mathcal{Z}}_q \boldsymbol{y} = \boldsymbol{\mathcal{H}}^{(q)}\boldsymbol{\mathcal{A}}'\check{\boldsymbol{\Pi}}^{-1}\boldsymbol{y} = \boldsymbol{\mathcal{H}}^{(q)}\boldsymbol{\mathcal{A}}'(\boldsymbol{\Pi} + \boldsymbol{\mathcal{A}}\boldsymbol{\mathcal{K}}_f\boldsymbol{\mathcal{A}}')^{-1}\boldsymbol{y},$$

which is a kind of weighted backprojection with a (perhaps somewhat unusual) postfilter. For $q < 0$, this postfilter is something like a regularized deconvolver. For $q = -1$ in particular, the Fisher reconstructor yields

$$\hat{\boldsymbol{f}} = \boldsymbol{\mathcal{H}}^{\dagger}\boldsymbol{\mathcal{A}}'\check{\boldsymbol{\Pi}}^{-1}\boldsymbol{y}, \tag{3.27}$$

so it is equivalent to the QPWLS estimator from (3.13) with $\boldsymbol{W} = \check{\boldsymbol{\Pi}}^{-1}$.

Interestingly, for any Fisher reconstructor, even an unregularized one, there is a corresponding linear observer that achieves optimal detection performance. We refer to this observer family as *Fisher observers*:

$$\boldsymbol{w}_p \triangleq \boldsymbol{\mathcal{H}}^{(p)}\bar{\boldsymbol{f}}_s \tag{3.28}$$

---

[1] We follow Definition 6.2.1 of the pseudo-inverse in [53]. An immediate result of this definition is that for nonsingular matrices the pseudo-inverse is the same as the inverse (Theorem 6.2.13 in [53].) In this notation, we have $\boldsymbol{\mathcal{H}}^{(-1)} = \boldsymbol{\mathcal{H}}^{\dagger} = \boldsymbol{\mathcal{H}}^{-1}$ whenever $\boldsymbol{\mathcal{H}}$ is nonsingular.

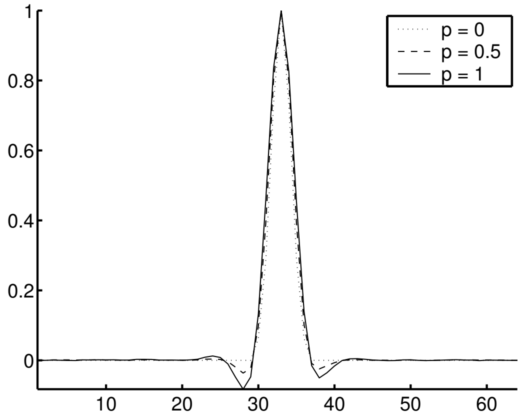

Figure 3.1: Profiles through the center of (normalized) Fisher observer templates $\boldsymbol{w}_p$ for $p = 0, 0.5, 1$. Note that $\boldsymbol{w}_0 = \bar{\boldsymbol{f}}_s$.

for some $p \in \mathbb{R}$. For a reconstructor of the family (3.26) and an observer of the family (3.28) we have

$$\boldsymbol{\mathcal{Z}}_q' \boldsymbol{w}_p = \check{\boldsymbol{\Pi}}^{-1} \boldsymbol{\mathcal{A}} \boldsymbol{\mathcal{H}}^{(q)} \boldsymbol{\mathcal{H}}^{(p)} \bar{\boldsymbol{f}}_s. \tag{3.29}$$

By comparing (3.29) with (3.20) for $\bar{\boldsymbol{y}} = \boldsymbol{\mathcal{A}} \bar{\boldsymbol{f}}_s$ we see that the choice $p = -q$ leads to the optimal SNR, *i.e.*, the observer $\boldsymbol{w}_{-q}$ achieves optimal SNR when applied to images produced by the reconstructor $\boldsymbol{\mathcal{Z}}_q$. This is true even for $\boldsymbol{\mathcal{R}} = \mathbf{0}$ and singular $\check{\boldsymbol{\mathcal{F}}}$, as we can show by using the fact that

$$\boldsymbol{\mathcal{B}}(\boldsymbol{\mathcal{B}}'\boldsymbol{\mathcal{B}})^{(q)}(\boldsymbol{\mathcal{B}}'\boldsymbol{\mathcal{B}})^{(-q)} = \boldsymbol{\mathcal{B}} \tag{3.30}$$

for any $\boldsymbol{\mathcal{B}}$, which we obtain by singular value decomposition of $\boldsymbol{\mathcal{B}}$. Substituting $\boldsymbol{\mathcal{B}} = \check{\boldsymbol{\Pi}}^{-1/2} \boldsymbol{\mathcal{A}}$ in (3.30) yields the desired result. Fig. 3.1 shows template profiles for some of the observers in (3.28). The profile shape for $p = 0.5$ especially is reminiscent of those estimated from human observers (*e.g.*, see [4]).

Fisher reconstructors with $\boldsymbol{\mathcal{R}} = \mathbf{0}$ appear to be largely devoid of regularization[2]. How-

---

[2] For $\boldsymbol{\mathcal{R}} = \mathbf{0}$ and $q > -1$, one could construe the Fisher reconstructor as being marginally regularized since it entails somewhat "less deconvolution" than the unregularized WLS reconstructor. However, this type of "regularization" does not improve the condition number in the case of singular $\check{\boldsymbol{\mathcal{F}}}$, and it is unlike most regularization methods described in the literature.

ever, even for these reconstructors, there are corresponding observers that can achieve optimal performance in the detection task considered here. As shown in the following sections, some well-known observers from the literature fall within this category. This means that, if the reconstructed images are meant to be interpreted by these observers, *regularization is not essential* in the reconstruction process, at least for the simple detection task at hand.

3.3.3 Hotelling Observer

Substituting the moments of $\hat{\boldsymbol{f}}$ from (3.10) and (3.11) into the HO template from (2.28) yields

$$\boldsymbol{w}_{\mathrm{HO}} = (\boldsymbol{\mathcal{Z}}\check{\boldsymbol{\Pi}}\boldsymbol{\mathcal{Z}}')^{\dagger}\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}} \tag{3.31}$$

and thus

$$\boldsymbol{\mathcal{Z}}'\boldsymbol{w}_{\mathrm{HO}} = \boldsymbol{\mathcal{Z}}'(\boldsymbol{\mathcal{Z}}\check{\boldsymbol{\Pi}}\boldsymbol{\mathcal{Z}}')^{\dagger}\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}} = \check{\boldsymbol{\Pi}}^{-1/2}\mathcal{P}_{\check{\boldsymbol{\Pi}}^{1/2}\boldsymbol{\mathcal{Z}}'}(\check{\boldsymbol{\Pi}}^{-1/2}\,\bar{\boldsymbol{y}}), \tag{3.32}$$

where $\mathcal{P}_{\check{\boldsymbol{\Pi}}^{1/2}\boldsymbol{\mathcal{Z}}'}(\cdot)$ denotes the orthogonal projection of a vector onto $\mathcal{R}_{\check{\boldsymbol{\Pi}}^{1/2}\boldsymbol{\mathcal{Z}}'}$, the range space of $\check{\boldsymbol{\Pi}}^{1/2}\boldsymbol{\mathcal{Z}}'$. By comparing (3.32) with (3.20) we find that the HO achieves optimal SNR for any $\boldsymbol{\mathcal{Z}}$ that satisfies $\check{\boldsymbol{\Pi}}^{-1/2}\,\bar{\boldsymbol{y}} \in \mathcal{R}_{\check{\boldsymbol{\Pi}}^{1/2}\boldsymbol{\mathcal{Z}}'}$ or, equivalently,

$$\check{\boldsymbol{\Pi}}^{-1}\,\bar{\boldsymbol{y}} \in \mathcal{R}_{\boldsymbol{\mathcal{Z}}'}. \tag{3.33}$$

There are a multitude of choices of $\boldsymbol{\mathcal{Z}}$ that satisfy the mild condition in (3.33). For example, any $\boldsymbol{\mathcal{Z}}_q$ of the Fisher reconstructor family (3.26) satisfies (3.33), as long as $\bar{\boldsymbol{y}}$ is in the range space of $\boldsymbol{\mathcal{A}}$ (*e.g.*, $\bar{\boldsymbol{y}} = \boldsymbol{\mathcal{A}}\,\bar{\boldsymbol{f}}_s$). This is true even for a $\boldsymbol{\mathcal{Z}}_q$ with $\boldsymbol{\mathcal{R}} = \mathbf{0}$ and singular $\check{\boldsymbol{\mathcal{F}}}$, following an argument similar to the previous section.

Thus the HO achieves the ideal SNR for any Fisher reconstructor, regularized or not. This result is consistent with the fact that linear transformations of the data do not affect the performance of the optimal linear observer [44, §10.2] (except when the transformation operator does not have a right inverse, in which case performance degrades). Note that for $\bar{\boldsymbol{y}} = \boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_s$, $\boldsymbol{\mathcal{R}} = \mathbf{0}$ and $q \neq 0$ the HO template in (3.31) becomes $\boldsymbol{w}_{\mathrm{HO}} = \check{\boldsymbol{\mathcal{F}}}^{(-q)}\bar{\boldsymbol{f}}_s$, which is the Fisher observer template in (3.28) with $p = -q$ and the optimality of $\boldsymbol{\mathcal{Z}}_q$ for this observer follows by the previous section.

3.3.4 Prewhitening Observer

Substituting the moments of $\hat{\boldsymbol{f}}$ from (3.10) and (3.11) into the PW template from (2.29) yields

$$\boldsymbol{w}_{\mathrm{PW}} = (\boldsymbol{\mathcal{Z}}\boldsymbol{\Pi}_0\boldsymbol{\mathcal{Z}}')^{\dagger}\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}}$$

and thus

$$\boldsymbol{\mathcal{Z}}'\boldsymbol{w}_{\mathrm{PW}} = \boldsymbol{\mathcal{Z}}'(\boldsymbol{\mathcal{Z}}\boldsymbol{\Pi}_0\boldsymbol{\mathcal{Z}}')^{\dagger}\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}} = \boldsymbol{\Pi}_0^{-1/2}\mathcal{P}_{\boldsymbol{\Pi}_0^{1/2}\boldsymbol{\mathcal{Z}}'}(\boldsymbol{\Pi}_0^{-1/2}\,\bar{\boldsymbol{y}}). \tag{3.34}$$

By comparing (3.34) with (3.20) we find that the PW observer achieves optimal SNR for any $\boldsymbol{\mathcal{Z}}$ that satisfies

$$\mathcal{P}_{\boldsymbol{\Pi}_0^{1/2}\boldsymbol{\mathcal{Z}}'}(\boldsymbol{\Pi}_0^{-1/2}\,\bar{\boldsymbol{y}}) \propto \boldsymbol{\Pi}_0^{1/2}\check{\boldsymbol{\Pi}}^{-1}\,\bar{\boldsymbol{y}},$$

which in turn implies that $\boldsymbol{\Pi}_0^{-1/2}\,\bar{\boldsymbol{y}} - c\boldsymbol{\Pi}_0^{1/2}\check{\boldsymbol{\Pi}}^{-1}\,\bar{\boldsymbol{y}}$ must be orthogonal to $\mathcal{R}_{\boldsymbol{\Pi}_0^{1/2}\boldsymbol{\mathcal{Z}}'}$ for any constant c. This finally leads to the requirement that

$$\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}} \propto \boldsymbol{\mathcal{Z}}\boldsymbol{\Pi}_0\check{\boldsymbol{\Pi}}^{-1}\,\bar{\boldsymbol{y}},$$

i.e., that the sinograms $\bar{\boldsymbol{y}}$ and $\boldsymbol{\Pi}_0\check{\boldsymbol{\Pi}}^{-1}\,\bar{\boldsymbol{y}}$ yield the same reconstructed image but for a scaling constant. Thus, in general there is no linear reconstructor $\boldsymbol{\mathcal{Z}}$ that leads to optimal SNR for the PW observer.

An exception to this is the SKE case, where $\check{\boldsymbol{\Pi}} = \boldsymbol{\Pi}_0$ and the PW observer is the same as the HO, so it can achieve optimal SNR for infinitely many choices of $\boldsymbol{\mathcal{Z}}$. The minimal dependence of the SNR on $\boldsymbol{\mathcal{Z}}$ for the HO is consistent with the observation of Qi *et al.*

that performance of the PW observer in the SKE task is independent of smoothing method in the MAP case [106].

3.3.5 Non-Prewhitening Observer

Substituting the moments of $\hat{\boldsymbol{f}}$ from (3.10) into the NPW template from (2.30) yields

$$\boldsymbol{w}_{\text{NPW}} = \boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}} \tag{3.35}$$

and thus by (3.20) the NPW observer achieves optimal SNR for any $\boldsymbol{\mathcal{Z}}$ that satisfies

$$\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}} \propto \check{\boldsymbol{\Pi}}^{-1}\,\bar{\boldsymbol{y}}\,. \tag{3.36}$$

For a $\boldsymbol{\mathcal{Z}}_q$ of the Fisher reconstructor family (3.26) and $\bar{\boldsymbol{y}} = \boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_s$, we have

$$\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}} = \check{\boldsymbol{\Pi}}^{-1}\boldsymbol{\mathcal{A}}\boldsymbol{\mathcal{H}}^{(2q)}\check{\boldsymbol{\mathcal{F}}}\bar{\boldsymbol{f}}_s,$$

so the optimal SNR is achieved when $\boldsymbol{\mathcal{R}} = 0$ and $q = -1/2$. This corresponds to the somewhat unusual Fisher reconstructor

$$\boldsymbol{\mathcal{Z}} = \check{\boldsymbol{\mathcal{F}}}^{(-1/2)}\boldsymbol{\mathcal{A}}'\check{\boldsymbol{\Pi}}^{-1}. \tag{3.37}$$

Whether there are other solutions that satisfy (3.36) is an open problem. (There is also the choice of $\boldsymbol{\mathcal{Z}} = \check{\boldsymbol{\Pi}}^{-1/2}$, which is not a reconstruction method. It is equivalent to the Hotelling observer for sinogram-based detection.) Note that for the reconstructor (3.37) the NPW template in (3.35) becomes $\boldsymbol{w}_{\text{NPW}} = \check{\boldsymbol{\mathcal{F}}}^{1/2}\bar{\boldsymbol{f}}_s$, which is the Fisher observer template in (3.28) with $p = 1/2$.

The reconstructor in (3.37) yields the estimated image

$$\hat{\boldsymbol{f}} = \check{\boldsymbol{\mathcal{F}}}^{(-1/2)}\boldsymbol{\mathcal{A}}'\check{\boldsymbol{\Pi}}^{-1}\boldsymbol{y} = \check{\boldsymbol{\mathcal{F}}}^{1/2}(\check{\boldsymbol{\mathcal{F}}}^{\dagger}\boldsymbol{\mathcal{A}}'\check{\boldsymbol{\Pi}}^{-1}\boldsymbol{y}), \tag{3.38}$$

which is the unregularized WLS estimate, given in (3.27) for $\boldsymbol{\mathcal{R}} = \mathbf{0}$, followed by the unusual shift-variant post-filter $\check{\boldsymbol{\mathcal{F}}}^{1/2}$. This estimator is impractical for two reasons. Firstly, even if $\boldsymbol{\mathcal{A}}$ happens to have full rank, $\check{\boldsymbol{\mathcal{F}}}$ is usually very ill-conditioned, so computing the WLS solution (3.27) will require a multitude of iterations for any practical iterative algorithm. Secondly, the shift-variant post-filter $\check{\boldsymbol{\mathcal{F}}}^{1/2}$ would be computationally intensive to implement for typical problem sizes.

3.3.6 Region-of-Interest Observer

From (2.31) the template of the ROI observer is given by

$$\boldsymbol{w}_{\mathrm{ROI}} = \bar{\boldsymbol{f}}_s.$$

Then by (3.20) the ROI observer achieves optimal SNR for any $\boldsymbol{\mathcal{Z}}$ that satisfies

$$\boldsymbol{\mathcal{Z}}' \bar{\boldsymbol{f}}_s \propto \check{\boldsymbol{\Pi}}^{-1} \bar{\boldsymbol{y}}\,.$$

Since the ROI template corresponds to the Fisher observer in (3.28) with $p = 0$, the optimal SNR for $\bar{\boldsymbol{y}} = \boldsymbol{\mathcal{A}} \bar{\boldsymbol{f}}_s$ is achieved by the Fisher reconstructor with $q = 0$, *i.e.*,

$$\boldsymbol{\mathcal{Z}} = \boldsymbol{\mathcal{A}}' \check{\boldsymbol{\Pi}}^{-1}. \tag{3.39}$$

Curiously, in this case

$$\hat{\boldsymbol{f}} = \boldsymbol{\mathcal{A}}' \check{\boldsymbol{\Pi}}^{-1} \boldsymbol{y}.$$

This is a very blurry estimate of $\boldsymbol{f}$, being simply unfiltered backprojection. Yet for the ROI observer it is optimal, and no amount of deconvolution will improve the SNR for this detection task, which is an indication that the task is too simple. The optimality of this blurry estimate is consistent with the demonstration in Qi *et al.* of the ROI observer (for a penalized-likelihood reconstructor with $\boldsymbol{\mathcal{R}} = \beta \boldsymbol{\mathcal{I}}$) approaching the PW observer's performance as $\beta \to \infty$ [106].

Alternatively to (2.31), the ROI template may be defined as containing 1's wherever $\bar{\boldsymbol{f}}_s$ is non-zero and 0's everywhere else. In this case, the optimal $\boldsymbol{\mathcal{Z}}$ would be the backprojector

Observer	q	$\boldsymbol{\mathcal{R}}$	Best estimator	Interpretation
Hotelling	$\mathbb{R}$	$\succeq \mathbf{0}$	$(\check{\boldsymbol{\mathcal{F}}} + \boldsymbol{\mathcal{R}})^{(q)} \boldsymbol{\mathcal{A}}' \check{\boldsymbol{\Pi}}^{-1} \boldsymbol{y}$	Any Fisher reconstructor (*e.g.*, WLS)
NPW	$-1/2$	$\mathbf{0}$	$\check{\boldsymbol{\mathcal{F}}}^{(-1/2)} \boldsymbol{\mathcal{A}}' \check{\boldsymbol{\Pi}}^{-1} \boldsymbol{y}$	Partly deconvolved backprojection
ROI	0	-	$\boldsymbol{\mathcal{A}}' \check{\boldsymbol{\Pi}}^{-1} \boldsymbol{y}$	Backprojection

Table 3.1: Known-location observer models and the corresponding optimal Fisher reconstructors. These reconstructors allow the respective observers to achieve the optimal SNR in known-location detection tasks.

in (3.39) followed by a diagonal operator that weighs the reconstructed image by the shape of $\bar{\boldsymbol{f}}_s$, *i.e.*, $\boldsymbol{\mathcal{Z}} = \mathsf{diag}\{\bar{\boldsymbol{f}}_s\}\, \boldsymbol{\mathcal{A}}' \check{\boldsymbol{\Pi}}^{-1}$. This weighting does not improve the resolution of the backprojection, so for either definition of the ROI template the conclusion remains that optimal performance is achieved with a very blurry estimator.

3.3.7 Summary

For three of the specific observer models considered above, at least one reconstructor of the Fisher family (3.26) was found to achieve the highest SNR possible for linear observers. Table 3.1 summarizes these findings, which show that the linear reconstructors $\boldsymbol{\mathcal{Z}}$ that yield optimal detection performance for the HO, NPW, and ROI models need not include any form of regularization. This is true even if the system operator $\boldsymbol{\mathcal{A}}$ is a matrix with less than full column rank. We conclude on theoretical grounds that regularization is not absolutely essential for the HO, NPW, and ROI observer models in the task of detecting a statistically varying signal of known location on a statistically varying background.

The optimality of blurry backprojections with respect to detection performance for the HO, NPW, and ROI models implies that spatial resolution is not important for these known-location tasks. This is consistent with the findings of Wagner *et al.*, who attempted to optimize imaging systems with respect to performance in such a task, only to find that a pinhole of very large size would be optimal [136]. Although Wagner *et al.* considered a SKE/BKE task, our analysis above shows that spatial resolution does not play an important role even in the presence of signal and background variability.

Furthermore, there is a strong dependence of the optimal reconstruction method on the specific observer model considered. This implies that there is no universally optimal reconstruction method, even for the simple detection task considered here. As discussed in chapter 2, the apparent premise of human-observer studies in the literature is that humans do not perform as well as the ideal observer. Therefore, the fact that there exist "simple" reconstruction methods that allow the HO, NPW, and ROI observer models to achieve the ideal linear-detection SNR (which is also the overall ideal SNR for SKE tasks) suggests that these observer models, the tasks, or both are somehow inappropriate. We proceed with analyzing the effect of regularization on the performance of channelized observer models, which are known to correlate well with human observers for known-location tasks.

3.4 Channelized Linear Observers

3.4.1 Overview

We now consider linear observers that extract their feature vector from the reconstructed image through a noisy, dimension-reducing operation, as in (2.32). Following convention, we assume that the internal noise vector $\boldsymbol{\varepsilon}_{\text{int}}$ is zero-mean Gaussian, with some covariance matrix $\boldsymbol{\Pi}_{\text{int}}$.

For a linear reconstruction method $\boldsymbol{\mathcal{Z}}$, combining the reconstruction moments from (3.10) and (3.11) with the channelized linear discriminant in (2.33) yields the following expressions for the moments of t_{ch}:

$$\mathsf{E}[t_{\text{ch}}|H_1] - \mathsf{E}[t_{\text{ch}}|H_0] = \boldsymbol{w}'\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}} \tag{3.40}$$

$$\mathsf{Var}\{t_{\text{ch}}|H_i\} = \boldsymbol{w}'\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\boldsymbol{\Pi}_i\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}\boldsymbol{w} + \boldsymbol{w}'\boldsymbol{\Pi}_{\text{int}}\boldsymbol{w}, \quad i = 0, 1. \tag{3.41}$$

Combining expressions (3.40) and (3.41) with (2.19) yields the SNR of the channelized observer:

$$\text{SNR}_{\text{ch}}^2 = \frac{(\boldsymbol{w}'\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}})^2}{\boldsymbol{w}'\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\check{\boldsymbol{\Pi}}\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}\boldsymbol{w} + \boldsymbol{w}'\boldsymbol{\Pi}_{\text{int}}\boldsymbol{w}}. \tag{3.42}$$

We would like to find a reconstructor $\boldsymbol{\mathcal{Z}}$ that optimizes the SNR in (3.42). However, if $\boldsymbol{\Pi}_{\text{int}} \neq \boldsymbol{0}$ is independent of $\boldsymbol{\mathcal{Z}}$, then a $\boldsymbol{\mathcal{Z}}$ of infinitely large norm would be optimal. Thus one would need some constraint on $\boldsymbol{\mathcal{Z}}$ to optimize the SNR in its most general form (3.42), but it is unclear what constraints would be suitable.

A special case occurs if we assume that the covariance of the internal noise in the channels is proportional to the covariance of $\boldsymbol{\mathcal{C}}'\hat{\boldsymbol{f}}$ due to imaging noise and object variability. Specifically, we assume

$$\boldsymbol{\Pi}_{\text{int}} = c_{\text{int}}\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}'\check{\boldsymbol{\Pi}}\boldsymbol{\mathcal{Z}}\boldsymbol{\mathcal{C}} \tag{3.43}$$

for some constant of proportionality c_{int}, which includes the special case of no internal noise for $c_{\text{int}} = 0$. The model (3.43) is similar to the one proposed in [1], except that (3.43) does not presume independence of the internal noise in individual channels. Under the assumption (3.43) the SNR of the channelized linear observer in (3.42) becomes

$$\text{SNR}_{\text{ch}}^2 = \frac{1}{1+c_{\text{int}}}\frac{(\boldsymbol{w}'\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}})^2}{\boldsymbol{w}'\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\check{\boldsymbol{\Pi}}\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}\boldsymbol{w}} \leq \frac{1}{1+c_{\text{int}}}\,\bar{\boldsymbol{y}}'\,\check{\boldsymbol{\Pi}}^{-1}\,\bar{\boldsymbol{y}} = \frac{\text{SNR}_{\text{lin}_\circ}^2}{1+c_{\text{int}}}. \tag{3.44}$$

The SNR upper bound in (3.44) is achieved for any $\boldsymbol{\mathcal{Z}}$ that satisfies

$$\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}\boldsymbol{w} \propto \check{\boldsymbol{\Pi}}^{-1}\,\bar{\boldsymbol{y}}, \tag{3.45}$$

similarly to the non-channelized version in (3.18) and (3.20).

3.4.2 Channelized Hotelling Observer

Combining the moments of $\hat{\boldsymbol{f}}$ from (3.10) and (3.11) with the channelized feature vector from (2.32) and substituting into the CHO template from (2.34) yields

$$\boldsymbol{w}_{\text{CHO}} = (\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\check{\boldsymbol{\Pi}}\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}} + \boldsymbol{\Pi}_{\text{int}})^{\dagger}\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}}\,. \tag{3.46}$$

(The filters involved in $\boldsymbol{\mathcal{C}}$ typically correspond to distinct frequency bands. If this is the case, then the covariance $\boldsymbol{\Pi}_{\hat{c}}$ can be assumed to be nonsingular even when $\boldsymbol{\Pi}_{\text{int}} = \mathbf{0}$, so its pseudo-inverse is the same as its inverse. Nevertheless, we use a pseudo-inverse in the interest of generality.) By substituting the CHO template from (3.46) into the SNR expression (3.42), we find this observer's SNR:

$$\text{SNR}^2_{\text{CHO}} = \bar{\boldsymbol{y}}' \, \boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}(\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\check{\boldsymbol{\Pi}}\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}} + \boldsymbol{\Pi}_{\text{int}})^{\dagger}\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}} \, \bar{\boldsymbol{y}} \, . \tag{3.47}$$

In the special case of the internal noise model (3.43) we have

$$\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}\boldsymbol{w}_{\text{CHO}} \propto \boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}(\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\check{\boldsymbol{\Pi}}\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}})^{\dagger}\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}} \, \bar{\boldsymbol{y}} = \check{\boldsymbol{\Pi}}^{-1/2}\mathcal{P}_{\check{\boldsymbol{\Pi}}^{1/2}\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}}(\check{\boldsymbol{\Pi}}^{-1/2} \, \bar{\boldsymbol{y}}). \tag{3.48}$$

By comparing (3.48) to the optimality condition (3.45), we find that the SNR upper bound in (3.44) is achieved when $\boldsymbol{\mathcal{Z}}$ satisfies $\check{\boldsymbol{\Pi}}^{-1/2} \, \bar{\boldsymbol{y}} \in \mathcal{R}_{\check{\boldsymbol{\Pi}}^{1/2}\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}}$ or, equivalently,

$$\check{\boldsymbol{\Pi}}^{-1} \, \bar{\boldsymbol{y}} \in \mathcal{R}_{\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}}, \tag{3.49}$$

which is similar to (3.33). To find a $\boldsymbol{\mathcal{Z}}$ that satisfies (3.49), consider a reconstructor of the form

$$\boldsymbol{\mathcal{Z}} = \boldsymbol{\mathcal{W}}\boldsymbol{\mathcal{C}}(\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{W}}\boldsymbol{\mathcal{C}})^{-1}\boldsymbol{\mathcal{G}}'\boldsymbol{\mathcal{A}}'\check{\boldsymbol{\Pi}}^{-1}, \tag{3.50}$$

where $\boldsymbol{\mathcal{W}}$ is any image-domain weighting operator and $\boldsymbol{\mathcal{G}}$ is a mapping from $\mathbb{C}^M$ to object space. For the $\boldsymbol{\mathcal{Z}}$ in (3.50), we have $\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}} = \check{\boldsymbol{\Pi}}^{-1}\boldsymbol{\mathcal{A}}\boldsymbol{\mathcal{G}}$, so the optimality condition (3.49) is satisfied if

$$\bar{\boldsymbol{y}} \in \mathcal{R}_{\boldsymbol{\mathcal{A}}\boldsymbol{\mathcal{G}}}. \tag{3.51}$$

As long as $\boldsymbol{\mathcal{G}}$ can be chosen to satisfy (3.51), the unregularized reconstructor $\boldsymbol{\mathcal{Z}}$ given in (3.50) allows the CHO to achieve the optimal SNR. For emission tomography, where $\bar{\boldsymbol{y}} = \boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_s$, there is always a way to choose $\boldsymbol{\mathcal{G}}$ to satisfy (3.51); choosing $\boldsymbol{\mathcal{G}}$ so that one of its "columns" is proportional to $\bar{\boldsymbol{f}}_s$ is the most obvious way.

The rather unconventional family of reconstructors in (3.50), although ensuring optimal SNR for the CHO under the internal noise model (3.43), does not produce what we usually consider to be reconstructed images. In essence, the HO template for detection in the raw measurement domain ($\boldsymbol{w} = \check{\mathbf{\Pi}}^{-1}\,\bar{\boldsymbol{y}}$) is hidden in one of the rows of such a reconstructor. Thus it is of interest to know whether reconstructors that are more conventional than (3.50) can be combined with the CHO to achieve SNR values close to the optimal and, if so, whether these reconstructors must include regularization. As shown below, such insight can be gained by using approximations to the SNR expression in (3.47) that are based on local Fourier-domain analysis. Such approximations also facilitate analysis of internal noise models other than (3.43). In the following, we first provide a brief review of local Fourier-domain approximations and we then use them to analyze the SNR of the CHO for one unregularized and one regularized reconstructor example.

3.4.3 Local Fourier Analysis of Detection Performance

For the purposes of the Fourier analysis that follows, we will assume to have a discrete representation of the object[3] $\boldsymbol{f} \in \mathbb{R}^{n_\mathrm{p}}$. Let $\boldsymbol{\mathcal{U}}$ be a discrete Fourier operator, mapping an object in $\mathbb{R}^{n_\mathrm{p}}$ to some vector in $\mathbb{C}^{n_\mathrm{p}}$ known as the *spectrum* of the object. The Fourier operator $\boldsymbol{\mathcal{U}}$ is linear and, due to the shift property of the Fourier transform, it can be defined through its action on an object $\boldsymbol{e}_0$ that consists of an impulse centered at the origin:

$$\boldsymbol{\mathcal{U}}\boldsymbol{e}_0 = \frac{1}{\sqrt{n_\mathrm{p}}}\mathbf{1}, \tag{3.52}$$

where $\mathbf{1}$ is the vector of n_p ones. Without loss of generality, we choose the origin of the Fourier transform to coincide with the location where the signal $\bar{\boldsymbol{f}}_s$ is centered. (The $1/\sqrt{n_\mathrm{p}}$ factor results simply from using the orthonormal form of the DFT.)

---

[3] Nevertheless, it is possible to generalize this analysis for an infinite-dimensional object space.

Let $\boldsymbol{\lambda} \in \mathbb{C}^{n_{\mathrm{p}}}$ be the *local frequency response* of the Fisher information operator $\boldsymbol{\mathcal{F}}$ at the origin, defined by

$$\boldsymbol{\lambda} \triangleq \sqrt{n_{\mathrm{p}}}\,\boldsymbol{\mathcal{U}}\boldsymbol{\mathcal{F}}\boldsymbol{e}_0. \tag{3.53}$$

In general, the operator $\boldsymbol{\mathcal{F}}$ is shift-variant, so its behavior when applied to an object other the impulse $\boldsymbol{e}_0$ cannot be fully defined through $\boldsymbol{\lambda}$. However, if the operator $\boldsymbol{\mathcal{F}}$ is applied to an object that is contained within a small spatial extent around the origin, we can approximate the behavior of $\boldsymbol{\mathcal{F}}$ using $\boldsymbol{\lambda}$, *i.e.*, we can approximate $\boldsymbol{\mathcal{F}}$ as shift-invariant within the spatial extent of the object. The closer the object comes to resembling an impulse, the greater the accuracy of this approximation. Thus we approximate the operator $\boldsymbol{\mathcal{F}}$ within a small spatial extent around the origin as

$$\boldsymbol{\mathcal{F}} \approx \boldsymbol{\mathcal{U}}^{-1}\boldsymbol{\Lambda}\,\boldsymbol{\mathcal{U}}, \tag{3.54}$$

where $\boldsymbol{\Lambda} \triangleq \mathsf{diag}\{\boldsymbol{\lambda}\}$. The elements of $\boldsymbol{\lambda}$, λ_k, $k = 1, \ldots, n_{\mathrm{p}}$, contain the local frequency response of $\boldsymbol{\mathcal{F}}$. Since $\boldsymbol{\mathcal{F}}$ is symmetric nonnegative definite, we force the λ_k's to be real and nonnegative by discarding imaginary parts and setting negatives to zero. Locally shift-invariant approximations of $\boldsymbol{\mathcal{F}}$ have proven to be useful and accurate in several cases [19, 36, 38, 40, 106, 108, 109, 124].

We also use a locally shift-invariant approximation of the object covariance:

$$\boldsymbol{\mathcal{K}}_f \approx \boldsymbol{\mathcal{U}}^{-1}\boldsymbol{N}\,\boldsymbol{\mathcal{U}}, \tag{3.55}$$

where $\boldsymbol{N} \triangleq \mathsf{diag}\{\nu_k, k = 1, \ldots, n_{\mathrm{p}}\}$. The ν_k's contain the local frequency response of $\boldsymbol{\mathcal{K}}_f$ (*i.e.*, the *local power spectrum* of the object $\boldsymbol{f}$) around the location of $\bar{\boldsymbol{f}}_s$. Although it may be reasonable to approximate the background $\boldsymbol{f}_b$ as locally stationary within the spatial extent of the signal, approximating the signal $\boldsymbol{f}_s$ itself as locally stationary would be less reasonable. That is, the approximation in (3.55) is much more accurate when the signal $\boldsymbol{f}_s$ is deterministic or when the signal covariance $\boldsymbol{\mathcal{K}}_s$ is negligible in comparison to the background covariance $\boldsymbol{\mathcal{K}}_b$. Therefore, the Fourier analysis that follows applies to detection tasks where the target signal is somewhat variable in its shape but remains weak and spatially localized. It does not apply to tasks with significant signal magnitude and/or width variability.

Using the approximations (3.54) and (3.55), we can start from (3.25) to derive the following locally shift-invariant approximation for the Fisher information operator $\check{\boldsymbol{\mathcal{F}}}$:

$$\begin{aligned} \check{\boldsymbol{\mathcal{F}}} &= \boldsymbol{\mathcal{F}}(\boldsymbol{\mathcal{I}} + \boldsymbol{\mathcal{K}}_f \boldsymbol{\mathcal{F}})^{-1} \\ &\approx \boldsymbol{\mathcal{U}}^{-1} \boldsymbol{\Lambda} (\boldsymbol{I} + \boldsymbol{N} \boldsymbol{\Lambda})^{-1} \boldsymbol{\mathcal{U}} = \boldsymbol{\mathcal{U}}^{-1} \check{\boldsymbol{\Lambda}} \, \boldsymbol{\mathcal{U}}, \end{aligned} \tag{3.56}$$

where $\check{\boldsymbol{\Lambda}} \triangleq \mathsf{diag}\{\check{\lambda}_k, k = 1, \ldots, n_{\mathrm{p}}\}$ and

$$\check{\lambda}_k \triangleq \frac{\lambda_k}{1 + \nu_k \lambda_k}, \tag{3.57}$$

which reduces to $\check{\boldsymbol{\Lambda}} = \boldsymbol{\Lambda}$ in the SKE/BKE case.

As in [41, 104, 145], we also use the fact that $\boldsymbol{\mathcal{C}}$ is a collection of filters to get its frequency-domain representation. Let $\boldsymbol{t}^m \in \mathbb{C}^{n_{\mathrm{p}}}$ denote the frequency response of the mth bandpass filter. Then the mth operator in $\boldsymbol{\mathcal{C}}$ has the form

$$\boldsymbol{\mathcal{C}}_m = \boldsymbol{\mathcal{U}}^{-1} \, \mathsf{diag}\{\boldsymbol{t}^m\} \boldsymbol{\mathcal{U}} \boldsymbol{e}_0 \tag{3.58}$$

and, combining (3.58) with (3.52), we get

$$\boldsymbol{\mathcal{C}} = \boldsymbol{\mathcal{U}}^{-1} \boldsymbol{T}, \qquad \boldsymbol{T} = \frac{1}{\sqrt{n_{\mathrm{p}}}} \left[\boldsymbol{t}^1 \; \ldots \; \boldsymbol{t}^M \right]. \tag{3.59}$$

3.4.4 CHO and Fisher Reconstructors

For an unregularized Fisher reconstructor $\boldsymbol{\mathcal{Z}}_q$, given by (3.26) with $\boldsymbol{\mathcal{R}} = \boldsymbol{0}$, we can use the Fourier decompositions (3.56) and (3.59) along with (3.6) to obtain the following approximations for the moments of the channel outputs:

$$\begin{aligned}\mathsf{E}[\hat{\boldsymbol{c}}|H_1] - \mathsf{E}[\hat{\boldsymbol{c}}|H_0] &= \boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}_q\bar{\boldsymbol{y}} \\ &= \boldsymbol{\mathcal{C}}'\check{\boldsymbol{\mathcal{F}}}^{(1+q)}\bar{\boldsymbol{f}}_s \approx \boldsymbol{T}'\check{\boldsymbol{\Lambda}}^{(1+q)}\boldsymbol{X} \qquad (3.60)\\ \boldsymbol{\Pi}_{\hat{\boldsymbol{c}}} &= \boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}_q\check{\boldsymbol{\Pi}}\boldsymbol{\mathcal{Z}}_q'\boldsymbol{\mathcal{C}} + \boldsymbol{\Pi}_{\text{int}} \\ &= \boldsymbol{\mathcal{C}}'\check{\boldsymbol{\mathcal{F}}}^{(1+2q)}\boldsymbol{\mathcal{C}} + \boldsymbol{\Pi}_{\text{int}} \approx \boldsymbol{T}'\check{\boldsymbol{\Lambda}}^{(1+2q)}\boldsymbol{T} + \boldsymbol{\Pi}_{\text{int}}, \qquad (3.61)\end{aligned}$$

where $\boldsymbol{X} = \boldsymbol{\mathcal{U}}\bar{\boldsymbol{f}}_s$ is the spectrum of the mean target signal. The approximation (3.60) assumes that $\bar{\boldsymbol{f}}_s$ is highly localized in space. The accuracy of the approximation (3.61) depends on how localized in space the channel responses are. However, when $\check{\boldsymbol{\Lambda}}$ is positive definite, the expression (3.61) is exact for $q = -1/2$, so it may also be reasonable for q near $-1/2$ (including $q = -1$ and $q = 0$, which are the two cases of greatest practical interest). Approximation error plots for some values of q were presented in [41].

Substituting the approximations (3.61) and (3.60) into (3.47) yields the following approximation for the SNR of the CHO when combined with an unregularized Fisher reconstructor:

$$\text{SNR}^2_{\text{CHO,F}} \approx \boldsymbol{X}'\check{\boldsymbol{\Lambda}}^{(1+q)}\boldsymbol{T}(\boldsymbol{T}'\check{\boldsymbol{\Lambda}}^{(1+2q)}\boldsymbol{T} + \boldsymbol{\Pi}_{\text{int}})^{\dagger}\boldsymbol{T}'\check{\boldsymbol{\Lambda}}^{(1+q)}\boldsymbol{X}. \qquad (3.62)$$

Channels with Disjoint Passbands

The CHO filters are sometimes assumed to be bandpass filters with disjoint passbands. In that case, the vectors $\boldsymbol{t}^m$ have disjoint nonzero entries and the $M \times M$ matrix $\boldsymbol{T}'\check{\boldsymbol{\Lambda}}^{(1+2q)}\boldsymbol{T}$ on the right-hand side of (3.61) is diagonal, so we have

$$\begin{aligned}[\boldsymbol{T}'\check{\boldsymbol{\Lambda}}^{(1+q)}\boldsymbol{X}]_m &= \frac{1}{\sqrt{n_{\text{p}}}}\sum_{k\in\mathcal{T}_m}(t_k^m)^*\check{\lambda}_k^{(1+q)}X_k \\ [\boldsymbol{T}'\check{\boldsymbol{\Lambda}}^{(1+2q)}\boldsymbol{T}]_{mm} &= \frac{1}{n_{\text{p}}}\sum_{k\in\mathcal{T}_m}|t_k^m|^2\check{\lambda}_k^{(1+2q)},\end{aligned}$$

where $\mathcal{T}_m = \{k : t_k^m \neq 0\}$ is the passband of the mth filter and "$*$" denotes complex conjugate.

If we furthermore combine the assumption of disjoint passbands with the internal noise model (3.43), the covariance $\boldsymbol{\Pi}_{\text{int}}$ becomes diagonal (to within the approximations used to derive (3.61)) and the model (3.43) coincides with the one proposed in [1]. Under these assumptions the SNR approximation in (3.62) simplifies to

$$\text{SNR}^2_{\text{CHO,F}} \approx \text{SNR}^2_1 \triangleq \frac{1}{1+c_{\text{int}}} \sum_{m=1}^{M} \frac{\left| \sum_{k \in \mathcal{T}_m} X_k (t_k^m)^* \check{\lambda}_k^{(1+q)} \right|^2}{\sum_{k \in \mathcal{T}_m} |t_k^m|^2 \check{\lambda}_k^{(1+2q)}}, \tag{3.63}$$

We now examine the conditions under which the approximate SNR in (3.63) can reach the SNR upper bound in (3.44).

Achievability of the Optimal SNR

To determine how closely the approximate SNR in (3.63) can approach the SNR upper bound, define vectors $\boldsymbol{u}^m$ and $\boldsymbol{v}^m$ with elements

$$\begin{aligned} u_k^m &\triangleq X_k \check{\lambda}_k^{1/2} 1_{\{k \in \mathcal{T}_m\}}, \\ v_k^m &\triangleq t_k^m \check{\lambda}_k^{(q+1/2)}, \end{aligned}$$

respectively. Then rewriting (3.63) and using Cauchy's inequality yields

$$\begin{aligned} \text{SNR}^2_1 &= \frac{1}{1+c_{\text{int}}} \sum_{m=1}^{M} \frac{|\langle \boldsymbol{u}^m, \boldsymbol{v}^m \rangle|^2}{\|\boldsymbol{v}^m\|^2} \\ &\le \frac{1}{1+c_{\text{int}}} \sum_{m=1}^{M} \|\boldsymbol{u}^m\|^2 = \frac{1}{1+c_{\text{int}}} \sum_{k \in \mathcal{T}} |X_k|^2 \check{\lambda}_k \\ &\le \frac{1}{1+c_{\text{int}}} \sum_{k} |X_k|^2 \check{\lambda}_k \approx \frac{1}{1+c_{\text{int}}} \bar{\boldsymbol{f}}_s' \check{\boldsymbol{\mathcal{F}}} \bar{\boldsymbol{f}}_s = \frac{\text{SNR}^2_{\text{lin}_\circ}}{1+c_{\text{int}}}, \end{aligned} \tag{3.64}$$

where $\mathcal{T} = \cup_{m=1}^{M} \mathcal{T}_m$ denotes the combined passband of all M channels. If the combined passband $\mathcal{T}$ does not contain all of the signal energy, then the SNR will be strictly less than the optimal SNR. This suboptimality is expected due to the dimensionality decrease caused by the channels.

Is there an unregularized Fisher reconstructor $\boldsymbol{\mathcal{Z}}_q$ for which the upper bound in (3.64) can be achieved? Suppose that each channel filter has a flat passband, *i.e.*, $t_k^m = 1_{\{k \in \mathcal{T}_m\}}$.

Then there are two obvious cases where the SNR upper bound in (3.64) is achieved, as can be verified by substitution or by using the requirement $\boldsymbol{u}^m \propto \boldsymbol{v}^m \, \forall m$:

- If the X_k's are constant over each passband, then $q = 0$ (the unfiltered backprojector) will be optimal.
- If the $\check{\lambda}_k$'s are also constant over each passband, then any $q \in \mathbb{R}$ will be optimal.

In practice, it may be unlikely that either the $\check{\lambda}_k$'s or the X_k's are *exactly* uniform over each channel's passband, but if the passbands are reasonably narrow, then it is likely that these spectra will be *approximately* uniform over each passband.

We conclude that to within the accuracy of the approximations considered above, one or more unregularized reconstructors will nearly achieve the highest SNR obtainable for the given CHO channels and internal noise model. Once again, regularization does not seem to play a crucial role, even for the CHO.

3.4.5 CHO and QPWLS Reconstructors

The analysis in the previous section showed some situations in which one or more unregularized reconstructors allow the CHO to achieve (approximately) optimal SNR in the detection task at hand. We next examine the QPWLS family of regularized reconstructors, described in section 3.1.3, to explore how closely the CHO can approach the optimal SNR with a practical regularized reconstruction method and to examine the optimal choice of regularizer.

The general form of a QPWLS reconstructor is given in (3.13) and the commonly used weighting matrix $\boldsymbol{W}$ is the one corresponding to MAP estimation, *i.e.*,

$$\boldsymbol{W} = (\mathsf{Cov}\{\boldsymbol{y}|\boldsymbol{f} = \bar{\boldsymbol{f}}\})^{-1} = \boldsymbol{\Pi}^{-1},$$

where $\bar{\boldsymbol{f}}$ is the unconditional mean of the object. For this $\boldsymbol{W}$, the reconstructor in (3.13) becomes

$$\boldsymbol{\mathcal{Z}} = (\boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{R}})^{-1}\boldsymbol{\mathcal{A}}'\boldsymbol{\Pi}^{-1}. \tag{3.65}$$

Note that the QPWLS reconstructor in (3.65) is the same as the Fisher reconstructor (3.26) with $q = -1$ only if $\boldsymbol{\mathcal{K}}_f = \mathbf{0}$, *i.e.*, only if there is no object variability. We assume throughout that the regularization operator $\boldsymbol{\mathcal{R}}$ is chosen such that $\boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{R}}$ is positive definite. We would like to design $\boldsymbol{\mathcal{R}}$ to optimize CHO detection performance in the QPWLS-reconstructed images.

To analyze CHO performance with QPWLS reconstruction, we assume that both $\boldsymbol{\mathcal{F}}$ and $\boldsymbol{\mathcal{R}}$ are diagonalized locally by a common operator (the Fourier operator $\boldsymbol{\mathcal{U}}$). Specifically, we use (3.54), (3.55), (3.59) and also assume that

$$\boldsymbol{\mathcal{R}} \approx \boldsymbol{\mathcal{U}}^{-1}\boldsymbol{\Omega}\,\boldsymbol{\mathcal{U}}, \tag{3.66}$$

where $\boldsymbol{\Omega} \triangleq \mathsf{diag}\{\omega_k, k = 1, \ldots, n_{\mathrm{p}}\}$. The real, nonnegative ω_k's contain the local frequency response of the regularizer around the location of $\bar{\boldsymbol{f}}_s$. Approximations of simultaneous diagonalization of $\boldsymbol{\mathcal{F}}$ and $\boldsymbol{\mathcal{R}}$ have been used by other investigators to analyze observer performance with penalized-likelihood reconstruction [19, 104, 145].

Substituting the QPWLS reconstructor (3.65) and the emission tomography measurement moments (3.6) and (3.22) into the SNR of the CHO from (3.47) yields

$$\mathrm{SNR}^2_{\mathrm{CHO,QPWLS}} = \bar{\boldsymbol{f}}_s'\boldsymbol{\mathcal{F}}\boldsymbol{\mathcal{H}}^{-1}\boldsymbol{\mathcal{C}}[\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{H}}^{-1}(\boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{F}}\boldsymbol{\mathcal{K}}_f\boldsymbol{\mathcal{F}})\boldsymbol{\mathcal{H}}^{-1}\boldsymbol{\mathcal{C}} + \boldsymbol{\Pi}_{\mathrm{int}}]^{\dagger}\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{H}}^{-1}\boldsymbol{\mathcal{F}}\bar{\boldsymbol{f}}_s, \tag{3.67}$$

where $\boldsymbol{\mathcal{H}} \triangleq \boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{R}}$. Similarly to section 3.4.4, we use the local Fourier approximations (3.54), (3.55), (3.66) to obtain the following expressions for the moments of the channel outputs:

$$\mathsf{E}[\hat{\boldsymbol{c}}|H_1] - \mathsf{E}[\hat{\boldsymbol{c}}|H_0] = \boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{H}}^{-1}\boldsymbol{\mathcal{F}}\bar{\boldsymbol{f}}_s \tag{3.68}$$

$$\approx \boldsymbol{T}'\boldsymbol{H}^{-1}\boldsymbol{\Lambda}\boldsymbol{X} \tag{3.69}$$

$$\boldsymbol{\Pi}_{\hat{\boldsymbol{c}}} = \boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{H}}^{-1}(\boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{F}}\boldsymbol{\mathcal{K}}_f\boldsymbol{\mathcal{F}})\boldsymbol{\mathcal{H}}^{-1}\boldsymbol{\mathcal{C}} + \boldsymbol{\Pi}_{\text{int}} \tag{3.70}$$

$$\approx \boldsymbol{T}'\boldsymbol{H}^{-1}(\boldsymbol{\Lambda} + \boldsymbol{\Lambda}^2\boldsymbol{N})\boldsymbol{H}^{-1}\boldsymbol{T} + \boldsymbol{\Pi}_{\text{int}}, \tag{3.71}$$

where $\boldsymbol{H} \triangleq \boldsymbol{\Lambda} + \boldsymbol{\Omega} = \mathsf{diag}\{\lambda_k + \omega_k, k = 1, \dots, n_{\text{p}}\}$. Substituting (3.69) and (3.71) into (3.67) yields the following approximation for the SNR of the CHO when combined with a QPWLS reconstructor:

$$\text{SNR}^2_{\text{CHO,QPWLS}} \approx \boldsymbol{X}'\boldsymbol{\Lambda}\boldsymbol{H}^{-1}\boldsymbol{T}[\boldsymbol{T}'\boldsymbol{H}^{-1}(\boldsymbol{\Lambda} + \boldsymbol{\Lambda}^2\boldsymbol{N})\boldsymbol{H}^{-1}\boldsymbol{T} + \boldsymbol{\Pi}_{\text{int}}]^{\dagger}\boldsymbol{T}'\boldsymbol{H}^{-1}\boldsymbol{\Lambda}\boldsymbol{X}. \tag{3.72}$$

Channels with Disjoint Passbands

Similarly to section 3.4.4, when the CHO channels have disjoint frequency responses, the $M \times M$ matrix $\boldsymbol{T}'\boldsymbol{H}^{-1}(\boldsymbol{\Lambda} + \boldsymbol{\Lambda}^2\boldsymbol{N})\boldsymbol{H}^{-1}\boldsymbol{T}$ in (3.71) becomes diagonal. We assume that its diagonal elements are nonzero, *i.e.*, the system has some nonzero λ_k for each passband. (If not, the noninformative passband could be eliminated.)

As in section 3.4.4, we combine the assumption of disjoint passbands with the internal noise model (3.43). Then the approximate SNR expression in (3.72) becomes

$$\text{SNR}^2_{\text{CHO,QPWLS}} \approx \text{SNR}^2_2 \triangleq \frac{1}{1 + c_{\text{int}}} \sum_{m=1}^{M} \frac{\left| \sum_{k \in \mathcal{T}_m} X_k (t_k^m)^* \frac{\lambda_k}{\lambda_k + \omega_k} \right|^2}{\sum_{k \in \mathcal{T}_m} |t_k^m|^2 \frac{\lambda_k^2}{\check{\lambda}_k (\lambda_k + \omega_k)^2}}. \tag{3.73}$$

We now examine the conditions under which the approximate SNR in (3.73) can reach the SNR upper bound in (3.44).

Achievability of the Optimal SNR

To determine how closely the approximate SNR in (3.73) can approach the SNR upper bound, define vectors $\boldsymbol{u}^m$ and $\boldsymbol{v}^m$ with elements

$$\begin{aligned} u_k^m &\triangleq X_k \check{\lambda}_k^{1/2} 1_{\{k \in \mathcal{T}_m\}}, \\ v_k^m &\triangleq t_k^m \lambda_k / \check{\lambda}_k^{1/2} (\lambda_k + \omega_k), \end{aligned}$$

respectively. Then rewriting (3.73) and using Cauchy's inequality yields

$$\begin{aligned} \mathrm{SNR}_2^2 &= \frac{1}{1+c_{\text{int}}} \sum_{m=1}^{M} \frac{|\langle \boldsymbol{u}^m, \boldsymbol{v}^m \rangle|^2}{\|\boldsymbol{v}^m\|^2} \\ &\le \frac{1}{1+c_{\text{int}}} \sum_{m=1}^{M} \|\boldsymbol{u}^m\|^2 = \frac{1}{1+c_{\text{int}}} \sum_{k \in \mathcal{T}} |X_k|^2 \check{\lambda}_k \qquad (3.74) \\ &\le \frac{1}{1+c_{\text{int}}} \sum_{k} |X_k|^2 \check{\lambda}_k \approx \frac{1}{1+c_{\text{int}}} \bar{\boldsymbol{f}}_s' \check{\boldsymbol{\mathcal{F}}} \bar{\boldsymbol{f}}_s = \frac{\mathrm{SNR}_{\mathrm{lin}_\circ}^2}{1+c_{\text{int}}}. \qquad (3.75) \end{aligned}$$

Once again, if the combined passband $\mathcal{T} = \cup_{m=1}^{M} \mathcal{T}_m$ does not contain all of the signal energy, then the SNR will be strictly less than the optimal SNR.

The intermediate inequality in (3.74) becomes an equality if $\boldsymbol{u}^m \propto \boldsymbol{v}^m \, \forall m.$ Suppose that each channel filter is an ideal bandpass filter over some frequency band, *i.e.*, $t_k^m = 1_{\{k \in \mathcal{T}_m\}}$. Then, for $\lambda_k \neq 0$, $X_k \neq 0$, the intermediate upper bound in (3.74) is achieved for example when

$$\omega_k = \alpha \frac{\lambda_k}{X_k \check{\lambda}_k} - \lambda_k, \qquad (3.76)$$

where the constant $\alpha \neq 0$ can be chosen arbitrarily. Choosing

$$\alpha > \max_k X_k \check{\lambda}_k$$

would keep the ω_k's positive. To within approximations (3.54), (3.55) and (3.66), the local frequency response in (3.76) corresponds to the following nonnegative definite regularizer:

$$\boldsymbol{\mathcal{R}} = \alpha(\boldsymbol{\mathcal{I}} + \boldsymbol{\mathcal{F}}\boldsymbol{\mathcal{K}}_f)(\boldsymbol{\mathcal{U}}^{-1} \, \mathsf{diag}\{\boldsymbol{X}\}^{-1} \boldsymbol{\mathcal{U}}) - \boldsymbol{\mathcal{F}}. \qquad (3.77)$$

The $\mathcal{R}$ in (3.77) usually has a high-pass characteristic, so it could be construed as a regularization operator, but it is quite different from standard forms of regularization studied in the literature. Furthermore, substituting the $\mathcal{R}$ from (3.77) into the QPWLS reconstructor in (3.65) yields the estimator

$$\hat{\boldsymbol{f}} = \frac{1}{\alpha}(\boldsymbol{\mathcal{U}}^{-1}\,\mathsf{diag}\{\boldsymbol{X}\}\,\boldsymbol{\mathcal{U}})\boldsymbol{\mathcal{A}}'\check{\boldsymbol{\Pi}}^{-1}\boldsymbol{y}. \tag{3.78}$$

Therefore the “optimal” choice of $\mathcal{R}$ in (3.77) results in an unregularized estimator that is simply a weighted backprojection followed by the application of a “matched” filter (convolution with $\bar{\boldsymbol{f}}_s$). This agrees with the conclusion from section 3.4.4 that regularization is not essential for the CHO, if the passbands of the CHO channels are flat, non-overlapping and with internal noise covariance of the form (3.43).

Similarly to section 3.4.4, a degenerate case occurs when the channel passbands are flat and the mean signal spectrum $\{X_k\}$, system spectrum $\{\lambda_k\}$, and object power spectrum $\{\nu_k\}$ are all constant over each channel’s passband. Then the first upper bound in (3.74) is achieved for *any* choice of regularization $\{\omega_k\}$ that is also constant over each passband, including $\omega_k = 0$. Once again, under the internal noise model (3.43), the choice of regularization is not important in this detection task if there is little within-passband variation of the mean signal spectrum, the system spectrum, the object power spectrum, and/or the channel response itself.

3.4.6 CHO and QPWLS with a Roughness Penalty

In the previous section, where we placed no constraints on the form of the regularizer $\mathcal{R}$, we ended up with an “optimal” $\mathcal{R}$ that yielded the unregularized estimator in (3.78). We now constrain the regularizer to the commonly used form of a uniform quadratic roughness penalty. In this case, the regularization operator $\boldsymbol{\mathcal{R}}$ in (2.14) is such that

$$R(\boldsymbol{f}) = -\beta \sum_{j=1}^{n_{\mathrm{p}}} \sum_{k \in \mathcal{N}_j} (f_j - f_k)^2, \tag{3.79}$$

where f_j is the jth sample of the object $\boldsymbol{f}$, $\mathcal{N}_j$ a neighborhood of pixels around the jth pixel, and $\beta \geq 0$ the regularization parameter. A regularization term of the form (3.79) penalizes differences between neighboring image pixels, thus favoring smoother images. For simplicity, we consider here a first-order neighborhood $\mathcal{N}_j$, consisting of the four closest (top, bottom, left and right) neighboring pixels. The only free parameter in the regularizer (3.79) is thus the regularization parameter β.

We now provide an example of how β affects the SNR of the CHO with overlapping or non-overlapping passbands in the presence of two different types of internal noise. In this example, $\boldsymbol{\mathcal{A}}$ corresponds to a 2-D PET system model with the characteristics of a CTI ECAT 931 scanner (matrix size 128×128, pixel size 4.7mm, 192 radial samples with 3.1mm spacing, 160 projection angles over 180^{o}), generated by the ASPIRE software package [34].

We assume that the target signal $\boldsymbol{f}_s$ has a known Gaussian shape with a FWHM of 4 pixels and amplitude 0.1, the background $\boldsymbol{f}_b$ has a Gaussian autocorrelation function with a FWHM of 8 pixels and a standard deviation of 0.05, and the mean background $\bar{\boldsymbol{f}}_b$ is the anthropomorphic phantom shown in Fig. 3.2, which corresponds to a slice of the Zubal phantom [157]. We determine measurement variance by assuming that the signal intensity is weak with respect to the background intensity, so we can approximate

$$\boldsymbol{\Pi} = \mathsf{Cov}\{\boldsymbol{y}|\bar{\boldsymbol{f}}\} \approx \mathsf{Cov}\{\boldsymbol{y}|\bar{\boldsymbol{f}}_b\} = \mathsf{diag}\{\boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_b + \boldsymbol{r}\}\,.$$

We scale $\boldsymbol{\mathcal{A}}$ to produce a total of 5×10^5 recorded counts.

We consider three different models for the CHO channels: Non-overlapping square channels with $M = 4$ (SQR) and overlapping difference-of-Gaussians channels with

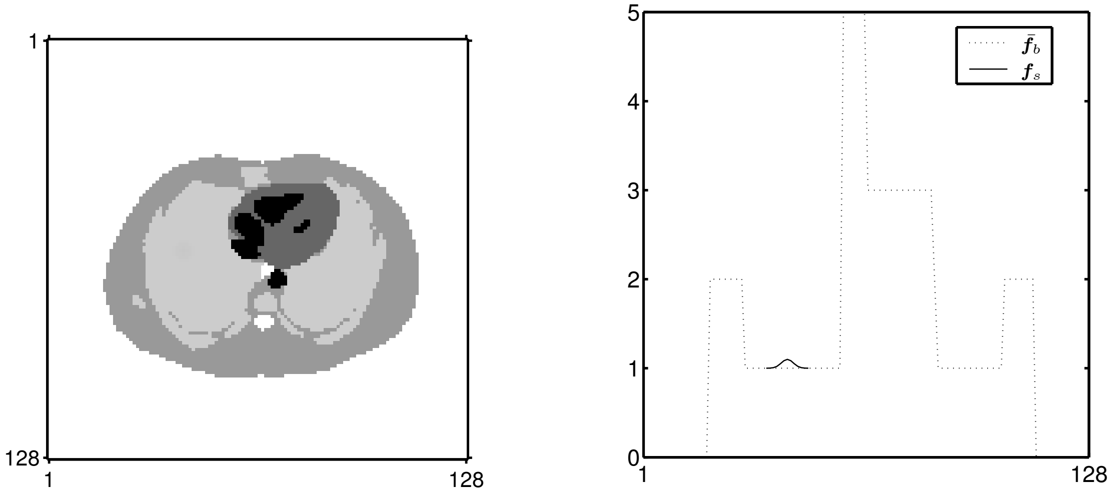

(a) Mean background $\bar{\boldsymbol{f}}_b$ and target signal $\boldsymbol{f}_s$

(b) Profile of $\bar{\boldsymbol{f}}_b$ and $\boldsymbol{f}_s$

Figure 3.2: Mean background and target signal for an example of a known-location task. The figure shows the mean background (left) and a profile through the mean background with the target signal superimposed (right).

$M = 3$ (S-DOG) and $M = 10$ (D-DOG), as defined by Abbey *et al.* [1] and plotted in Fig. 2.6. With each of these three channel sets, we consider two different models for the internal noise: *(i)* White internal noise, *i.e.*, $\boldsymbol{\Pi}_{\text{int}} = \sigma^2 \boldsymbol{I}$, where we choose the constant, reconstructor-independent variance of the channels to be $\sigma^2 = 0.005$, or *(ii)* The proportional internal noise covariance of (3.43). To facilitate comparison, we choose the proportionality constant c_{int} in (3.43) so that the peak SNR achieved with either of the two internal noise models is the same for a particular channel set.

Fig. 3.3 shows plots of the SNR of the CHO for QPWLS reconstruction with the roughness penalty in (3.79) and various values of β. The plots on the left are produced with the white internal noise model, whereas the ones on the right are produced with the proportional internal noise model. All the SNR values in the plots are normalized with respect to the ideal SNR in (3.23). The abscissa of the plots represents the resolution of the QPWLS reconstruction, defined as the FWHM (in pixels) of the reconstructed image when the true object is a noiseless impulse at the location of the target signal. This local FWHM is a measure of the amount of smoothing imposed by QPWLS. The resolution equals 1 pixel for $\beta = 0$, which corresponds to unregularized WLS, and it increases as β increases.

The plots show similar SNR behavior when comparing different channel sets under the same internal noise model. However, when the internal noise model changes, so does the behavior of the SNR versus resolution. Therefore, internal noise is an important factor to consider when optimizing regularization methods with respect to detectability. One should use the model that correlates best with humans for the task of interest. Since establishing correlation with humans is beyond the scope of this thesis, we present results with both models.

Fig. 3.3(a), 3.3(c), and 3.3(e) show that, with a constant, reconstructor-independent internal noise variance, a fast SNR drop occurs for very large amounts of regularization. This is due to the fact that, with a $\mathbf{\Pi}_{\text{int}} \neq \mathbf{0}$ that is independent of β, the SNR in (3.67) vanishes as $\beta \rightarrow \infty$. Similar behavior was reported by Qi with this type of internal noise [104]. However, no similarly fast drop in SNR occurs for very small amounts of regularization. In other words, with this type of internal noise and any of the three channel sets, the peak SNR achieved by QPWLS is only slightly higher than the SNR achieved by unregularized WLS (β = 0). We also investigated different amounts of imaging noise and/or background variability and obtained very similar results (not shown here). That is, unregularized WLS led to SNR very close to the peak SNR attainable with QPWLS and the roughness penalty in (3.79). Once again, we have a situation in which regularization is of minimal benefit in the known-location detection task – even with a CHO model that includes overlapping channel passbands and internal noise.

Fig. 3.3(b), 3.3(d), and 3.3(f) show that, with a proportional internal noise covariance, the SNR versus resolution plots are relatively flat. Under this model, the internal noise covariance decreases as β increases. As a result, the SNR does not drop fast for large

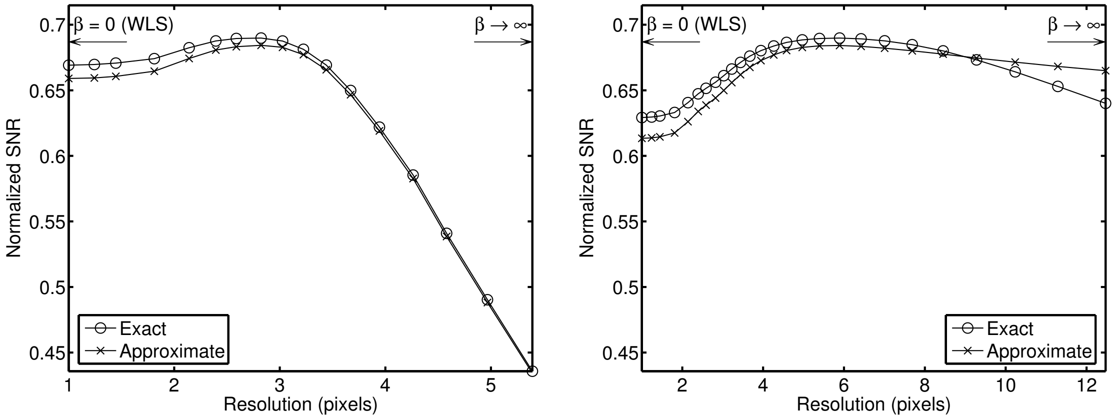

(a) SQR channels, const. internal variance

(b) SQR channels, prop. internal variance

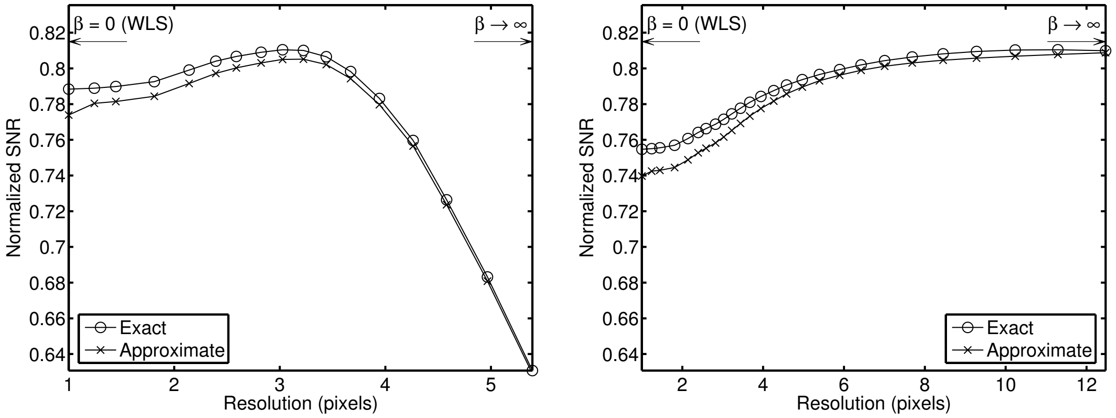

(c) S-DOG channels, const. internal variance

(d) S-DOG channels, prop. internal variance

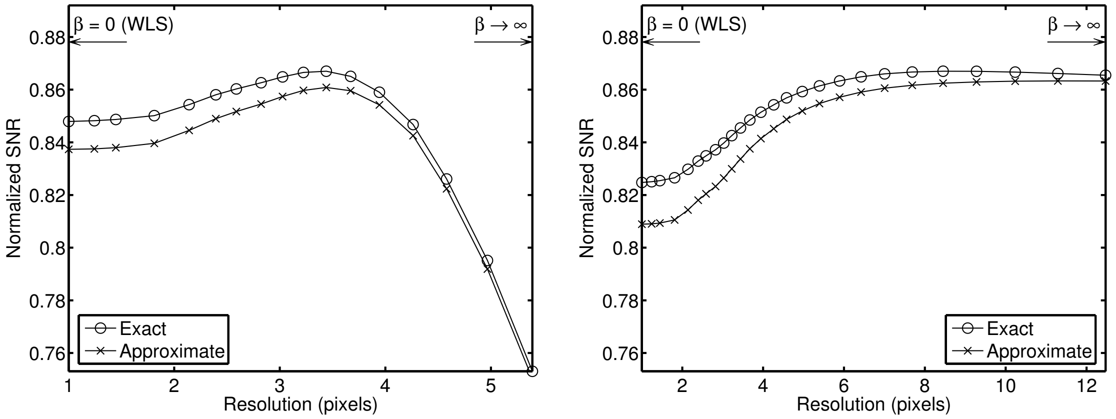

(e) D-DOG channels, const. internal variance

(f) D-DOG channels, prop. internal variance

Figure 3.3: SNR of the CHO versus QPWLS reconstruction resolution. Results are shown for three different channel sets and two different internal noise models. The exact and approximate SNR is computed from (3.67) and (3.72) respectively. All SNR values are normalized with respect to the upper bound (ideal SNR for the internal-noise-free observer) in (3.23).

amounts of regularization as it did with the constant internal covariance, where the internal noise dominated as β increased. With the proportional model, peak or near-peak SNR can be achieved with very large amounts of regularization. The SNR achieved with unregularized WLS is not much lower than the peak, although the difference is slightly more pronounced than with the constant internal covariance. Note that under the proportional internal covariance model, as seen from (3.44), changing the amount of internal noise would only scale the SNR plots, not altering their variation with resolution. This includes the case where there is no internal observer noise ($c_{\text{int}} = 0$). The optimality of a large amount of regularization, as seen in Fig. 3.3(b), 3.3(d), and 3.3(f), is consistent with the optimality of blurry backprojections, as shown for the CHO with the proportional internal covariance model in sections 3.4.4 and 3.4.5.

Note that there is a third possibility for the internal noise covariance, namely assuming that it is proportional to the *diagonal* of the left-hand side of (3.43), as Abbey and Barrett do in [1]. With this choice of internal noise covariance, the SNR drop that occurs for large amounts of regularization is faster than with the proportional covariance in Fig. 3.3(b), 3.3(d), and 3.3(f), but slower than with the white, reconstruction-independent covariance in Fig. 3.3(a), 3.3(c), and 3.3(e).

The plots in Fig. 3.3 also compare the exact SNR, computed from (3.67), to the approximate SNR, computed from (3.72). There is reasonable agreement between the two in all cases. This confirms the accuracy of the Fourier-domain approximations on which our results in sections 3.4.4 and 3.4.5 were based.

3.4.7 Channelized Non-Prewhitening Observer

Combining the moments of $\hat{\boldsymbol{f}}$ from (3.10) with the channelized feature vector from (2.32) and substituting into the CNPW template from (2.35) yields

$$\boldsymbol{w}_{\text{CNPW}} = \boldsymbol{\mathcal{C}}' \boldsymbol{\mathcal{Z}}\, \bar{\boldsymbol{y}}\,. \tag{3.80}$$

By substituting the CNPW template from (3.80) into the SNR expression (3.42), we find this observer's SNR:

$$\text{SNR}^2_{\text{CNPW}} = \frac{(\bar{\boldsymbol{y}}'\, \boldsymbol{\mathcal{Z}}' \boldsymbol{\mathcal{C}} \boldsymbol{\mathcal{C}}' \boldsymbol{\mathcal{Z}}\, \bar{\boldsymbol{y}})^2}{\bar{\boldsymbol{y}}'\, \boldsymbol{\mathcal{Z}}' \boldsymbol{\mathcal{C}} (\boldsymbol{\mathcal{C}}' \boldsymbol{\mathcal{Z}} \check{\boldsymbol{\Pi}} \boldsymbol{\mathcal{Z}}' \boldsymbol{\mathcal{C}} + \boldsymbol{\Pi}_{\text{int}}) \boldsymbol{\mathcal{C}}' \boldsymbol{\mathcal{Z}}\, \bar{\boldsymbol{y}}}. \tag{3.81}$$

It is straightforward to show that the SNR of the CNPW observer is bounded above by the SNR of the CHO. Using

$$\bar{\boldsymbol{c}} = \boldsymbol{\mathcal{C}}' \boldsymbol{\mathcal{Z}}\, \bar{\boldsymbol{y}}, \qquad \boldsymbol{\Pi}_{\hat{\boldsymbol{c}}} = \boldsymbol{\mathcal{C}}' \boldsymbol{\mathcal{Z}} \check{\boldsymbol{\Pi}} \boldsymbol{\mathcal{Z}}' \boldsymbol{\mathcal{C}} + \boldsymbol{\Pi}_{\text{int}} \tag{3.82}$$

and the SNR expressions from (3.47) and (3.81), we have

$$\begin{aligned} \text{SNR}^2_{\text{CNPW}} &= \frac{|\langle \bar{\boldsymbol{c}},\, \bar{\boldsymbol{c}} \rangle|^2}{\left\| \boldsymbol{\Pi}_{\hat{\boldsymbol{c}}}^{(-1/2)} \bar{\boldsymbol{c}} \right\|^2} = \frac{\left| \langle \boldsymbol{\Pi}_{\hat{\boldsymbol{c}}}^{(-1/2)} \bar{\boldsymbol{c}},\, \boldsymbol{\Pi}_{\hat{\boldsymbol{c}}}^{1/2} \bar{\boldsymbol{c}} \rangle \right|^2}{\left\| \boldsymbol{\Pi}_{\hat{\boldsymbol{c}}}^{1/2} \bar{\boldsymbol{c}} \right\|^2} && (3.83) \\ &\leq \left\| \boldsymbol{\Pi}_{\hat{\boldsymbol{c}}}^{(-1/2)} \bar{\boldsymbol{c}} \right\|^2 = \text{SNR}^2_{\text{CHO}}, && (3.84) \end{aligned}$$

where we use (3.30) with $\boldsymbol{\mathcal{B}} = \check{\boldsymbol{\Pi}}^{1/2} \boldsymbol{\mathcal{Z}}' \boldsymbol{\mathcal{C}}$ to get the right-hand side of (3.83) and we use the Cauchy-Schwarz inequality to get (3.84). The upper bound in (3.84) is achieved when

$$\boldsymbol{\Pi}_{\hat{\boldsymbol{c}}}^{\dagger} \bar{\boldsymbol{c}} \propto \bar{\boldsymbol{c}}, \tag{3.85}$$

in which case the CHO and CNPW templates become equivalent. This holds when the channel outputs are white.

We now examine the existence of a reconstructor that allows the CNPW observer to achieve the optimal SNR in (3.44). In the special case of the internal noise model (3.43), using (3.80), we have

$$\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}\boldsymbol{w}_{\mathrm{CNPW}} \propto \boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}}\,. \tag{3.86}$$

By comparing (3.86) to the optimality condition (3.45), we find that the SNR upper bound in (3.44) is achieved when $\boldsymbol{\mathcal{Z}}$ satisfies

$$\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}} \propto \check{\boldsymbol{\Pi}}^{-1}\,\bar{\boldsymbol{y}}, \tag{3.87}$$

which is similar to (3.36). To find a $\boldsymbol{\mathcal{Z}}$ that satisfies (3.87), consider once again a reconstructor of the form (3.50). For the $\boldsymbol{\mathcal{Z}}$ in (3.50), with any image-domain weighting operator $\boldsymbol{\mathcal{W}}$, the optimality condition (3.87) is satisfied if

$$\boldsymbol{\mathcal{A}}\boldsymbol{\mathcal{G}}\boldsymbol{\mathcal{G}}'\boldsymbol{\mathcal{A}}'\check{\boldsymbol{\Pi}}^{-1}\,\bar{\boldsymbol{y}} \propto \bar{\boldsymbol{y}}\,. \tag{3.88}$$

As long as $\boldsymbol{\mathcal{G}}$ can be chosen to satisfy (3.88), the unregularized reconstructor $\boldsymbol{\mathcal{Z}}$ given in (3.50) allows the CNPW observer to achieve the optimal SNR. For emission tomography, where $\bar{\boldsymbol{y}} = \boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_s$, the condition (3.88) is satisfied if

$$\boldsymbol{\mathcal{G}}\boldsymbol{\mathcal{G}}'\check{\boldsymbol{\mathcal{F}}}\bar{\boldsymbol{f}}_s \propto \bar{\boldsymbol{f}}_s. \tag{3.89}$$

Choosing $\boldsymbol{\mathcal{G}}$ so that one of its "columns" is proportional to $\bar{\boldsymbol{f}}_s$ and the remaining columns, $\boldsymbol{\mathcal{G}}_m$, $m = 2, \dots, M$, satisfy $\boldsymbol{\mathcal{G}}_m'\check{\boldsymbol{\mathcal{F}}}\bar{\boldsymbol{f}}_s = 0$, $m = 2, \dots, M$, is one way to satisfy (3.89). Since $\mathrm{SNR}^2_{\mathrm{CNPW}} \le \mathrm{SNR}^2_{\mathrm{CHO}} \le \frac{1}{1+c_{\mathrm{int}}}\mathrm{SNR}^2_{\mathrm{lin}_\circ}$, it is reasonable that the requirements on $\boldsymbol{\mathcal{G}}$ derived in section 3.4.2 for the CHO are a subset of those derived here for the CNPW.

Therefore, under the internal noise model (3.43), the family of reconstructors in (3.50) with an appropriately chosen $\boldsymbol{\mathcal{G}}$ can ensure optimal SNR for the CNPW observer. Note that, with a reconstructor of the form (3.50) and $\boldsymbol{\mathcal{G}}$ chosen as discussed in the previous paragraph, the moments of the channel outputs become:

$$\begin{aligned} \boldsymbol{w}_{\mathrm{CNPW}} &= \mathsf{E}[\hat{\boldsymbol{c}}|H_1] - \mathsf{E}[\hat{\boldsymbol{c}}|H_0] = \boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\,\bar{\boldsymbol{y}} \\ &= \boldsymbol{\mathcal{G}}'\check{\boldsymbol{\mathcal{F}}}\bar{\boldsymbol{f}}_s = [\bar{\boldsymbol{f}}_s'\check{\boldsymbol{\mathcal{F}}}\bar{\boldsymbol{f}}_s\ 0\ 0\ \ldots\ 0]' \qquad (3.90) \\ \boldsymbol{\Pi}_{\hat{\boldsymbol{c}}} &= (1+c_{\mathrm{int}})\boldsymbol{\mathcal{C}}'\boldsymbol{\mathcal{Z}}\check{\boldsymbol{\Pi}}\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}} = (1+c_{\mathrm{int}})\boldsymbol{\mathcal{G}}'\check{\boldsymbol{\mathcal{F}}}\boldsymbol{\mathcal{G}}. \qquad (3.91) \end{aligned}$$

As seen on the right-hand side of (3.91), the reconstructor at hand does not necessarily result in white channel outputs. However, the reconstructor transforms the data so that the output of a single channel (the first channel in the above notation) becomes uncorrelated from the outputs of the remaining channels. Then detection can be performed with a template that, as seen on the right-hand side of (3.90), uses only this uncorrelated channel, disregarding the rest. Then the second-order statistics of the outputs of the remaining channels become irrelevant. Thus the CNPW observer is not handicapped by its lack of knowledge about these second-order statistics and it achieves the same performance as the CHO.

As discussed in section 3.4.2, the reconstructor family (3.50) does not produce what we usually consider to be reconstructed images. Thus it is of interest to know whether more conventional reconstructors can be combined with the CNPW observer to achieve SNR values close to the optimal and, if so, whether these reconstructors must include regularization. As before, we do this with the use of approximations to the SNR expression in (3.81) obtained through local Fourier-domain analysis. In the following, we apply such approximations to analyze the SNR of the CNPW observer with regularized reconstruction.

3.4.8 CNPW Observer and QPWLS Reconstructors

As with the CHO in section 3.4.5, we now explore how closely the CNPW observer can approach the optimal SNR with a practical regularized reconstruction method and examine the optimal choice of regularizer. We consider the QPWLS reconstructor given in (3.65), assuming as usual that the regularization operator $\boldsymbol{\mathcal{R}}$ is chosen such that $\boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{R}}$ is positive definite.

This time we would like to design $\boldsymbol{\mathcal{R}}$ to optimize the CNPW observer's detection performance in the QPWLS-reconstructed images. Substituting the QPWLS reconstructor (3.65) and the emission tomography measurement moments (3.6) and (3.22) into the SNR of the CNPW observer from (3.81) yields

(3.92)

$$\mathrm{SNR}^2_{\mathrm{CNPW,QPWLS}} = \frac{(\bar{\boldsymbol{f}}_s' \boldsymbol{\mathcal{F}} \boldsymbol{\mathcal{H}}^{-1} \boldsymbol{\mathcal{C}} \boldsymbol{\mathcal{C}}' \boldsymbol{\mathcal{H}}^{-1} \boldsymbol{\mathcal{F}} \bar{\boldsymbol{f}}_s)^2}{\bar{\boldsymbol{f}}_s' \boldsymbol{\mathcal{F}} \boldsymbol{\mathcal{H}}^{-1} \boldsymbol{\mathcal{C}} [\boldsymbol{\mathcal{C}}' \boldsymbol{\mathcal{H}}^{-1} (\boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{F}} \boldsymbol{\mathcal{K}}_f \boldsymbol{\mathcal{F}}) \boldsymbol{\mathcal{H}}^{-1} \boldsymbol{\mathcal{C}} + \boldsymbol{\Pi}_{\mathrm{int}}] \boldsymbol{\mathcal{C}}' \boldsymbol{\mathcal{H}}^{-1} \boldsymbol{\mathcal{F}} \bar{\boldsymbol{f}}_s},$$

where $\boldsymbol{\mathcal{H}} \triangleq \boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{R}}$.

Similarly to section 3.4.5, we assume that the operators $\boldsymbol{\mathcal{F}}$, $\boldsymbol{\mathcal{K}}_f$, and $\boldsymbol{\mathcal{R}}$ are approximately diagonalized by the Fourier transform around the location of the target signal according to (3.54), (3.55), and (3.66) respectively. This allows us to use the approximations (3.69) and (3.71) for the moments of the channel outputs. Substituting (3.69) and (3.71) into (3.92) yields the following approximation for the SNR of the CNPW when combined with a QPWLS reconstructor:

(3.93)

$$\mathrm{SNR}^2_{\mathrm{CNPW,QPWLS}} \approx \frac{(\boldsymbol{X}' \boldsymbol{\Lambda} \boldsymbol{H}^{-1} \boldsymbol{T} \boldsymbol{T}' \boldsymbol{H}^{-1} \boldsymbol{\Lambda} \boldsymbol{X})^2}{\boldsymbol{X}' \boldsymbol{\Lambda} \boldsymbol{H}^{-1} \boldsymbol{T} [\boldsymbol{T}' \boldsymbol{H}^{-1} (\boldsymbol{\Lambda} + \boldsymbol{\Lambda}^2 \boldsymbol{N}) \boldsymbol{H}^{-1} \boldsymbol{T} + \boldsymbol{\Pi}_{\mathrm{int}}] \boldsymbol{T}' \boldsymbol{H}^{-1} \boldsymbol{\Lambda} \boldsymbol{X}},$$

where $\boldsymbol{H} \triangleq \boldsymbol{\Lambda} + \boldsymbol{\Omega} = \mathsf{diag}\{\lambda_k + \omega_k, k = 1, \ldots, n_{\mathrm{p}}\}$.

Channels with Disjoint Passbands

Based on (3.69) and (3.71), the moments of individual channel outputs can be approximated as

$$[\bar{\boldsymbol{c}}]_m \approx c_m \triangleq \frac{1}{\sqrt{n_{\mathrm{p}}}} \sum_{k \in \mathcal{T}_m} X_k (t_k^m)^* \frac{\lambda_k}{\lambda_k + \omega_k}, \quad m = 1, \dots, M \tag{3.94}$$

$$[\boldsymbol{\Pi}_{\hat{\mathbf{c}}}]_{mm} \approx s_m^2 \triangleq \frac{1}{n_{\mathrm{p}}} \sum_{k \in \mathcal{T}_m} |t_k^m|^2 \frac{\lambda_k^2}{\check{\lambda}_k (\lambda_k + \omega_k)^2} + \sigma_m^2, \quad m = 1, \dots, M, \tag{3.95}$$

where $\sigma_m^2 \triangleq [\boldsymbol{\Pi}_{\mathrm{int}}]_{mm}$ is the internal noise variance of the mth channel.

Similarly to section 3.4.5, we consider the case where the channels of the CNPW observer have disjoint frequency responses, and the internal noise covariance has the proportional form of (3.43). In this special case, the approximate covariance matrix in (3.71) becomes diagonal. Then the approximate SNR expression in (3.93) can be written as

$$\mathrm{SNR}^2_{\mathrm{CNPW,QPWLS}} \approx \mathrm{SNR}_3^2 \triangleq \frac{1}{1 + c_{\mathrm{int}}} \frac{\left(\sum_{m=1}^M |c_m|^2 \right)^2}{\sum_{m=1}^M |c_m|^2 s_m^2} \tag{3.96}$$

with c_m and s_m^2 as defined in (3.94) and (3.95) respectively. We now examine the conditions under which the approximate SNR in (3.96) can reach the SNR upper bound in (3.44).

Achievability of the Optimal SNR

To determine how closely the approximate SNR of the CNPW observer in (3.96) can approach the SNR upper bound, we first examine the conditions under which (3.96) can approach the approximate SNR of the CHO in (3.73).

As usual, we assume that the system has some nonzero λ_k's for each channel passband. (If not, the noninformative passband could be eliminated.) If the mean signal $\bar{\boldsymbol{f}}_s$ is well-localized in space, we can assume similarly that there is no passband within which the X_k's are all zero. This implies that the elements of $\bar{\boldsymbol{c}}$, as approximated in (3.94), are generally non-zero for QPWLS reconstruction.

When all elements of $\bar{\boldsymbol{c}}$ are non-zero and $\boldsymbol{\Pi}_{\hat{c}}$ is diagonal, the only way to satisfy (3.85) is with $\boldsymbol{\Pi}_{\hat{c}} \propto \boldsymbol{I}$. Therefore, assuming disjoint channel passbands and to within approximations (3.94) and (3.95), the CNPW observer achieves the same SNR as the CHO only if the channel outputs are white:

$$s_m^2 = \text{const.}, \quad m = 1, \ldots, M. \tag{3.97}$$

Consider the case where the internal noise variance is either proportional to the variance due to the channel inputs, as in (3.43), or input-independent but constant, *i.e.*, $\boldsymbol{\Pi}_{\text{int}} = \sigma^2 \boldsymbol{I}$. Then, due to (3.95), the requirement (3.97) is equivalent to

$$\sum_{k \in \mathcal{T}_m} \frac{|t_k^m|^2 \lambda_k^2}{\check{\lambda}_k (\lambda_k + \omega_k)^2} = \text{const.}, \quad m = 1, \ldots, M. \tag{3.98}$$

A choice of ω_k that satisfies the requirement (3.98) is

$$\omega_k = \alpha \xi_k \frac{\lambda_k}{\check{\lambda}_k^{1/2}} - \lambda_k, \tag{3.99}$$

where

$$\xi_k \triangleq \begin{cases} \|\boldsymbol{t}^m\|, & k \in \mathcal{T}_m \\ \xi_k^{\text{o}}, & k \notin \mathcal{T} \triangleq \cup_{m=1}^{M} \mathcal{T}_m \end{cases}$$

and $\alpha \neq 0$, $\xi_k^{\text{o}} \neq 0$ can be chosen arbitrarily. Choosing

$$\alpha > \max_k \frac{\check{\lambda}_k^{1/2}}{\xi_k}, \qquad \xi_k^{\text{o}} > 0$$

would keep the ω_k's positive.

To within approximations (3.54), (3.55) and (3.66), the local frequency response in (3.99) corresponds to the following nonnegative definite regularizer:

$$\boldsymbol{\mathcal{R}} = \alpha \boldsymbol{\mathcal{F}} \check{\boldsymbol{\mathcal{F}}}^{(-1/2)} (\boldsymbol{\mathcal{U}}^{-1} \Xi \boldsymbol{\mathcal{U}}) - \boldsymbol{\mathcal{F}}, \tag{3.100}$$

where $\Xi \triangleq \mathsf{diag}\{\xi_k, k = 1, \ldots, n_{\text{p}}\}$. The $\boldsymbol{\mathcal{R}}$ in (3.100) may have a high-pass characteristic, but, other than that, it is quite different from standard forms of regularization in the literature. Substituting the $\boldsymbol{\mathcal{R}}$ from (3.100) into the QPWLS reconstructor in (3.65) yields the estimator

$$\hat{\boldsymbol{f}} = \frac{1}{\alpha}(\boldsymbol{\mathcal{U}}^{-1}\Xi^{-1}\boldsymbol{\mathcal{U}})\check{\boldsymbol{\mathcal{F}}}^{1/2}(\boldsymbol{\mathcal{F}}^{\dagger}\boldsymbol{\mathcal{A}}'\boldsymbol{\Pi}^{-1}\boldsymbol{y}). \tag{3.101}$$

Therefore the "optimal" choice of $\boldsymbol{\mathcal{R}}$ in (3.100) results in an unregularized WLS estimator followed by the postfilter $\check{\boldsymbol{\mathcal{F}}}^{1/2}$, similarly to the non-channelized NPW case in (3.38), and then convolution with a shift-invariant filter whose frequency response Ξ^{-1} depends on the channel filters. In typical channel models, such as the ones shown in Fig. 2.6, higher-frequency channels $\boldsymbol{t}^m$ have wider passbands, therefore $\|\boldsymbol{t}^m\|$ increases with frequency. This means that within the combined channel passband $\mathcal{T}$ (or everywhere, if the ξ_k^{o}'s are chosen accordingly) the frequency response Ξ corresponds to a high-pass filter, therefore Ξ^{-1} in (3.101) corresponds to a low-pass filter.

Because it satisfies (3.97), the regularizer choice in (3.100) ensures that the channel outputs are white and therefore (for disjoint channel passbands) that the SNR of the CNPW observer becomes equal to the SNR of the CHO. Under the internal noise model (3.43), in particular, (3.92) becomes equal to (3.67). Furthermore, as discussed in the last paragraph of section 3.4.5, if the channel passbands are flat and the mean signal spectrum $\{X_k\}$, system spectrum $\{\lambda_k\}$, and object power spectrum $\{\nu_k\}$ are all constant over each channel's passband, the SNR in (3.67) achieves the first upper bound in (3.74) for any choice of regularization $\{\omega_k\}$ that is also constant over each passband. This includes the $\{\omega_k\}$ in (3.99), when there is little within-passband variation of the mean signal spectrum, the system spectrum, and the object power spectrum. In this case, the SNR of the CNPW observer also reaches the overall optimal SNR attainable with the given channel passbands and internal noise. As usual, if the combined channel passband $\mathcal{T}$ does not contain all of the signal energy, then this SNR will be strictly less than the optimal SNR in (3.75).

We conclude there are some situations where regularization in the usual sense seems to not be essential for the CNPW. In these cases, the CNPW observer can (approximately) achieve optimal SNR in the detection task at hand with an unregularized WLS reconstructor. However, because the optimal reconstructor must now compensate for the prewhitening that the CNPW observer cannot perform, these cases are more limited than the ones in which the CHO achieves the same SNR without regularization.

3.4.9 CNPW Observer and QPWLS with a Roughness Penalty

As in section (3.4.6), we now constrain the regularizer to the practical form of the uniform quadratic roughness penalty in (3.79), which penalizes differences between neighboring image pixels. Once again, we consider a first-order pixel neighborhood $\mathcal{N}_j$. We are interested in tuning the only free parameter in the regularizer (3.79), the regularization parameter β, so that the SNR of the CNPW observer is maximized.

The first-order quadratic penalty in (3.79) corresponds to

$$\omega_k = 2\beta\Big[2 - \cos(2\pi k_x/n_{\mathrm{p}}) - \cos(2\pi k_y/n_{\mathrm{p}})\Big], \tag{3.102}$$

where (k_x, k_y) are the discrete 2-D spatial frequency coordinates of the kth sample of the frequency response. As is the case with typical roughness penalties, (3.102) is of a high-pass nature. Because of the white channel output condition that needs to be met in order to achieve optimality for the CNPW observer, one should expect the value of β to have more of an impact on the SNR of this observer than it has on the SNR of the CHO.

For a simple illustration of this property, consider the case where the detection task is SKE/BKE ($\check{\lambda}_k = \lambda_k$) and the channel passbands are disjoint and flat ($t_k^m = 1_{\{k \in \mathcal{T}_m\}}$). Then the whitening condition in (3.98) becomes

$$\sum_{k \in \mathcal{T}_m} \frac{\lambda_k}{(\lambda_k + \omega_k)^2} = \text{const.}, \quad m = 1, \ldots, M. \tag{3.103}$$

In the unregularized case where $\omega_k = 0\,\forall k$, this becomes

$$\sum_{k \in \mathcal{T}_m} \frac{1}{\lambda_k} = \text{const.}, \quad m = 1, \ldots, M. \tag{3.104}$$

However, not only is the system frequency response $\{\lambda_k\}$ low-pass, but also the passbands $\mathcal{T}_m$ of higher-frequency channels are typically wider. Therefore, both the values and the number of the positive summands on the left-hand side of (3.104) increase with increasing m. As a result, the whitening condition cannot be satisfied with $\omega_k = 0\,\forall k$. The presence of appropriately chosen, high-pass, non-zero ω_k's decreases the values of the higher-frequency summands on the left-hand side of (3.103) so that the whitening condition can be satisfied.

The analysis of CNPW performance in the previous section showed that, if we place no constraints on the form of the regularizer $\boldsymbol{\mathcal{R}}$, we end up with the "optimal" $\boldsymbol{\mathcal{R}}$ in (3.100), which may be high-pass and non-zero, but yields the unregularized estimator in (3.101). However, if we constrain the regularizer to the form of a roughness penalty such as the one in (3.79), we expect, due to the previous discussion, that regularized reconstruction with $\beta > 0$ will lead to better CNPW performance than the unregularized $\beta = 0$.

We provide here an example of how β affects the SNR of the CNPW with overlapping or non-overlapping passbands in the presence of two different types of internal noise. This example follows the same assumptions for the system model $\boldsymbol{\mathcal{A}}$, the target signal $\boldsymbol{f}_s$, the background $\boldsymbol{f}_b$, and the imaging noise level (number of counts) as the CHO example of section 3.4.6.

We also consider the same three models for the CNPW channels, *i.e.*, the non-overlapping SQR and overlapping S-DOG and D-DOG channels from [1]. With each of these channel sets, we apply the two models for the internal noise covariance that we discussed in section 3.4.6, *i.e.*, white or proportional.

In Fig. 3.4 we show plots of the SNR of the CNPW observer for QPWLS reconstruction with the roughness penalty in (3.79) and various values of β. The plots on the left are produced with the white internal noise model, whereas the ones on the right are produced with the proportional internal noise model. All the SNR values in these plots are normalized with respect to the ideal SNR in (3.23). The abscissa of the plots represents the resolution of the QPWLS reconstruction at the location of the target signal, as defined in section 3.4.6.

As was the case with the CHO, Fig. 3.4 shows that the three different channel sets that we consider lead to similar SNR behavior for the CNPW observer. Compared to the CHO, however, the CNPW observer has a more consistent behavior across the two different internal noise models, in that it exhibits a clear SNR peak with either noise model. The existence of this peak is related to the fact that the inability to prewhiten leads the CNPW observer to suboptimal performance and the reconstructor needs to compensate for this inability.

For no regularization ($\beta = 0$), the channel outputs are not white, since the outputs of higher-frequency channels have higher variance, so the CPNW observer performs suboptimally. As the amount of regularization increases, the variances of the channel outputs decrease, with those of higher-frequency channels decreasing faster. For some intermediate amount of regularization the variances will become closer to uniform, mitigating the effect of the observer's inability to prewhiten and leading to improved SNR. However, for very large amounts of regularization, as the variances of higher-frequency channels decrease excessively, the channel variances become non-uniform again. Furthermore, as the amount of regularization increases, so do the correlations between channel outputs. Thus the channel outputs veer away from whiteness for too much regularization, leading to the SNR drop seen in all plots of Fig. 3.4 as $\beta \to \infty$.

Whether the internal noise covariance is constantly white or proportional affects this

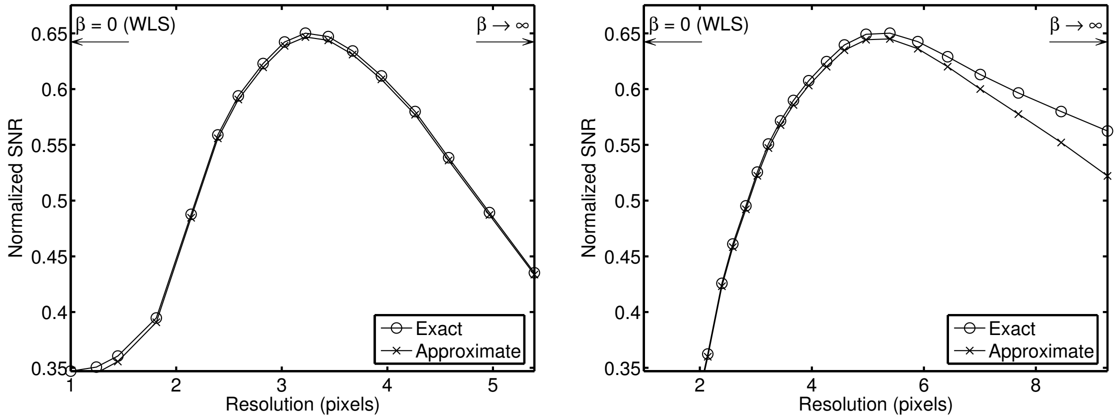

(a) SQR channels, const. internal variance

(b) SQR channels, prop. internal variance

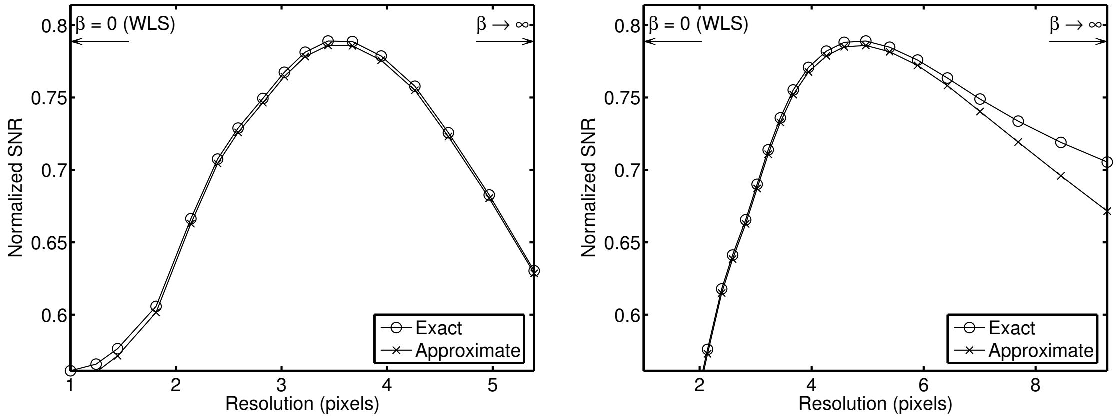

(c) S-DOG channels, const. internal variance

(d) S-DOG channels, prop. internal variance

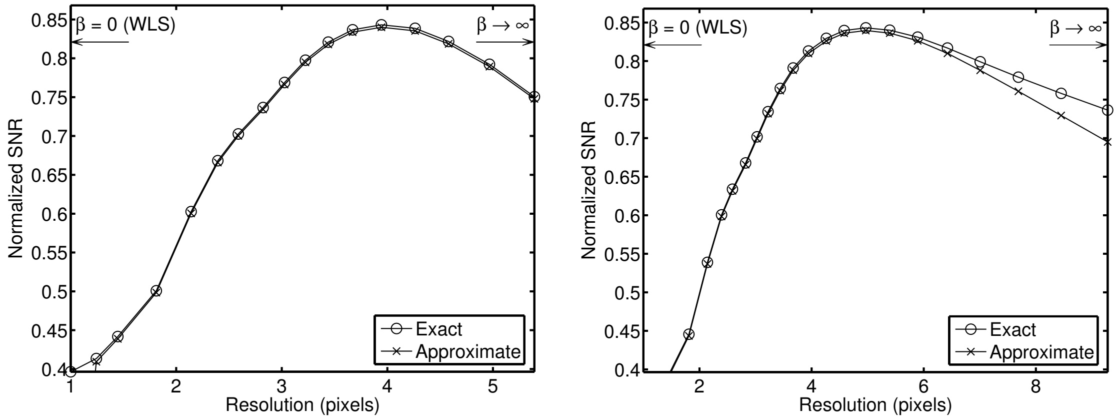

(e) D-DOG channels, const. internal variance

(f) D-DOG channels, prop. internal variance

Figure 3.4: SNR of the CNPW observer versus QPWLS reconstruction resolution. Results are shown for three different channel sets and two different internal noise models. The exact and approximate SNR is computed from (3.92) and (3.93) respectively. All SNR values are normalized with respect to the upper bound (ideal SNR for the internal-noise-free observer) in (3.23).

process as well. The SNR drops faster under the white internal covariance model, as the internal noise dominates for $\beta \to \infty$ in this case. As a result, the SNR peak occurs at a lower β in Fig. 3.4(a), 3.4(c), and 3.4(e) than it does under the proportional internal covariance model in Fig. 3.4(b), 3.4(d), and 3.4(f) respectively.

The plots in Fig. 3.4 also compare the exact SNR, computed from (3.92), to the approximate SNR, computed from (3.93). The good agreement between the two confirms the accuracy of the Fourier-domain approximations on which our results in section 3.4.8 were based.

In general, we find here that regularization can be beneficial in a known-location detection task for the CNPW observer, much more so than for the CHO. This applies for non-overlapping or overlapping channel passbands, with or without internal noise (as the proportional internal covariance model encompasses the case $c_{\text{int}} = 0$).

Thus an important factor to consider when optimizing regularization with respect to observer performance is whether the observer of interest can do prewhitening or not. As discussed in section 2.3.9, human observers have been found to perform more like the CHO than the CNPW observer in known-location tasks. This implies that human observers have some knowledge of the second-order statistics of the reconstructed images and are able to use this knowledge when performing known-location tasks. However, one can imagine that this knowledge may be imperfect. Furthermore, since the MaCNPW observer (see sections 2.3.8, 2.3.9) has been found to be well-correlated with humans in unknown-location tasks, it is plausible that humans are less capable of incorporating their knowledge of second-order statistics in some tasks than in others. Finally, when applying the CHO model itself to images produced from real data, the second-order statistics may be difficult to estimate accurately. For all these reasons, one may be interested in the effect of regularized reconstruction on observers that have imperfect prewhitening capabilities.

We consider an example of such an observer model in the following section.

3.4.10 Channelized Observers with Imperfect Prewhitening

To illustrate the effect of an imperfect knowledge of second-order statistics on the observer side, we define here a *channelized partially prewhitening* (CPPW) observer with template

$$\boldsymbol{w}_{\mathrm{CPPW}} \triangleq \Big[(1-\gamma)\boldsymbol{I} + \gamma\boldsymbol{\Pi}_{\hat{\boldsymbol{c}}}\Big]^{\dagger}\Big(\mathsf{E}[\hat{\boldsymbol{c}}|H_1] - \mathsf{E}[\hat{\boldsymbol{c}}|H_0]\Big), \tag{3.105}$$

for some $\gamma \in [0,1]$. For $\gamma = 0$ the above becomes equal to the CNPW template in (2.35) and for $\gamma = 1$ it becomes equal to the CHO template in (2.34). Intermediate values of γ yield intermediate prewhitening accuracy.

The expression in (3.105) is only one of several ways in which a prewhitening deficiency could possibly be introduced in the channelized observer's template. We use (3.105) here simply to illustrate our point, without making any claims about correlation with human observers.

The SNR of a CPPW observer applied to QPWLS-reconstructed images can be computed by substituting the reconstructor (3.65) and the channel moments (3.68), (3.70) into the template (3.105) and the channelized SNR expression (3.42). We revisit here the example from sections 3.4.6 and 3.4.9, this time examining the SNR of CPPW observers with various values of γ in the same task. In Fig. 3.5 we plot the SNR of some of these observers for QPWLS reconstruction with the roughness penalty in (3.79) and various values of β. The SNR is plotted for CPPW observers with S-DOG channels and the two internal noise models considered in sections 3.4.6 and 3.4.9.

With both internal noise models, as seen in Fig. 3.5(a) and Fig. 3.5(b), the effect of imperfect prewhitening is similar. For unregularized WLS ($\beta = 0$), the SNR varies widely with the level of prewhitening performed by the observer. For some intermediate range of

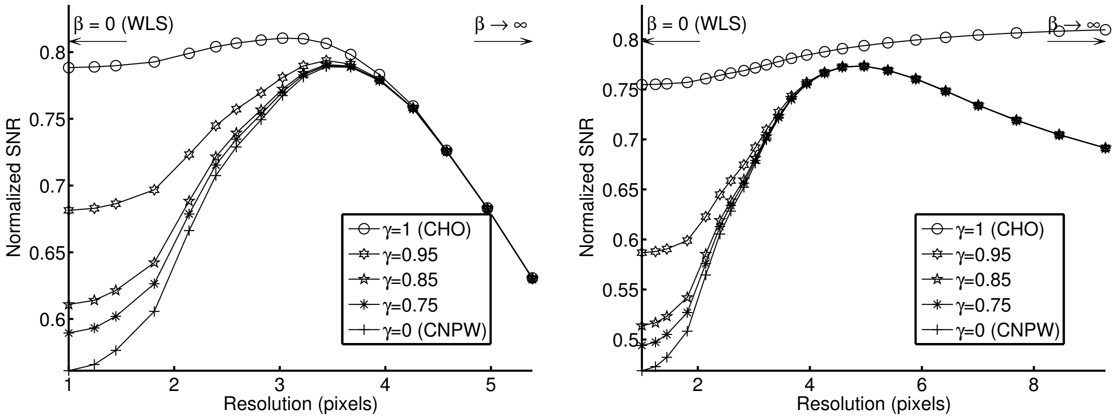

(a) S-DOG channels, const. internal variance

(b) S-DOG channels, prop. internal variance

Figure 3.5: SNR of CPPW observers versus QPWLS reconstruction resolution. Results are shown for five different degrees of prewhitening accuracy and two different internal noise models. All SNR values are computed from the exact expression (3.42) and normalized with respect to the upper bound (ideal SNR for the internal-noise-free observer) in (3.23).

β's, albeit a different range with each internal noise model, the SNR curves for different CPPW observers all lie close to each other and to the peak SNR attainable by the CHO. For this range of β's the observer's SNR performance in this known-location task is more robust to imperfect prewhitening by the observer.

Note that, in each of Fig. 3.5(a) and Fig. 3.5(b), the peak SNR occurs at the same β for all imperfectly prewhitening observers. Therefore, one could optimize regularization for the CNPW observer without the need to know the exact γ of an observer of interest. This invariance of the optimal β is a consequence of the specific form adopted for the CPPW template in (3.105) and may not hold exactly for other forms of prewhitening errors.

3.4.11 Rapid Computation of the Local Frequency Response of $\boldsymbol{\mathcal{F}}$

The expression (3.53) implies that the local frequency response $\boldsymbol{\lambda}$ of the Fisher information operator $\boldsymbol{\mathcal{F}}$ is computed by applying the operator directly to an object e_0 consisting of an impulse centered at the desired location. However, for large object sizes, this approach can become time-consuming if one wants to repeat it for a number of different locations (e.g., if one wants to optimize regularization locally for every pixel in the reconstructed image).

However, approximations of $\boldsymbol{\lambda}$ have been developed that are faster to compute and that involve a factorization of the frequency response. A key component of this factorization is a local certainty $\kappa(x_0, y_0)$ that depends on the location (x_0, y_0) of $\boldsymbol{e}_0$ [40]. In [19, 41, 104, 106, 145], where such analyses are incorporated into the calculation of detectability measures, they are performed entirely in the discrete domain. These analyses utilize a factorization of $\boldsymbol{\lambda}$ that involves uniform (angle-independent) local certainties $\kappa(x_0, y_0)$.

One can also adopt the strategy followed in [36] for roughness penalty design, *i.e.*, perform the analysis entirely in the continuous domain and discretize the final result. The continuous approach leads to a refinement of the fast approximation of $\boldsymbol{\lambda}$ that is not possible with the discrete approach. Alternatively, one can adopt a hybrid formulation and extend the analysis of [36] to the most realistic scenario where the true object is continuous but the measurements are discrete. We present this extension in Appendix B. This analysis derives a factorization of $\boldsymbol{\lambda}$ that involves angle-dependent local certainties $\kappa_{\varphi_j}(x_0, y_0), j = 1, \ldots, n_\varphi$. Since a practical implementation would employ DFT's, one would use a discretized version of the frequency response derived in Appendix B.

Fig. 3.6 compares these approximations for the PET example of previous sections. Since all approximations show very good agreement when they are used to compute the SNR of the CNPW observer, we show here an example where they are used to compute the SNR of the CHO (with D-DOG channels and both internal noise models). We obtain the approximate SNR from (3.72) as usual, but now we compute $\boldsymbol{\lambda}$ in three different ways: *(i)* Using (3.53) directly and applying $\boldsymbol{\mathcal{F}}$ to an impulse, as we did in previous examples, *(ii)* Using the angle-independent local certainty κ_0 from (B.3), as in other work [41, 104, 145], and *(iii)* Using the angle-dependent local certainties $\kappa_{\varphi_j}, j = 1, \ldots, n_\varphi$, from (B.2).

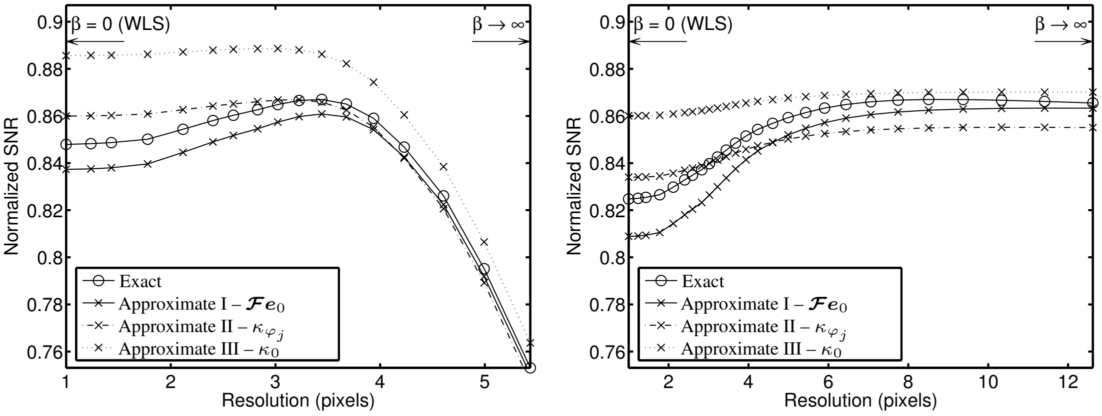

(a) D-DOG channels, const. internal variance

(b) D-DOG channels, prop. internal variance

Figure 3.6: Various approximations for the SNR of the CHO versus QPWLS reconstruction resolution. Results are shown for two different internal noise models. The exact and approximate SNR is computed from (3.67) and (3.72) respectively. Three different approximations for the local frequency response $\boldsymbol{\lambda}$ are used. All SNR values are normalized with respect to the upper bound (ideal SNR for the internal-noise-free observer) in (3.23).

As can be seen in Fig. 3.6, there is good agreement for all approximations, but approximation I follows the contour of the exact SNR curve somewhat more accurately. Approximation II generally leads to a reduction of the approximation error when compared with approximation III, while the two involve the same amount of computation, as can be seen by comparing (B.2) to (B.3).

3.5 Discussion

We have analyzed the performance of various commonly used mathematical observer models in the task of *known-location* signal detection. Our analysis shows that there are unregularized reconstruction methods that can lead to exactly optimal SNR in this task for several linear observer models from the literature (HO, NPW, ROI). Even for the CHO and the CNPW observer with internal noise, we have found (somewhat unconventional) unregularized reconstructors that lead exactly to the optimal SNR attainable with a given internal noise level. We also showed that, under some assumptions, some conventional unregularized reconstructors also lead approximately to this optimal SNR for these channelized observers.

When reconstruction is constrained to a penalized-likelihood method with a conventional roughness penalty, the CHO and the CNPW observer behave differently as the amount of regularization is varied. In the case of the CHO, unregularized WLS can lead to SNR very close to the peak SNR attainable with QPWLS for a given channel set and internal noise level. Therefore, optimizing regularized reconstruction does not lead to any significant improvement of SNR performance in comparison to unregularized reconstruction for the CHO in known-location detection tasks. This conclusion is in agreement with experimental results showing that the post-smoothing of noisy images does not lead to significant benefit for human observers, as far as performance in a known-location task is concerned [1].

This is in contrast to the case of the CNPW observer, where unregularized WLS leads to SNR significantly lower than the peak attainable with optimal QPWLS reconstruction. More generally, when regularization is constrained to a roughness penalty, channelized observers that perform less than perfect prewhitening can benefit from optimizing regularization.

Our conclusions about the CHO can also be contrasted to the results of Qi and Leahy, who analyzed the effect of regularization on the contrast-to noise ratio (CNR) [108]. As opposed to the SNR of the CHO, the CNR exhibits a pronounced peak for an intermediate amount of regularization. This may be explained by the fact that the CNR is equivalent to the SNR of an observer that utilizes an impulse as its template. Since this template performs no smoothing of its own, it is reasonable to expect that the optimal reconstructor should perform some smoothing. On the other hand, the template of the CHO is an optimally prewhitened version of the mean signal $\bar{\boldsymbol{f}}_s$ (typically not an impulse), so it smoothes the reconstructed image in an optimal manner itself. As a result, there is little benefit in terms of the SNR of the CHO from additional smoothing performed by the reconstructor.

The relatively small significance of regularization for many of the observer models throughout our analysis indicates that detection performance in tasks where the target signal location is known exactly is of limited use as a criterion for optimizing the free parameters of regularized reconstruction methods. Furthermore, the finding that some form or another of unfiltered backprojection is an optimal reconstructor for several of the observer models considered here, shows that resolution is not an essential image quality for these observers as far as known-location detectability is concerned.

Since known-location tasks are easy enough for several observers to perform optimally with poor resolution, it is important to examine whether there are more complicated tasks where regularization has a more significant effect. A relevant area of investigation is the analysis of regularized image reconstruction methods with respect to detectability in tasks with location uncertainty, which is discussed in the following chapter.

CHAPTER 4

Unknown Signal Location

We now turn our attention to tasks where the location of the target signal is unknown to the observer. This location uncertainty reflects medical practice more realistically, but it complicates the mathematical analysis of model observer performance significantly. As in the previous chapter, we begin by establishing our notation for the objects, measurements, and reconstruction methods that we will be considering. We then analyze the detection performance of the observer models for unknown-location tasks that were described in chapter 2, with the objective of optimizing regularized reconstruction methods with respect to this performance. We propose using tail probability approximations to facilitate analytical evaluation of the detection performance of these observers. We illustrate how these approximations can be used to evaluate the probability of detection (for low probability of false alarm operating points) for observer models that use channelized linear local test statistics.

4.1 The Detection Task

As in chapter 3, we adopt an additive model for the background $\boldsymbol{f}_b$ and the target signal. We assume that, when the target signal is present within the object, it is centered at one of a finite set of locations $\ell = 1, \ldots, n_{\mathrm{L}}$. We denote the target signal centered at location ℓ by $\boldsymbol{f}_{s,\ell}$. The detection/localization task at hand is thus a decision among the following $n_{\mathrm{L}} + 1$

hypotheses:

$$
\begin{aligned}
&H_0 \;:\; \boldsymbol{f} = \boldsymbol{f}_b && \text{(signal absent)} \\
&H_\ell \;:\; \boldsymbol{f} = \boldsymbol{f}_b + \boldsymbol{f}_{s,\ell} && \text{(signal present at location } \ell,\ \ell = 1, \ldots, n_{\mathrm{L}}\text{).} \qquad (4.1)
\end{aligned}
$$

The background $\boldsymbol{f}_b$ and the signal $\boldsymbol{f}_{s,\ell}$ are random and statistically independent under any of the signal present hypotheses H_ℓ, $\ell = 1, \ldots, n_{\mathrm{L}}$. We denote the expectations of $\boldsymbol{f}_b$ and $\boldsymbol{f}_{s,\ell}$ by $\bar{\boldsymbol{f}}_b$ and $\bar{\boldsymbol{f}}_{s,\ell}$ respectively. We denote their covariances by $\boldsymbol{\mathcal{K}}_b$ and $\boldsymbol{\mathcal{K}}_{s,\ell}$ respectively. The covariances $\boldsymbol{\mathcal{K}}_{s,\ell}$, $\ell = 1, \ldots, n_{\mathrm{L}}$ express variability in the profile of the target signal and not in its location.

We focus as usual on observer models that are applied on the reconstructed image $\hat{\boldsymbol{f}}$. We assume that $\hat{\boldsymbol{f}}$ is reconstructed from a measurement $\boldsymbol{y}$ that is acquired by a tomographic imaging system. Thus the moments of $\boldsymbol{y}$ under some hypothesis H_i are given by (3.4) and (3.5). Once again, we assume that the reconstructed image is given by $\hat{\boldsymbol{f}}(\boldsymbol{y}) = \boldsymbol{\mathcal{Z}}\boldsymbol{y}$ as in (3.9), for some linear reconstructor $\boldsymbol{\mathcal{Z}}$. As discussed in 3.1.3, the linearity assumption holds either exactly or approximately for several common tomographic reconstruction techniques.

Following the literature, we consider observers whose decision rule relies on computing a scalar local test statistic $t_\ell = t_\ell(\hat{\boldsymbol{f}}(\boldsymbol{y}))$ for each of the candidate locations $\ell = 1, \ldots, n_{\mathrm{L}}$ and then comparing the maximum test statistic $t_{\max} = \max_{\ell=1,\ldots,n_{\mathrm{L}}} t_\ell$ from (2.37) to a data-independent threshold τ. If $t_{\max} > \tau$, it is decided that the signal is present, otherwise it is decided that the signal is absent.

Our goal is to optimize the reconstructor $\boldsymbol{\mathcal{Z}}$ with respect to the detection performance of such an observer. Since we are ultimately interested in optimizing reconstruction locally, we quantify the observer's signal detection performance by a ROC curve defined around a given location $\ell \in \{1, \ldots, n_{\mathrm{L}}\}$. This would be a plot of the probability of detection (deciding that the signal is present when it is actually present at ℓ),

$$P_{\mathrm{D}}(\tau) \triangleq \mathsf{P}\{t_{\max} \geq \tau | H_\ell\}, \text{ for some } \ell \in \{1, \ldots, n_{\mathrm{L}}\}\,, \tag{4.2}$$

versus the probability of false alarm,

$$P_{\mathrm{FA}}(\tau) \triangleq \mathsf{P}\{t_{\max} \geq \tau | H_0\}\,. \tag{4.3}$$

The curve is traced by varying the decision threshold τ.

4.2 Threshold-Exceeding Probabilities of the Maximum Test Statistic

To trace a ROC or LROC curve for the test statistic $t_{\max}$, one must be able to calculate the CDF of $t_{\max}$, from which threshold-exceeding probabilities such as (4.2) and (4.3) can then be obtained.

The CDF of $t_{\max}$ is straightforward to compute if the local test statistics t_ℓ, $\ell = 1, \ldots, n_{\mathrm{L}}$, are statistically independent, because then the CDF of their maximum is given simply by the product of their individual CDFs. Khurd and Gindi used this fact to calculate the area under the LROC curve for observer models of the form (2.37)-(2.38) and images reconstructed with penalized-likelihood methods, under the assumption of independent, Gaussian-distributed local test statistics [69].

Swensson investigated an assumption weaker than the independence of the t_ℓ's. Specifically, he assumed that the maximum of the t_ℓ's over all locations that do not contain the signal is *(i)* Independent of the presence or absence of the signal in the object and *(ii)* Independent of the value of t_ℓ at a location containing the signal [126]. Under these conditions, Swensson derived the probabilities of false alarm, detection, and correct localization and showed that there is a point-to-point correspondence between the last two. Thus, maximizing the probability of detection is equivalent to maximizing the probability of correct localization. Consequently, maximizing the area under the ROC curve of $t_{\max}$ is equivalent to maximizing the area under its LROC curve.

Typical image reconstruction methods, and regularized ones in particular, produce images where the intensities at neighboring locations are correlated. Thus local test statistics at neighboring locations are also correlated, with more regularization leading to a wider autocorrelation function. Analyses that use independence assumptions apply only to tasks where the candidate locations are at a distance from each other that is greater than the autocorrelation width. For tasks where all the pixels within some search area are candidate locations, correlations between the local test statistics must be taken into account.

Threshold-exceeding probabilities such as (4.2) and (4.3) are difficult to obtain in closed form when the t_ℓ's are correlated, even if their joint distribution is available, since the exact distribution of the maximum of correlated random variables has the form of a multiple integral. When the CDF of $t_{\max}$ is not easy to compute, the "brute-force" approach to evaluating these probabilities is to perform a large number of time-consuming tomographic reconstructions of Monte Carlo simulated projection data and produce realizations of $t_{\max}$ from the reconstructed images. When optimizing some reconstruction parameter with respect to detection performance, this simulation has to be repeated for every value of the parameter of interest.

To avoid performing multiple reconstructions, investigators have proposed analytical approximations of the moments of the local test statistics and used these moments to directly produce Monte Carlo simulated realizations of the local test statistics. This strategy has been used to trace the LROC of $t_{\max}$ for images reconstructed from tomographic data using penalized-likelihood methods [70, 107]. It results in significant time savings when compared to the brute-force method, but it still requires performing simulations that add up in terms of computation when they are repeated for various values of one or more parameters of interest.

We propose an alternative approach to evaluating the performance of the maximum observer. Although a closed-form expression for the threshold-exceeding probabilities of $t_{\max}$ is generally not available for correlated Gaussian random fields, approximations of these probabilities for high values of the threshold τ have been developed. We use them here to trace a portion of the ROC curve.

By analyzing the Euler characteristic of excursion sets, Adler has derived approximations for the distribution tails of the maximum of a correlated random field [5, 6]. In particular, if $t_{\max} = \max_{\boldsymbol{x} \in S} T(\boldsymbol{x})$ is the maximum value of a 2-D stationary random field $T(\boldsymbol{x}) = T(x_1, x_2)$ over a set S, then the probability of $t_{\max}$ exceeding a high threshold τ is approximately

$$\mathsf{P}\{t_{\max} \geq \tau\} \approx \sum_{d=0}^{2} R_d(S)\rho_d(\tau), \tag{4.4}$$

where the factors $R_d(S),\ d = 0, 1, 2,$ depend on the geometry of the search area S and the functions $\rho_d(\tau),\ d = 0, 1, 2,$ depend on the distribution of $T(\boldsymbol{x})$.

If the search area S is a disk of radius r, then

$$R_0(S) = 1, \qquad R_1(S) = \pi r, \qquad R_2(S) = \pi r^2.$$

The approximation (4.4) is most accurate for search areas that are convex with sufficiently smooth boundaries, as discussed in [6]. If the stationary field $T(\boldsymbol{x})$ is Gaussian-distributed with *zero mean,* variance σ_T^2, and autocovariance function $R_T(\boldsymbol{x}) = R_T(x_1, x_2)$, then we can denote the dependence of $\rho_d(\cdot)$ on the moments of the field by writing

$$\rho_d(\tau) = \rho_d(\tau; \sigma_T, \Lambda_T),$$

where Λ_T is the matrix with ij-th element equal to

$$\{\Lambda_T\}_{ij} = -\partial^2 R_T(0,0)/\partial x_i \partial x_j, \;\; i, j = 1, 2.$$

In particular, we have

$$\rho_0(\tau;\sigma_T,\Lambda_T) \triangleq 1-\Phi(\tau/\sigma_T) \tag{4.5}$$

$$\rho_1(\tau;\sigma_T,\Lambda_T) \triangleq \frac{|\det\Lambda_T|^{1/4}}{2\pi\sigma_T}\, e^{-\tau^2/2\sigma_T^2} \tag{4.6}$$

$$\rho_2(\tau;\sigma_T,\Lambda_T) \triangleq \frac{|\det\Lambda_T|^{1/2}}{(2\pi)^{3/2}\sigma_T^3}\,\tau\, e^{-\tau^2/2\sigma_T^2}, \tag{4.7}$$

where $\Phi(\cdot)$ is the standard normal CDF. Approximations of the form (4.4) have been applied to the problem of detecting activation in functional neuroimaging [143, 144]. According to Worsley *et al.*, they lead to satisfactory accuracy for tail probabilities as high as 0.2 [144].

The analysis leading to the approximation in (4.4) assumes a continuous random field. When the field is defined solely on a lattice, the results hold asymptotically as the lattice becomes finer [6]. Thus, the expression above can be applied to approximate the tail distribution of the maximum test statistic in (2.37) when the discrete local test statistics $t_\ell, \ell = 1, \ldots, n_{\mathrm{L}}$ can be considered stationary and their mean and autocovariance is known.

Under the signal-absent hypothesis, the t_ℓ's may be considered stationary to within the accuracy of some local shift invariance approximations discussed in the following section. The same can be said in terms of the second-order statistics of the t_ℓ's under the signal-present hypothesis, assuming that the contribution of the signal profile variability to the covariance of the t_ℓ's is insignificant. However, the mean of the t_ℓ's cannot be considered constant throughout the search area in the presence of a spatially localized target signal like the ones that we are interested in. Thus, (4.4) is not appropriate for the signal-present hypothesis.

An alternative approach for approximating threshold-exceeding probabilities in the signal-present case follows the argument of Siegmund and Worsley in [119]. This approach decomposes the probability of detection at some location $\boldsymbol{x}_\ell$ as

$$\begin{aligned}
&\mathsf{P}\{t_{\max} \geq \tau | H_\ell\} \\
&\quad = \mathsf{P}\{T(\boldsymbol{x}_\ell) \geq \tau | H_\ell\} + \mathsf{P}\{t_{\max} \geq \tau, T(\boldsymbol{x}_\ell) < \tau | H_\ell\} \\
&\quad = 1 - \Phi\Big(\frac{\tau - \mu_T(\boldsymbol{x}_\ell)}{\sigma_T}\Big) \\
&\quad + \int_0^\infty \mathsf{P}\big\{t_{\max} \geq \tau \big| T(\boldsymbol{x}_\ell) = \tau - s, H_\ell\big\}\, \phi\Big(\frac{\tau - s - \mu_T(\boldsymbol{x}_\ell)}{\sigma_T}\Big) ds,
\end{aligned} \tag{4.8}$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ are the standard normal CDF and PDF respectively, and $\mu_T(\boldsymbol{x}) \triangleq \mathsf{E}[T(\boldsymbol{x})|H_\ell]$. Assuming that the maximum is most likely to occur near $\boldsymbol{x}_\ell$, *i.e.*, near the center of the target signal, the integrand in (4.8) will be non-negligible only for small values of s.

To derive an expression for the conditional probability inside the integral, Siegmund and Worsley also assume that the field $T(\boldsymbol{x})$, in the signal-present case and in the immediate neighborhood of the target signal, can be approximated as quadratic in $\boldsymbol{x}$. Under these assumptions, they show that the conditional probability of interest is approximately

$$\begin{aligned}
&\mathsf{P}\big\{t_{\max} \geq \tau \big| T(\boldsymbol{x}_\ell) = \tau - s\big\} \\
&\quad \approx \mathsf{P}\Big\{-\dot{T}(\boldsymbol{x}_\ell)' \,\mathsf{E}\Big[\ddot{T}(\boldsymbol{x}_\ell)\big| T(\boldsymbol{x}_\ell) = \tau - s\Big]^{-1} \dot{T}(\boldsymbol{x}_\ell) \geq 2s\Big\},
\end{aligned} \tag{4.9}$$

where $\dot{T}(\cdot)$ and $\ddot{T}(\cdot)$ denote the gradient and Hessian respectively. The conditional expectation in (4.9) can be obtained using

$$\mathsf{E}\Big[\ddot{T}(\boldsymbol{x})\big|T(\boldsymbol{x})\Big] = \mathsf{E}\Big[\ddot{T}(\boldsymbol{x})\Big] + \mathsf{Cov}\Big\{\ddot{T}(\boldsymbol{x}), T(\boldsymbol{x})\Big\}\, \frac{T(\boldsymbol{x}) - \mathsf{E}[T(\boldsymbol{x})]}{\mathsf{Var}\{T(\boldsymbol{x})\}}$$

and the linearity of derivatives. If this conditional expectation is proportional to the identity matrix, then (4.9) corresponds to a χ^2 CDF. Substituting this CDF in (4.8) results in the following approximation for the probability of detection:

$$\begin{aligned}
\mathsf{P}\{t_{\max} > \tau | H_\ell\} &\approx 1 - \Phi\Big(\frac{\tau - \mu_T(\boldsymbol{x}_\ell)}{\sigma_T}\Big) \\
&\quad + \phi\Big(\frac{\tau - \mu_T(\boldsymbol{x}_\ell)}{\sigma_T}\Big) \frac{1}{\sigma_T} \Big[\frac{\partial^2 R_T(\mathbf{0})}{\partial^2 x_i}\Big] \Big/ \Big[\frac{\partial^2 \mu_T(\boldsymbol{x}_\ell)}{\partial^2 x_i}\Big].
\end{aligned} \tag{4.10}$$

In the following, we will combine (4.4) and (4.10) with locally shift-invariant approximations to compute, respectively, the probabilities of false alarm and detection for observers with channelized linear local test statistics.

4.3 Channelized Linear Local Test Statistics

As mentioned in section 2.3.8, the observer models commonly used in the literature to predict human performance in unknown-location tasks utilize a linear channelized local test statistic. Similarly to channelized linear observers for known-location tasks, such observer models include a set of M bandpass filters, attempting to mimic the human visual system, and some template that typically has the form of a matched filter.

For every location of interest ℓ, the observer samples the output of the M bandpass filters to obtain a new local feature vector $\hat{\boldsymbol{c}}_\ell \in \mathbb{R}^M$, to which the corresponding *local template* $\boldsymbol{w}_\ell \in \mathbb{R}^M$ is then applied:

$$t_\ell = \boldsymbol{w}_\ell' \hat{\boldsymbol{c}}_\ell, \tag{4.11}$$

$$\hat{\boldsymbol{c}}_\ell = \boldsymbol{\mathcal{C}}_\ell' \Big(\hat{\boldsymbol{f}} - \mathsf{E}\Big[\hat{\boldsymbol{f}}_b \Big] \Big) + \boldsymbol{\varepsilon}_\ell, \tag{4.12}$$

where $\boldsymbol{\mathcal{C}}_\ell = [\boldsymbol{\mathcal{C}}_{1,\ell}, \ldots, \boldsymbol{\mathcal{C}}_{M,\ell}]$ consists of M operators. The mth of these operators applies the impulse response of the mth bandpass filter and samples the output at location ℓ. The *internal noise* vector $\boldsymbol{\varepsilon}_\ell \in \mathbb{R}^M$ models inherent uncertainty in the observer's decisions. Similarly to [70, 107], the mean of the reconstructed background,

$$\hat{\boldsymbol{f}}_b \triangleq \boldsymbol{\mathcal{Z}}(\boldsymbol{\mathcal{A}} \boldsymbol{f}_b + \boldsymbol{r}),$$

is subtracted from the reconstructed image in (4.12) to signify that the observer determines the most suspicious location by comparing intensities relative to the background, rather than by considering absolute intensities.

From (4.11) and (4.12), we can also view the local test statistic t_ℓ as the output of an image-domain local template $\boldsymbol{v}_\ell$ applied to the background-corrected reconstructed image and corrupted by noise:

$$t_\ell = \boldsymbol{v}_\ell' \Big(\hat{\boldsymbol{f}} - \mathsf{E}\Big[\hat{\boldsymbol{f}}_b\Big] \Big) + \delta_\ell, \tag{4.13}$$

where the image-domain template, a vector in the same space as $\hat{\boldsymbol{f}}$, is simply a linear combination of the channel responses:

$$\boldsymbol{v}_\ell = \boldsymbol{\mathcal{C}}_\ell \boldsymbol{w}_\ell, \tag{4.14}$$

and the noise is a linear combination of the internal noise of individual channels:

$$\delta_\ell = \boldsymbol{w}_\ell' \boldsymbol{\varepsilon}_\ell.$$

We can now write (4.13) in matrix-vector form as

$$\boldsymbol{t} = \boldsymbol{\mathcal{V}}' \Big(\hat{\boldsymbol{f}} - \mathsf{E}\Big[\hat{\boldsymbol{f}}_b\Big] \Big) + \boldsymbol{\delta}, \tag{4.15}$$

where the vector $\boldsymbol{t} \in \mathbb{R}^{n_\mathrm{L}}$ contains the local test statistics t_ℓ, $\ell = 1, \ldots, n_\mathrm{L}$, the vector $\boldsymbol{\delta} \in \mathbb{R}^{n_\mathrm{L}}$ contains the noise samples δ_ℓ, $\ell = 1, \ldots, n_\mathrm{L}$, and $\boldsymbol{\mathcal{V}} = [\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{n_\mathrm{L}}]$ contains the n_L local templates.

Typically, the internal observer noise vectors $\boldsymbol{\varepsilon}_\ell$, $\ell = 1, \ldots, n_\mathrm{L}$, are assumed to have Gaussian-distributed entries, in which case $\boldsymbol{\delta}$ is also Gaussian. Regardless of the exact distribution of the data, the term $\boldsymbol{\mathcal{V}}' \Big(\hat{\boldsymbol{f}} - \mathsf{E}\Big[\hat{\boldsymbol{f}}_b\Big] \Big)$ can be considered approximately Gaussian due to its linearity and the central limit theorem. Thus the local test statistics can be considered approximately Gaussian. Assuming that the t_ℓ's form a correlated Gaussian random field, we will use the approximation in (4.4) to evaluate the probabilities of detection and false alarm in (4.2), (4.3).

The approximation in (4.4) assumes stationarity. For a typical shift-variant tomographic system, the t_ℓ's are not globally stationary. However, as in chapter (3), we will consider the system to be locally shift-invariant and the object background to be locally stationary over a small area around each of the candidate locations $\ell = 1, \ldots, n_{\mathrm{L}}$. This implies that the t_ℓ's are approximately stationary locally around each of the candidate locations.

4.3.1 Moments of the local test statistics

The internal observer noise vectors $\boldsymbol{\varepsilon}_\ell$, $\ell = 1, \ldots, n_{\mathrm{L}}$, are typically modeled as zero-mean and statistically independent of each other. Following these assumptions on the internal noise, we find the mean $\boldsymbol{\mu}_t$ and covariance $\boldsymbol{\Pi}_t$ of the test statistic vector $\boldsymbol{t}$ in (4.15) to be, respectively,

$$\boldsymbol{\mu}_t = \boldsymbol{\mathcal{V}}' \left(\mathsf{E}\left[\hat{\boldsymbol{f}}\right] - \mathsf{E}\left[\hat{\boldsymbol{f}}_b\right] \right) \tag{4.16}$$

$$\boldsymbol{\Pi}_t = \boldsymbol{\mathcal{V}}' \, \mathsf{Cov}\left\{\hat{\boldsymbol{f}}\right\} \boldsymbol{\mathcal{V}} + \boldsymbol{\Pi}_{\boldsymbol{\delta}}, \tag{4.17}$$

where

$$\boldsymbol{\Pi}_{\boldsymbol{\delta}} = \mathsf{diag}\{\boldsymbol{w}_\ell' \boldsymbol{\Pi}_{\boldsymbol{\varepsilon}_\ell} \boldsymbol{w}_\ell, \; \ell = 1, \ldots, n_{\mathrm{L}}\}$$

is the $n_{\mathrm{L}} \times n_{\mathrm{L}}$ covariance matrix of $\boldsymbol{\delta}$ and $\boldsymbol{\Pi}_{\boldsymbol{\varepsilon}_\ell}$ is the $M \times M$ covariance matrix of the internal noise vector $\boldsymbol{\varepsilon}_\ell$.

Combining (3.4) and (3.5) with the linear reconstructor (3.9) and the assumption that $\boldsymbol{f}_b$ and $\boldsymbol{f}_{s,\ell}$ are independent, yields the moments of the reconstructed image $\hat{\boldsymbol{f}}$ for emission tomography under the each of the hypotheses in (4.1):

$$\mathsf{E}\left[\hat{\boldsymbol{f}} \middle| H_0\right] = \boldsymbol{\mathcal{Z}}(\boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_b + \boldsymbol{r}) = \mathsf{E}\left[\hat{\boldsymbol{f}}_b\right] \tag{4.18}$$

$$\mathsf{E}\left[\hat{\boldsymbol{f}} \middle| H_\ell\right] = \mathsf{E}\left[\hat{\boldsymbol{f}}_b\right] + \boldsymbol{\mathcal{Z}}\boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_{s,\ell}, \;\; \ell = 1, \ldots, n_{\mathrm{L}} \tag{4.19}$$

$$\mathsf{Cov}\left\{\hat{\boldsymbol{f}} \middle| H_\ell\right\} = \boldsymbol{\mathcal{Z}}\boldsymbol{\Pi}_\ell\boldsymbol{\mathcal{Z}}', \;\; \ell = 0, \ldots, n_{\mathrm{L}}, \tag{4.20}$$

where $\boldsymbol{\Pi}_\ell \triangleq \mathsf{Cov}\{\boldsymbol{y}|H_\ell\}$. In particular,

$$\boldsymbol{\Pi}_0 = \mathsf{diag}\left\{\boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_b + \boldsymbol{r}\right\} + \boldsymbol{\mathcal{A}}\boldsymbol{\mathcal{K}}_b\boldsymbol{\mathcal{A}}' \tag{4.21}$$

$$\boldsymbol{\Pi}_\ell = \boldsymbol{\Pi}_0 + \mathsf{diag}\left\{\boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_{s,\ell}\right\} + \boldsymbol{\mathcal{A}}\boldsymbol{\mathcal{K}}_{s,\ell}\boldsymbol{\mathcal{A}}', \;\; \ell = 1, \ldots, n_{\mathrm{L}}. \tag{4.22}$$

Combining (4.18) and (4.19) with (4.16) yields the following forms for the mean of the test statistic vector under the signal absent and present hypotheses respectively:

$$\boldsymbol{\mu}_{t|H_0} = \mathbf{0} \tag{4.23}$$

$$\boldsymbol{\mu}_{t|H_\ell} = \boldsymbol{\mathcal{V}}'\boldsymbol{\mathcal{Z}}\boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_{s,\ell}, \;\; \ell = 1, \ldots, n_{\mathrm{L}}. \tag{4.24}$$

Similarly, combining (4.20) with (4.17) yields the following form for the covariance of the test statistic vector under any of the $n_{\mathrm{L}} + 1$ hypotheses:

$$\boldsymbol{\Pi}_{t|H_\ell} = \boldsymbol{\mathcal{V}}'\boldsymbol{\mathcal{Z}}\boldsymbol{\Pi}_\ell\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{V}} + \boldsymbol{\Pi}_{\boldsymbol{\delta}}, \;\; \ell = 0, \ldots, n_{\mathrm{L}}. \tag{4.25}$$

For typical problem sizes, applying the operators $\boldsymbol{\mathcal{Z}}\boldsymbol{\mathcal{A}}$ and $\boldsymbol{\mathcal{Z}}\boldsymbol{\Pi}_\ell\boldsymbol{\mathcal{Z}}'$ explicitly to compute the moments of $\boldsymbol{t}$ in (4.24) and (4.25) is time-consuming for common statistical image reconstruction methods. Therefore, this approach is not practical when one needs these moments to compute measures of detectability for many values of some reconstruction parameter that is to be optimized.

As in section 3.4, we use locally shift-invariant analysis once again, this time to derive Fourier-domain approximations to the moments in (4.24) and (4.25). A similar approach to computing the moments of the local test statistics has been followed in [70, 107]. We will eventually combine these moment approximations with (4.4) to speed up the computation of the ROC-related probabilities (4.2) and (4.3), much like we did in section 3.4 to speed up the computation of the SNR in known-location tasks.

We illustrate the use of this approach in the case where the reconstructor $\boldsymbol{\mathcal{Z}}$ belongs to the QPWLS family in (3.65). Since we are interested in detecting small perturbations on the object background, we assume that the signal intensity is weak with respect to the background intensity. Then we can approximate

$$\boldsymbol{\Pi} \triangleq \mathsf{diag}\{\boldsymbol{\mathcal{A}}\,\mathsf{E}[\boldsymbol{f}] + \boldsymbol{r}\} \approx \mathsf{diag}\{\boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_b + \boldsymbol{r}\}\,. \tag{4.26}$$

For a $\boldsymbol{\mathcal{Z}}$ of the form (3.65) and under the weak-signal assumption (4.26), the operators of interest become

$$\boldsymbol{\mathcal{Z}}\boldsymbol{\mathcal{A}} = (\boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{R}})^{-1}\boldsymbol{\mathcal{F}} \tag{4.27}$$

$$\boldsymbol{\mathcal{Z}}\boldsymbol{\Pi}_0\boldsymbol{\mathcal{Z}}' \approx (\boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{R}})^{-1}(\boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{F}}\boldsymbol{\mathcal{K}}_b\boldsymbol{\mathcal{F}})(\boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{R}})^{-1} \tag{4.28}$$

$$\boldsymbol{\mathcal{Z}}\boldsymbol{\Pi}_\ell\boldsymbol{\mathcal{Z}}' \approx (\boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{R}})^{-1}[\boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{F}}(\boldsymbol{\mathcal{K}}_b + \boldsymbol{\mathcal{K}}_{s,\ell})\boldsymbol{\mathcal{F}}](\boldsymbol{\mathcal{F}} + \boldsymbol{\mathcal{R}})^{-1},\ \ell = 1, \ldots, n_{\mathrm{L}}. \tag{4.29}$$

In the following, we will apply locally shift-invariant approximations to the three operators above. As in section 3.4.3, we will assume to have a discrete representation of the object $\boldsymbol{f} \in \mathbb{R}^{n_{\mathrm{p}}}$ for the purposes of our Fourier-domain analysis.

4.3.2 Signal and Channel Spectra

We consider tasks where the mean target signal $\bar{\boldsymbol{f}}_{s,\ell}$ has the same shape at all candidate locations $\ell = 1, \ldots, n_{\mathrm{L}}$. Then $\bar{\boldsymbol{f}}_{s,\ell}$ is a copy of a common mean signal profile $\bar{\boldsymbol{f}}_s$ shifted to location ℓ. Similarly, we assume that the responses of the observer's channels at different locations in the image are shifted copies of each other.

The above assumptions imply that

$$\bar{\boldsymbol{f}}_{s,\ell} = \boldsymbol{\mathcal{U}}^{-1}\boldsymbol{E}_\ell\boldsymbol{X}_0 \tag{4.30}$$

$$\boldsymbol{C}_\ell = \boldsymbol{\mathcal{U}}^{-1}\boldsymbol{E}_\ell\boldsymbol{T}_0, \tag{4.31}$$

where $\boldsymbol{\mathcal{U}}$ is the discrete Fourier operator from (3.52),

$$\boldsymbol{E}_\ell \triangleq \sqrt{n_{\mathrm{p}}}\,\mathsf{diag}\{\boldsymbol{\mathcal{U}}\boldsymbol{e}_\ell\}$$

consists of the complex exponential that corresponds to a circular shift from the origin $(0, 0)$ to location ℓ, and $\boldsymbol{e}_\ell$ is an impulse centered at location ℓ. Finally, $\boldsymbol{X}_0$ is the spectrum of $\bar{\boldsymbol{f}}_s$, when the latter is assumed centered at $(0, 0)$, and the M columns of $\boldsymbol{T}_0$ contain the frequency responses of the observer's M channels, when their impulse responses are assumed centered at $(0, 0)$.

4.3.3 Locally Shift-Invariant Approximations

We now assume that the mean target signal profile $\bar{\boldsymbol{f}}_s$ is well-localized in space. Similarly to section 3.4.3, we approximate the system as locally shift-invariant around each of the locations $\ell = 1, \ldots, n_{\mathrm{L}}$. We also assume that, under any single hypothesis, the object can be approximated as locally stationary around each of the locations $\ell = 1, \ldots, n_{\mathrm{L}}$. The latter approximation is more likely to be accurate when the signal profile is deterministic under all signal present hypotheses (*i.e.*, $\boldsymbol{\mathcal{K}}_{s,\ell} = \mathbf{0}$, $\ell = 1, \ldots, n_{\mathrm{L}}$) or when the signal profile variability is negligible when compared to the background variability.

Thus we approximate, within some small spatial extent, the Fisher operator $\boldsymbol{\mathcal{F}}$ by (3.54), the regularization operator $\boldsymbol{\mathcal{R}}$ by (3.66), and the object covariance by

$$\boldsymbol{\mathcal{K}}_b \approx \boldsymbol{\mathcal{U}}^{-1} \boldsymbol{N}_0 \boldsymbol{\mathcal{U}} \tag{4.32}$$

$$\boldsymbol{\mathcal{K}}_b + \boldsymbol{\mathcal{K}}_{s,\ell} \approx \boldsymbol{\mathcal{U}}^{-1} \boldsymbol{N}_\ell \boldsymbol{\mathcal{U}}, \;\; \ell = 1, \ldots, n_{\mathrm{L}}, \tag{4.33}$$

where $\boldsymbol{N}_\ell \triangleq \mathsf{diag}\{\nu_k^\ell, k = 1, \ldots, n_{\mathrm{P}}\}$. The ν_k^ℓ's contain the local power spectrum of the object under hypothesis H_ℓ.

Applying the locally shift-invariant approximations (3.54), (3.66), (4.32) and (4.33) to the right-hand sides of (4.27) and (4.28) yields, respectively,

$$\boldsymbol{\mathcal{Z}}\boldsymbol{\mathcal{A}} \approx \boldsymbol{\mathcal{U}}^{-1} \boldsymbol{H}^{-1} \boldsymbol{\Lambda} \boldsymbol{\mathcal{U}} \tag{4.34}$$

$$\boldsymbol{\mathcal{Z}} \boldsymbol{\Pi}_\ell \boldsymbol{\mathcal{Z}}' \approx \boldsymbol{\mathcal{U}}^{-1} \boldsymbol{H}^{-2} (\boldsymbol{\Lambda} + \boldsymbol{\Lambda}^2 \boldsymbol{N}_\ell) \boldsymbol{\mathcal{U}}, \;\; \ell = 0, \ldots, n_{\mathrm{L}}, \tag{4.35}$$

where $\boldsymbol{H} \triangleq \boldsymbol{\Lambda} + \boldsymbol{\Omega}$.

4.3.4 Local Templates

As mentioned in section 2.3.8, observer models that compute the maximum of CHO- or CNPW-type local test statistics are the ones whose correlation with humans in unknown-location tasks has been investigated in the literature. These are the observer models on which we focus hereafter.

Maximum Channelized Hotelling Observer

The MaCHO applies a local template of the CHO form (2.34) at each candidate location within the search area and then compares the maximum test statistic to a threshold. The local template applied by this observer at some location ℓ is thus

$$\begin{aligned} \boldsymbol{w}_{\mathrm{CHO},\ell} &\triangleq \Big(\frac{1}{2}\,\mathsf{Cov}\{\hat{c}_\ell|H_\ell\} + \frac{1}{2}\,\mathsf{Cov}\{\hat{c}_\ell|H_0\}\Big)^{\dagger}\Big(\mathsf{E}[\hat{\boldsymbol{c}}_\ell|H_\ell] - \mathsf{E}[\hat{\boldsymbol{c}}_\ell|H_0]\Big) \\ &= (\boldsymbol{\mathcal{C}}_\ell'\boldsymbol{\mathcal{Z}}\check{\boldsymbol{\Pi}}_\ell\boldsymbol{\mathcal{Z}}'\boldsymbol{\mathcal{C}}_\ell + \boldsymbol{\Pi}_{\boldsymbol{\varepsilon}_\ell})^{\dagger}\boldsymbol{\mathcal{C}}_\ell'\boldsymbol{\mathcal{Z}}\boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_{s,\ell}, \end{aligned} \tag{4.36}$$

where $\check{\boldsymbol{\Pi}}_\ell \triangleq \frac{1}{2}\boldsymbol{\Pi}_\ell + \frac{1}{2}\boldsymbol{\Pi}_0$.

In the case of a reconstructor $\boldsymbol{\mathcal{Z}}$ of the QPWLS family (3.65), applying the locally-shift invariant approximations of (4.34) and (4.35) yields the following approximation to the MaCHO local template in (4.36):

$$\boldsymbol{w}_{\mathrm{CHO},\ell} \approx \Big[\boldsymbol{T}_0'\boldsymbol{H}^{-2}(\boldsymbol{\Lambda} + \boldsymbol{\Lambda}^2\check{\boldsymbol{N}}_\ell)\boldsymbol{T}_0 + \boldsymbol{\Pi}_{\boldsymbol{\varepsilon}_\ell}\Big]^{\dagger}\boldsymbol{T}_0'\boldsymbol{H}^{-1}\boldsymbol{\Lambda}\boldsymbol{X}_0, \tag{4.37}$$

where $\check{\boldsymbol{N}}_\ell \triangleq \frac{1}{2}\boldsymbol{N}_\ell + \frac{1}{2}\boldsymbol{N}_0$.

Maximum Channelized Non-Prewhitening Observer

The MaCNPW observer applies a local template of the CNPW form (2.35) at each candidate location within the search area and then compares the maximum test statistic to a threshold. The local template applied by this observer at a location ℓ is thus

$$\begin{aligned} \boldsymbol{w}_{\mathrm{CNPW},\ell} &\triangleq \mathsf{E}[\hat{\boldsymbol{c}}_\ell|H_{1,\ell}] - \mathsf{E}[\hat{\boldsymbol{c}}_\ell|H_0] \\ &= \boldsymbol{\mathcal{C}}_\ell'\boldsymbol{\mathcal{Z}}\boldsymbol{\mathcal{A}}\bar{\boldsymbol{f}}_{s,\ell}. \end{aligned} \tag{4.38}$$

In the case of a reconstructor $\boldsymbol{\mathcal{Z}}$ of the QPWLS family (3.65), applying the locally-shift invariant approximation of (4.34) yields the following approximation to the MaCNPW local template in (4.38):

$$\boldsymbol{w}_{\mathrm{CNPW},\ell} \approx \boldsymbol{T}_0' \boldsymbol{H}^{-1} \boldsymbol{\Lambda} \boldsymbol{X}_0. \tag{4.39}$$

4.3.5 Internal Observer Noise

Similarly to the known-location case in chapter 3, we consider here two different models for the observer's internal noise. For the first model,

$$\boldsymbol{\Pi}_{\boldsymbol{\varepsilon}_\ell} = \sigma_{\mathrm{int}}^2 \boldsymbol{\mathcal{I}}, \tag{4.40}$$

that is, some constant internal noise variance σ_{int}^2 is assumed for all channels and all locations ℓ.

For the second model,

$$\boldsymbol{\Pi}_{\boldsymbol{\varepsilon}_\ell} = c_{\mathrm{int}} \boldsymbol{T}_\ell' \boldsymbol{\mathcal{Z}} \check{\boldsymbol{\Pi}}_\ell \boldsymbol{\mathcal{Z}}' \boldsymbol{T}_\ell, \tag{4.41}$$

that is, the internal noise covariance is assumed to be proportional to the covariance in the channel outputs due to imaging noise and object variability, for some constant of proportionality c_{int}. In this case, we can also apply the locally shift-invariant approximation in (4.35) to rewrite the covariance (4.41) as

$$\boldsymbol{\Pi}_{\boldsymbol{\varepsilon}_\ell} \approx c_{\mathrm{int}} \boldsymbol{T}_0' \boldsymbol{H}^{-2} (\boldsymbol{\Lambda} + \boldsymbol{\Lambda}^2 \check{\boldsymbol{N}}_\ell) \boldsymbol{T}_0. \tag{4.42}$$

4.3.6 Frequency-Domain Approximations to Test Statistic Moments

Substituting the frequency-domain representations (4.30), (4.31), (4.34) and (4.35) into the moments (4.24) and (4.25) yields

$$\boldsymbol{\mu}_{t|H_\ell} \approx \boldsymbol{\mathcal{V}}' \boldsymbol{\mathcal{U}}^{-1} \boldsymbol{E}_\ell \boldsymbol{H}^{-1} \boldsymbol{\Lambda} \boldsymbol{X}_0, \;\; \ell = 1, \ldots, n_{\mathrm{L}},, \tag{4.43}$$

$$\boldsymbol{\Pi}_{t|H_\ell} \approx \boldsymbol{\mathcal{V}}' \boldsymbol{\mathcal{U}}^{-1} \boldsymbol{H}^{-2} (\boldsymbol{\Lambda} + \boldsymbol{\Lambda}^2 \boldsymbol{N}_\ell) \boldsymbol{\mathcal{U}} \boldsymbol{\mathcal{V}} + \boldsymbol{\Pi}_{\boldsymbol{\delta}}, \;\; \ell = 0, \ldots, n_{\mathrm{L}}. \tag{4.44}$$

In the approximations (4.43) and (4.44), we take $\boldsymbol{\Lambda}$, $\boldsymbol{N}_\ell$, and $\boldsymbol{\Omega}$ to contain the local frequency responses of the respective operators at location ℓ_c, where $\ell_c \in \{1, \ldots, n_\mathrm{L}\}$ corresponds to the center of the observer's search area. We assume that these frequency responses capture the approximate behavior of the respective operators throughout the search area, if the latter is not too big.

We can simplify further by using locally shift-invariant approximations for the templates $\boldsymbol{w}_\ell$ as well, such as (4.37) for MaCHO or (4.39) for MaCNPW. If we use the same versions of $\boldsymbol{\Lambda}$, $\boldsymbol{N}_\ell$, and $\boldsymbol{\Omega}$ (*i.e.*, the ones corresponding to the center ℓ_c) in these expressions for every $\boldsymbol{w}_\ell$, $\ell = 1, \ldots, n_\mathrm{L}$, then we are approximating the templates as location-independent throughout the search area, *i.e.*, $\boldsymbol{w}_\ell \approx \boldsymbol{w}_{\ell_c}$, $\ell = 1, \ldots, n_\mathrm{L}$. In the case of the MaCHO, for the templates to be location-independent, the internal noise covariance $\boldsymbol{\Pi}_{\boldsymbol{\varepsilon}_\ell}$ must be so as well. This is exactly true in the case (4.40) or holds approximately in the case (4.42). So we can write

$$\begin{aligned} \boldsymbol{v}_\ell \approx \boldsymbol{\mathcal{C}}_\ell \boldsymbol{w}_{\ell_c} &= \boldsymbol{\mathcal{U}}^{-1} \boldsymbol{E}_\ell \boldsymbol{T}_0 \boldsymbol{w}_{\ell_c} = \boldsymbol{\mathcal{U}}^{-1} \boldsymbol{E}_\ell \boldsymbol{V}_0 \\ &= \sqrt{n_\mathrm{p}}\, \boldsymbol{\mathcal{U}}^{-1} \,\mathsf{diag}\{\boldsymbol{V}_0\} \boldsymbol{\mathcal{U}} \boldsymbol{e}_\ell, \end{aligned}$$

where $\boldsymbol{V}_0 \triangleq \boldsymbol{T}_0 \boldsymbol{w}_{\ell_c}$. This implies

$$\boldsymbol{\mathcal{V}} \approx \sqrt{n_\mathrm{p}}\, \boldsymbol{\mathcal{U}}^{-1} \,\mathsf{diag}\{\boldsymbol{V}_0\} \boldsymbol{\mathcal{U}} \boldsymbol{\mathcal{I}}_\mathrm{SA}, \tag{4.45}$$

where $\boldsymbol{\mathcal{I}}_\mathrm{SA} \triangleq [\boldsymbol{e}_1 | \ldots | \boldsymbol{e}_{n_\mathrm{L}}]$ has as its columns the impulses $\boldsymbol{e}_\ell$, $\ell = 1, \ldots, n_\mathrm{L}$, corresponding to the n_L locations in the search area.

We can now combine (4.45) with the moment approximations in (4.43) and (4.44) to write

$$\boldsymbol{\mu}_{t|H_\ell} \approx n_\mathrm{p} \boldsymbol{\mathcal{I}}'_\mathrm{SA} \boldsymbol{\mathcal{U}}^{-1} \boldsymbol{\Phi}\, \boldsymbol{\mathcal{U}} \boldsymbol{e}_\ell, \;\; \ell = 1, \ldots, n_\mathrm{L}, \tag{4.46}$$

$$\boldsymbol{\Pi}_{t|H_\ell} \approx n_\mathrm{p} \boldsymbol{\mathcal{I}}'_\mathrm{SA} \boldsymbol{\mathcal{U}}^{-1} \boldsymbol{\Psi}\, \boldsymbol{\mathcal{U}} \boldsymbol{\mathcal{I}}_\mathrm{SA} + (\boldsymbol{w}'_{\ell_c} \boldsymbol{\Pi}_{\boldsymbol{\varepsilon}_{\ell_c}} \boldsymbol{w}_{\ell_c}) \boldsymbol{\mathcal{I}}, \;\; \ell = 0, \ldots, n_\mathrm{L}, \tag{4.47}$$

where

$$\begin{aligned}\boldsymbol{\Phi} &\triangleq \mathsf{diag}\left\{\frac{V_k^* X_k \lambda_k}{\lambda_k + \omega_k},\ k = 1, \ldots, n_{\mathrm{p}}\right\}\\ \boldsymbol{\Psi} &\triangleq \mathsf{diag}\left\{\frac{|V_k|^2 \lambda_k (1 + \lambda_k \nu_k^\ell)}{(\lambda_k + \omega_k)^2},\ k = 1, \ldots, n_{\mathrm{p}}\right\},\end{aligned}$$

and V_k, X_k are the kth elements of $\boldsymbol{V}_0$, $\boldsymbol{X}_0$ respectively.

4.4 Approximations of Detection Performance

We use the approximation in (4.4) to calculate the probability of false alarm (4.3) for channelized linear local test statistics and high detection thresholds τ. In particular, we write:

$$P_{\mathrm{FA}}(\tau) \approx \sum_{d=0}^{2} R_d(S) \rho_d(\tau; \sigma_{T|0}, \Lambda_{T|0}), \tag{4.48}$$

where we obtain $\sigma_{T|\ell}$, $\Lambda_{T|\ell}$, $\ell = 0$, from the local Fourier approximation (4.47):

$$\sigma_{T|\ell}^2 \approx \sum_{k=1}^{n_{\mathrm{p}}} \frac{|V_k|^2 \lambda_k (1 + \lambda_k \nu_k^\ell)}{(\lambda_k + \omega_k)^2} + \boldsymbol{w}_{\ell_c}' \boldsymbol{\Pi}_{\boldsymbol{\varepsilon}_{\ell_c}} \boldsymbol{w}_{\ell_c}. \tag{4.49}$$

In the special case where the target signal is variable in its location only and not in its shape, *i.e.*, $\boldsymbol{\mathcal{K}}_{s,\ell_c} = \mathbf{0}$, we have $\nu_k^{\ell_c} = \nu_k^0$. As a result, $\sigma_{T|\ell_c} = \sigma_{T|0}$ and $\Lambda_{T|\ell_c} = \Lambda_{T|0}$.

It would be possible, as in [148], to construct a similar approximation for the probability of detection (4.2), *i.e.*, to write

$$P_{\mathrm{D}}(\tau) \approx \sum_{d=0}^{2} R_d(S) \rho_d(\tau - \mu_{T|\ell_c}; \sigma_{T|\ell_c}, \Lambda_{T|\ell_c}), \tag{4.50}$$

obtaining $\sigma_{T|\ell_c}$, $\Lambda_{T|\ell_c}$ as above and $\mu_{T|\ell_c}$ from the local Fourier approximation (4.46):

$$\mu_{T|\ell_c} \approx \sum_{k=1}^{n_{\mathrm{p}}} \frac{V_k^* X_k \lambda_k}{\lambda_k + \omega_k}. \tag{4.51}$$

The approximation in (4.50) would be satisfactory if the mean of the test statistic field, $\mu_T(\boldsymbol{x}) \triangleq \mathsf{E}[T(\boldsymbol{x})|H_\ell]$, were constant throughout the search area under each of the signal present hypotheses. However, as mentioned previously, the mean cannot be considered constant throughout the search area in the presence of a spatially localized target signal. In this case, using a constant $\mu_T(\boldsymbol{x}) = \mu_T(\boldsymbol{x}_{\ell_c})$, as implied by (4.50), would lead to an overestimation of the probability of detection. The right-hand side of (4.50) is thus an upper bound for $P_{\mathrm{D}}(\tau)$, but it is by no means a tight one.

Thus, instead of using (4.50) for the probability of detection, we follow the approximation (4.10) for threshold-exceeding probabilities in the signal-present case, combining it with locally shift-invariant approximations of the local test statistic moments:

$$\begin{aligned} P_{\mathrm{D}}(\tau) \;&\approx\; 1 - \Phi\Big(\frac{\tau - \mu_{T|\ell_c}}{\sigma_{T|\ell_c}}\Big) \\ &+\; \phi\Big(\frac{\tau - \mu_{T|\ell_c}}{\sigma_{T|\ell_c}}\Big) \frac{1}{\sigma_{T|\ell_c}} \Big[\frac{\partial^2 R_{T|\ell_c}(\mathbf{0})}{\partial^2 x_i}\Big] \Big/ \Big[\frac{\partial^2 \mu_{T|\ell_c}(\boldsymbol{x}_{\ell_c})}{\partial^2 x_i}\Big]. \end{aligned} \tag{4.52}$$

We use (4.51) and (4.49) for $\mu_{T|\ell_c}$ and $\sigma_{T|\ell_c}$ respectively. Finally, we obtain the two derivatives in (4.52) by applying finite differences to (4.47) and (4.51) respectively.

4.5 Results

We provide an example of using the above approximations to evaluate the effect of regularization on the detection performance of the MaCHO and the MaCNPW observer, when applied to QPWLS-reconstructed images. We consider the same 2-D PET system model and mean background as in the examples of the previous chapter. The background $\boldsymbol{f}_b$ has a Gaussian autocorrelation function with a FWHM of 12 pixels and a standard deviation of 0.075. The target signal $\boldsymbol{f}_{s,\ell}$ has a known Gaussian profile with a FWHM of 4 pixels and an amplitude of 0.3, but unknown location ℓ. The search area is assumed to be a disk. Any pixel inside the disk is a candidate location for the target signal. Fig. 4.1 shows the mean background, the signal profile, and the largest of the concentric disk search areas that we consider here.

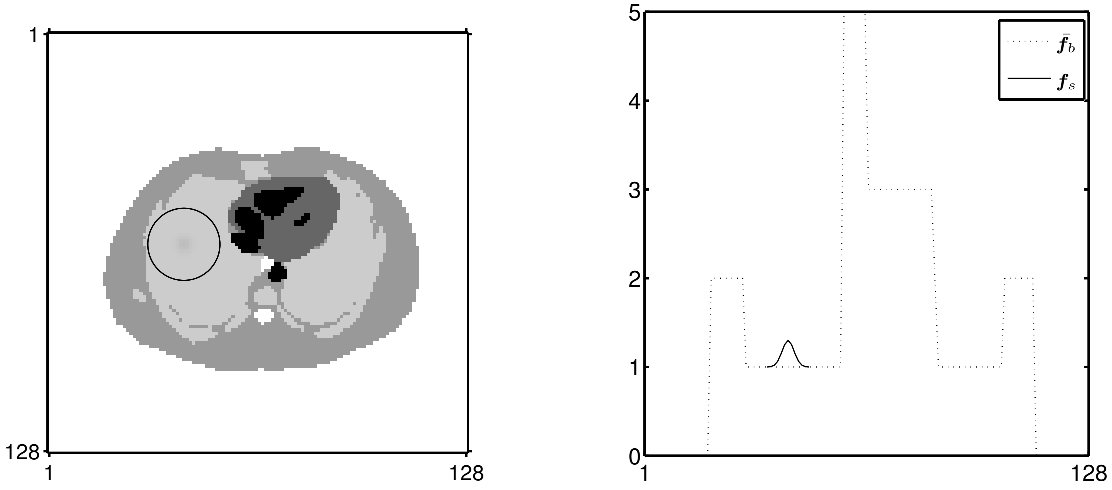

(a) Mean background $\bar{\boldsymbol{f}}_b$, target signal $\boldsymbol{f}_s$, and widest search area

(b) Profile of $\bar{\boldsymbol{f}}_b$ and $\boldsymbol{f}_s$

Figure 4.1: Mean background and target signal for an example of an unknown-location task. The figure shows the mean background (left) and a profile through the mean background (right) with the target signal superimposed. The largest of the search areas considered (diameter of 23 pixels) is also shown on the left.

We investigate observer models that combine CHO or CNPW local templates with one of the three channel sets (SQR, S-DOG, or D-DOG) and one of the two internal noise models (constant-variance or proportional-variance) discussed in section 4.3.5. We assume a variance of $\sigma^2_{\text{int}} = 0.005$ with the model in (4.40) and a proportionality constant of $c_{\text{int}} = 1.5$ with the model in (4.41). Detection performance is evaluated for QPWLS reconstruction with the uniform quadratic regularizer in (3.79) and various values of the regularization parameter β.

We first show plots of the AUC, to illustrate that the conclusions from our analysis of the SNR for the known-location task (chapter 3) generalize to the AUC for the unknown-location task.

4.5.1 Area Under the ROC Curve

The analytical approximations (4.48) and (4.52) apply only to high thresholds τ, so they cannot be used to trace the entire ROC and find the AUC. As a result, the AUC has to be computed by simulations.

Instead of reconstructing simulated data sets, we obtain the moments of the local test statistics for a range of different β's from the local shift-invariant approximations in (4.46) and (4.47). We then produce 5×10^5 realizations of the local test statistics, t_ℓ, $\ell = 1, \ldots, n_\mathrm{L}$, under each of the signal present and signal absent hypotheses, drawing from a Gaussian distribution with the respective moments. This yields multiple realizations of t_max under each hypothesis. For each β, we compare the realizations of t_max to a range of different thresholds τ to estimate $P_\mathrm{D}(\tau)$ and $P_\mathrm{FA}(\tau)$. We then compute the AUC for each β by numerical integration of $P_\mathrm{D}(\tau)$ versus $P_\mathrm{FA}(\tau)$.

Khurd and Gindi followed a similar approach to compute the area under the LROC curve, using local shift-invariant approximations of the local test statistic moments for HO, CHO, NPW, and CNPW local templates [70]. They compared the results to those produced by applying the observer templates to reconstructions of multiple noisy data sets. They reported very good agreement (for a range of signal contrasts) and a 10^3-fold reduction in computation time when using local shift-invariant analysis. Thus, we will not repeat the comparison of this method to reconstructions of simulated data sets.

We perform simulations to find the AUC of the MaCHO and the MaCNPW observer versus the amount of regularization. We also consider MaCPPW observers, *i.e.*, maximum-statistic observers that use a local test statistic of the CPPW form in (3.105), for three intermediate values of γ. Fig. 4.2 shows the AUC for each of these observers, with a search area diameter of 9 pixels and QPWLS reconstruction with the roughness penalty in (3.79) and various values of β. The plots on the left correspond to the constant internal vari-

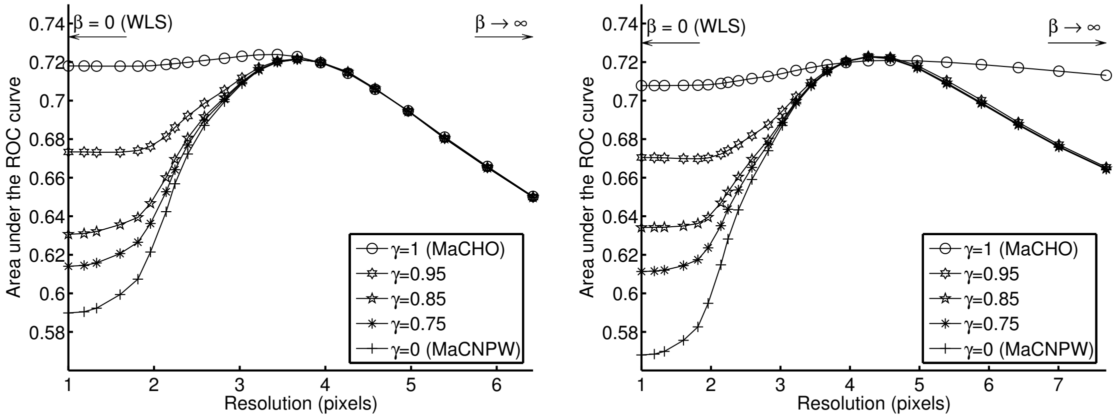

(a) SQR channels, const. internal variance

(b) SQR channels, prop. internal variance

Figure 4.2: AUC of MaCPPW observers versus QPWLS reconstruction resolution. Results are shown for five different degrees of prewhitening accuracy and two different internal noise models. The search area is a disk with a diameter of 9 pixels. The AUC is computed by simulations, using the moment approximations in (4.46) and (4.47).

ance model, whereas the ones on the right correspond to the proportional internal variance model. The abscissa of the plots represents the resolution of the QPWLS reconstruction at the center of the search area, as defined in section 3.4.6. We show only plots for the SQR channel set, but the results are similar with the S-DOG and D-DOG channel sets.

These plots show a pattern similar to the one seen in the SNR plots of section 3.4.10. That is, for the MaCHO observer, the peak AUC achieved with regularized QPWLS is very close to the AUC achieved with unregularized WLS. However, as the prewhitening capabilities of the observer deteriorate, the improvement in the AUC achieved by the optimal amount regularization becomes more and more significant.

The improvement afforded by optimal regularization for observers with imperfect prewhitening varies somewhat with the size of the search area. Fig. 4.3 shows the maximum increase in the AUC of the MaCNPW observer versus the diameter of the search area, for the SQR channel set and two different internal noise models. The increase is defined as the ratio $(\text{AUC}^* - \text{AUC}^0)/\text{AUC}^0$, where AUC^* and AUC^0 denote the AUC achieved by the

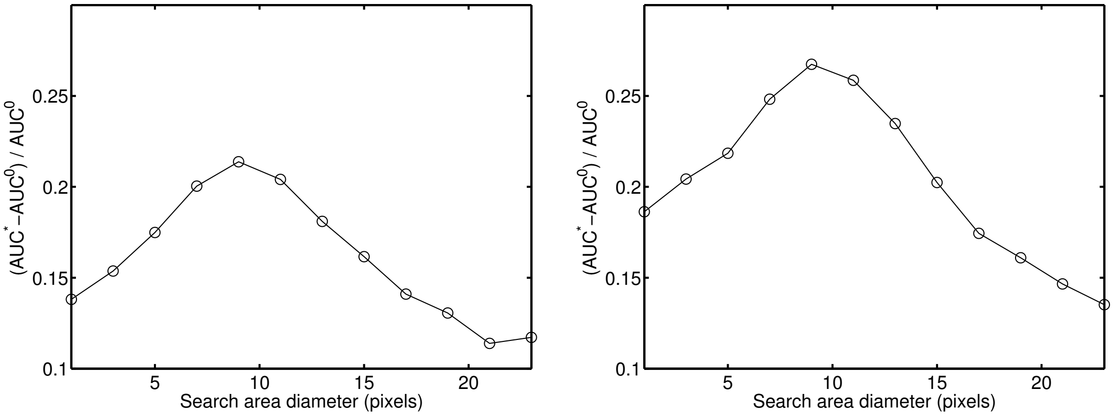

(a) SQR channels, const. internal variance

(b) SQR channels, prop. internal variance

Figure 4.3: AUC improvement of the MaCNPW observer versus search area diameter. Results are shown for two different internal noise models. The improvement is quantified as the ratio $(\mathrm{AUC}^* - \mathrm{AUC}^0)/\mathrm{AUC}^0$, where AUC^* and AUC^0 denote the AUC achieved with optimally regularized QPWLS and unregularized WLS respectively. The AUC is computed by simulations, using the moment approximations in (4.46) and (4.47).

MaCNPW observer with optimally regularized QPWLS and unregularized WLS respectively. The plots show that, with both internal noise models, the maximum improvement is achieved for some intermediate search diameter (in this case, 9 pixels).

Fig. 4.4 shows how the optimal amount of regularization varies with the size of the search area. Considering the range of QPWLS resolutions that achieve near-peak AUC for the MaCNPW observer, as seen in Fig. 4.2, the variation in optimal QPWLS resolution versus search diameter, as seen in Fig. 4.4, is not dramatic. Thus, a moderate uncertainty about the observer's search area would not result in a significantly different conclusion about the optimal amount of regularization.

The optimal QPWLS resolution in Fig. 4.4 ranges between 3 and 5 pixels, in most cases being closer to 4 pixels. Fig. 4.5 shows examples of QPWLS reconstructions with resolutions of 3, 4, and 5 pixels, all from the same data set with Poisson noise and the signal present in the center of the search area.

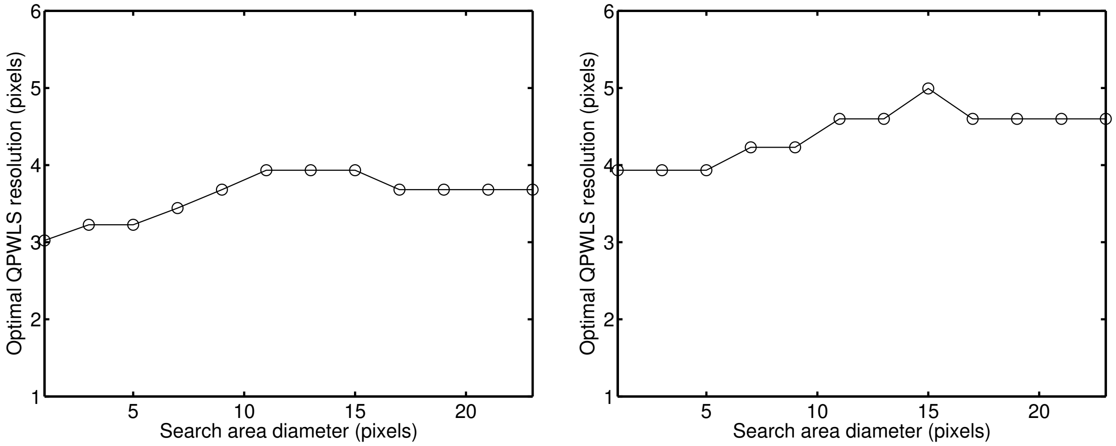

(a) SQR channels, const. internal variance

(b) SQR channels, prop. internal variance

Figure 4.4: Optimal QPWLS resolution for the MaCNPW observer versus search area diameter. Resolution is optimized in a maximum-AUC sense. Results are shown for two different internal noise models. The AUC is computed by simulations, using the moment approximations in (4.46) and (4.47).

4.5.2 High-Threshold Segment of the ROC Curve

Besides the area under the entire ROC curve, a partial area under the curve or even a single point on the curve (*i.e.*, the probability of detection at a fixed probability of false alarm) can be used as alternative figures of merit for detection performance. For example, one may be interested only in the observer's performance at low probabilities of false alarm. In the following, we use both the simulation method outlined in the previous section and the analytical approximations (4.52), (4.48) to plot the probability of detection at a fixed probability of false alarm versus the amount of regularization. Both the empirical and analytical methods perform a fast computation of the moments of the local test statistics based on local shift-invariance approximations. Thus the comparison herein focuses on the approximation introduced by (4.52) and (4.48), rather than the approximation introduced by assuming a shift-variant system to be locally shift-invariant. (As mentioned above, the latter approximation is investigated in [70].)

As an example, we consider the MaCNPW observer with the SQR channel set and no

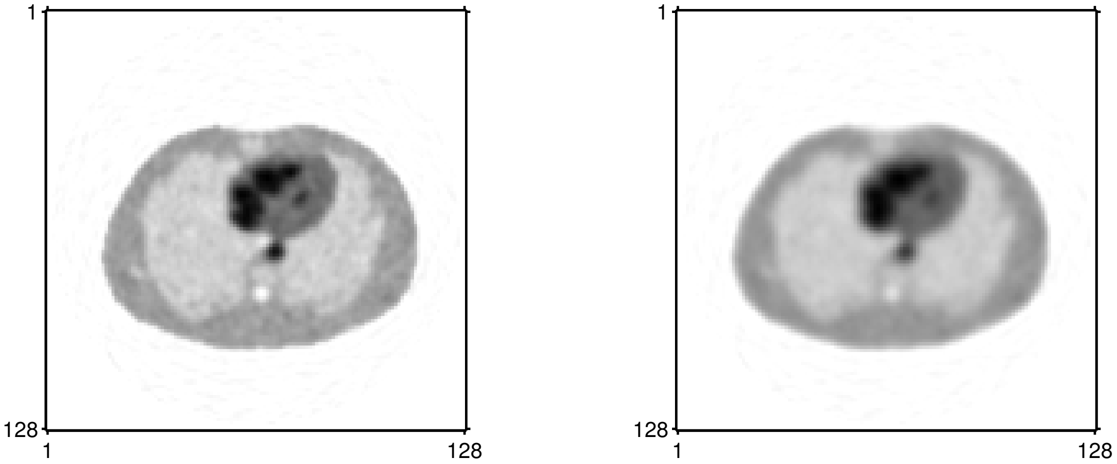

(a) QPWLS resolution: 3 pixels

(b) QPWLS resolution: 4 pixels

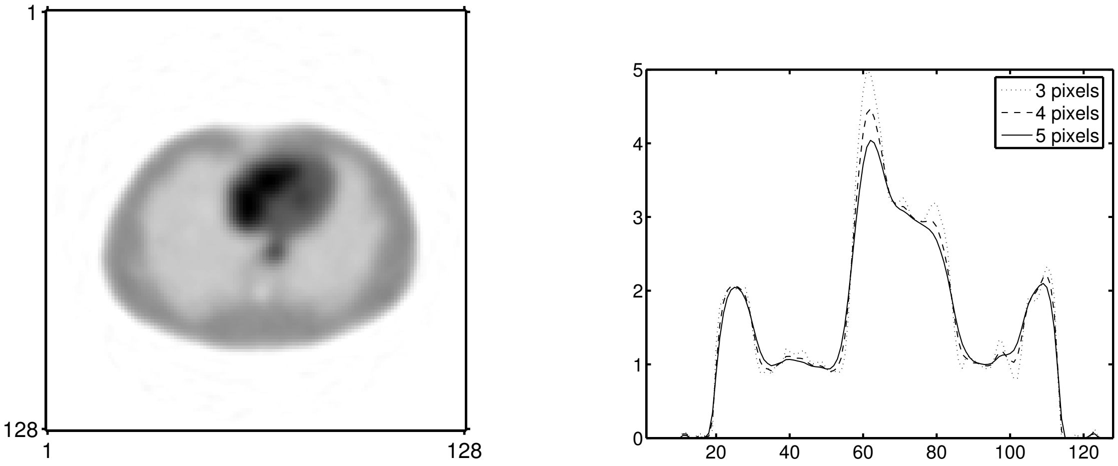

(c) QPWLS resolution: 5 pixels

(d) QPWLS reconstruction profiles

Figure 4.5: QPWLS reconstructions of a noisy Poisson data set with a resolution of 3, 4, or 5 pixels. The signal is present in the center of the search area.

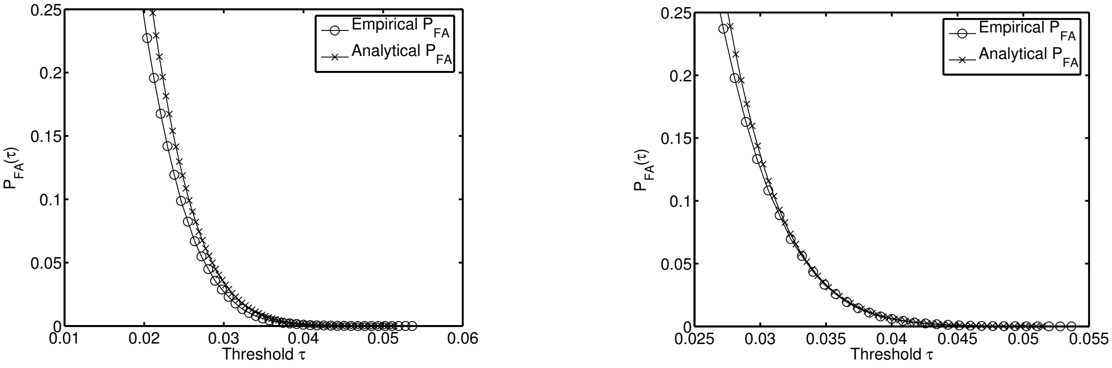

(a) Search diameter: 7 pixels

(b) Search diameter: 15 pixels

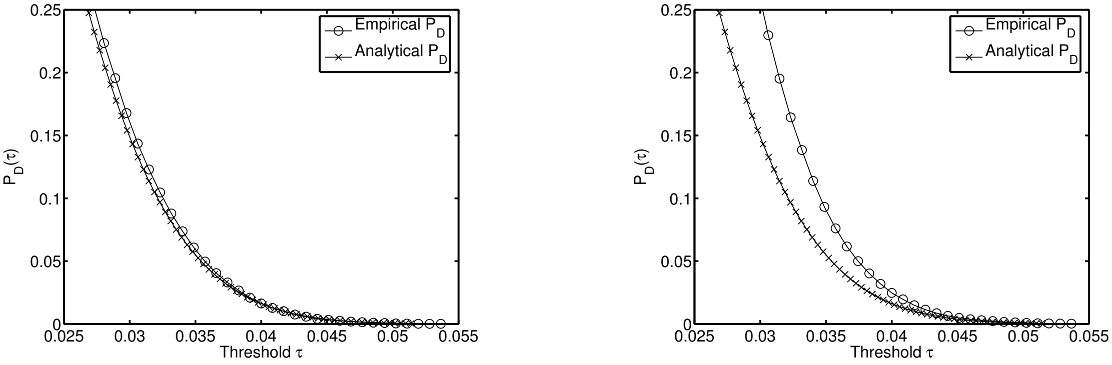

(c) Search diameter: 7 pixels

(d) Search diameter: 15 pixels

Figure 4.6: Empirical and analytical probabilities of false alarm and detection versus detection threshold. Results are shown for the MaCNPW observer with the SQR channel set, without internal noise, and for two different search area sizes.

internal noise. Using the analytical approximations (4.52) and (4.48), we calculate $P_{\mathrm{D}}(\tau)$ and $P_{\mathrm{FA}}(\tau)$ for different values of the threshold τ. We repeat this for each value of β. Fig. 4.6 compares these results to the empirical $P_{\mathrm{D}}(\tau)$ and $P_{\mathrm{FA}}(\tau)$ obtained from the simulations, for an intermediate QPWLS resolution of 3 pixels and for two different search area sizes. Fig. 4.6(a) and Fig. 4.6(c) show the probabilities of false alarm and detection, respectively, for a search diameter of 7 pixels. Fig. 4.6(b) and Fig. 4.6(d) show the probabilities of false alarm and detection, respectively, for a search diameter of 15 pixels.

The approximation (4.52) to the probability of detection assumes that the test statistic field is quadratic around the target signal location. This quadratic approximation is more

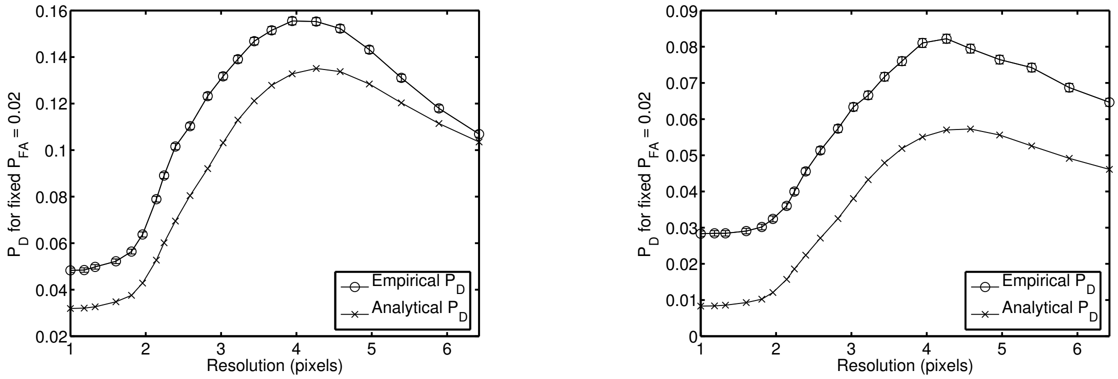

(a) Search diameter: 7 pixels

(b) Search diameter: 15 pixels

Figure 4.7: Plot of P_{D} (with error bars) versus QPWLS resolution for a fixed $P_{\mathrm{FA}} = 0.02$. Results are shown for the MaCNPW observer with the SQR channel set and without internal noise, for two different search area sizes.

accurate for a search diameter of 7 pixels, as shown in Fig. 4.6(c), since this search area size is comparable to the support of the target signal. For a larger search area, however, this approximation breaks down, as shown in Fig. 4.6(d).

We can now apply interpolation to these results, to find the probability of detection at a fixed value of the probability of false alarm for every β. Fig. 4.7 shows the linearly interpolated P_{D} for a fixed $P_{\mathrm{FA}} = 0.02$ versus QPWLS resolution, as obtained separately from the empirical and analytical results.

Even if the analytical approximations are not always accurate at predicting the value of the probability of detection, they follow the shape of the empirical curves and are maximized at nearly the same QWPLS resolution. Thus, analytically computed plots like the ones in Fig. 4.7 can be useful for choosing the regularization parameter β to optimize the probability of detection for a given low probability of false alarm within a small search area around some pixel. By increasing the probability of detection at some probability of false alarm, the area under the entire ROC curve is also increased.

For the small search diameter of 7 pixels, the analytical approach offers a 10-fold re-

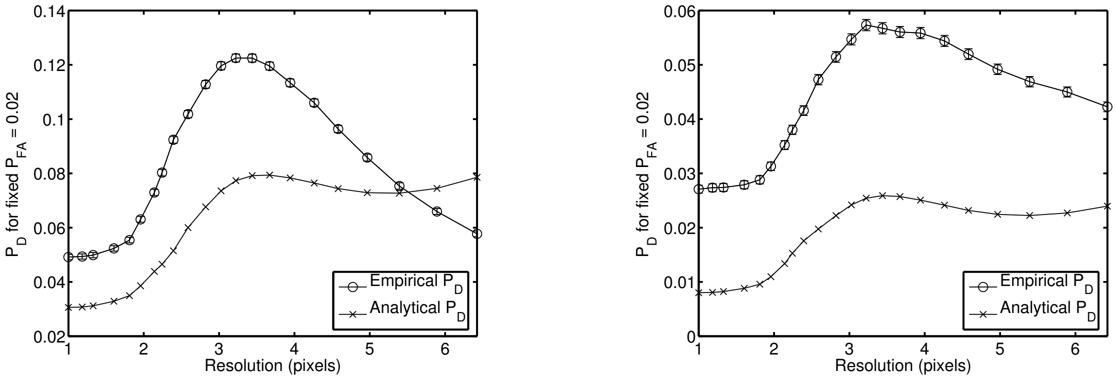

(a) Search diameter: 7 pixels

(b) Search diameter: 15 pixels

Figure 4.8: Plot of P_{D} (with error bars) versus QPWLS resolution for a fixed $P_{\mathrm{FA}} = 0.02$. Results are shown for the MaCNPW observer with the SQR channel set and the constant internal variance model, for two different search area sizes.

duction in computation, when compared to the empirical approach of simulating the local test statistics from the Fourier approximations of their moments. The computational savings increase with the size of the search area. They also increase with the number of realizations produced for the empirical approach.

For the example presented above, we did not include internal noise in the observer model. The reason is that internal noise would manifest itself as an impulse at the center $R_T(\mathbf{0})$ of the autocovariance function. In this case, the analytical approximations would break down, because they require the autocovariance function to be smooth so that its second derivative can be computed. Fig. 4.8 plots the interpolated P_{D} for a fixed $P_{\mathrm{FA}} = 0.02$ versus QPWLS resolution, for the same observer as above but with the addition of internal noise with constant covariance and $\sigma^2_{\mathrm{int}} = 0.005$.

Fig. 4.8 shows that the analytical approximation breaks down, especially for large amounts of regularization, where the internal noise dominates. The amount of internal noise determines whether the approximation is still usable for finding the optimal QPWLS resolution. This issue arises because, following convention, we assumed that the internal observer noise is uncorrelated between different candidate locations. Whether this assumption is crucial, however, must ultimately be validated by human observer studies.

4.6 Discussion

As indicated by the results of this chapter, conclusions drawn from our analysis of the SNR for the known-location task can be extended to the AUC for the unknown-location task. In other words, the observer's prewhitening capabilities are important in determining whether optimized regularization can improve detectability over unregularized reconstruction. However, the size of the search area also plays a role, with regularization being somewhat less beneficial for observers with very small or very large search areas.

We have considered here observers that search for the target signal within a local search area, rather than within the entire image. We have made this choice for two reasons. First, our locally shift invariant approximations to the moments of the local test statistics are accurate within a local region. If the system is shift-variant and/or the object background is not stationary, these approximations are not accurate within the entire image. Second, our ultimate goal is to use the analyses above to optimize the amount of regularization locally at each pixel in the image. It seems reasonable for the optimal amount of regularization at some pixel to be determined by the image statistics within a local area around that pixel.

However, it is not clear whether an observer model that bases its decisions on a well-delineated search area, disregarding pixels beyond this area, is an appropriate model for a human. An unintuitive characteristic of such models is that when the search area is small in comparison to the width of the autocorrelation function, the maximum $t_{\max}$ is most likely to occur along the boundary of the search area. The nature of the search performed by human observers in unknown-location tasks needs to be further investigated to improve these models.

Using the analytical approximations proposed in this chapter, one can rank and optimize image reconstruction methods with respect to unknown-location detectability without the need for time-consuming Monte Carlo simulations. Improving the accuracy of these approximations for a wider range of search area sizes and for non-zero internal noise are topics for future investigation. Furthermore, since the threshold-exceeding probability approximations used here have been derived for continuous random fields, the analysis can be applied naturally to tomographic reconstruction methods that yield continuous-space images $\hat{\boldsymbol{f}}$.

CHAPTER 5

Conclusion

5.1 Summary

Our analysis of tasks where the location of the target signal is known a priori shows that regularization is not necessary for the HO, NPW, and ROI observers. These observers can achieve exactly optimal detection performance in the known-location task on images reconstructed by some form of unfiltered backprojection. This result holds even when the background and/or the target signal profile is random.

The conclusion extends to the CHO and CNPW observers with a proportional internal covariance model. The difference is that the unregularized reconstructors that ensure exactly optimal performance for these channelized observer models are more unconventional than simple backprojection. In the case of the CHO, however, more conventional unregularized reconstructors can be shown to approximately achieve the optimal performance attainable with a certain channel set and internal noise level, to within the accuracy of some approximations of local shift invariance.

Besides seeking image reconstruction methods that optimize the detection performance of channelized observer models, we have also investigated the impact of a family of practical regularized reconstruction methods on this performance. We have shown that optimizing the amount of regularization applied by such a method has limited impact on the CHO, which can achieve similar performance with the unregularized reconstructor in this family and with the optimally regularized one. This is not the case for the CNPW observer, however. Because of its inability to prewhiten, this observer can achieve performance comparable to that of the CHO only when the channel outputs are close to white. As we have illustrated, this occurs only for some intermediate amount of regularization.

We have extended this result to a family of observers that can only partially prewhiten the channel outputs. Our conclusion is that the less accurate the observer's prior knowledge of second-order image statistics, the more significant the impact of choosing the optimal amount of regularization for detection performance.

Other than the observer's prewhitening capabilities, the statistics of its internal noise also play an important role. The two internal noise models that we have considered, one with covariance that changes with the reconstruction method and one with constant covariance, have led to different behavior with respect to regularization. On the other hand, the choice of channel profiles seems to be less important in this regard. The channel sets that we have considered, either with non-overlapping or with overlapping radially symmetric passbands, and with different numbers of channels, have led to similar general trends in the variation of detection performance versus regularization.

Several of our arguments regarding the impact of regularization on signal detectability are based on local shift invariance approximations. Comparisons of the approximate and the exact expressions for the SNR indicate that these approximations can be trusted in a variety of cases. Because the exact expressions become particularly cumbersome when they must be applied to multiple values of a reconstruction parameter and multiple locations within the object, the frequency-domain approximations derived here facilitate the theoretical evaluation and ranking of reconstruction methods significantly.

Finally, several of the conclusions we have drawn from analyzing detectability versus regularization for known-location tasks, also apply to unknown-location ones. This illustrates the importance of the simple known-location task, which is more amenable to analysis, in providing intuition about more complicated tasks. We have also investigated analytical approximations which should facilitate the study of model observer performance in tasks with location variability,

5.2 Future Work

Our analysis indicates that there is no universally optimal reconstruction method, since the optimal choice varies with the observer and detection task at hand. In other words, choosing the reconstruction method to optimize detection performance translates to choosing the observer model and the detection task for which the reconstructed images will be optimal. A necessary step towards applying our analyses to tune regularized reconstruction methods is establishing appropriate choices for the observer and task. Also, it would be of interest to investigate the sensitivity of the resulting optimal regularization parameters to these choices.

When reconstructed images are to be viewed by human observers, the appropriate observer model is the one that best predicts human performance. Note that the model does not need to capture every possible aspect of the human visual system, but it needs to behave similarly to humans in response to varying amounts of regularization in the reconstruction. Our analysis provides several clues regarding the sensitivity of the optimal regularization to the choice of observer model. In particular, our analysis illustrates the important effect of the observer's prewhitening capabilities and internal noise model on the optimal choice of regularization. In tasks with location variability, the size of the search area also appears to be important. Given a detection task, one should make sure that these aspects of the observer model agree with the ones exhibited by humans in the specific task. Although we have focused here on 2-D observer models, our analysis can be extended to models that act on 3-D reconstructed images. The analysis can serve as a guideline for observer properties to investigate when designing experiments that test the correlation of model and human observers.

The choice of an appropriate task involves various factors, such as target signal profile and location variability, background variability, and imaging system properties. Appropriate choices for these factors depend on the application of interest. Using our analysis, one can investigate how the optimal regularization parameters vary with each of these factors, *e.g.*, with the width of the mean signal profile or with the resolution of the imaging system. Various models for the background variability of medical images have been incorporated in model observer templates either through ensemble statistics or through sample statistics with the use of simulated images [2, 18, 113]. It would also be of interest to know whether these different choices result in significantly different optimal regularization parameters.

Furthermore, it would be interesting to explore whether imperfect knowledge of any properties of the target signal, background, and/or imaging system on the side of the observer or the reconstructor results in an increased importance of optimizing regularization. For example, we have already shown in section 3.4.10 that imperfect knowledge of the second-order image statistics on the side of the observer makes it more important to choose an optimal, non-zero value for the regularization parameter.

Another example of this would be imperfect knowledge about the imaging system on the side of the reconstructor. We have shown that a simple backprojection is the optimal reconstruction method for several observers and tasks, but this holds for a backprojection that uses the exact adjoint $\boldsymbol{\mathcal{A}}'$ of the system operator $\boldsymbol{\mathcal{A}}$ that produced the data. In practice, even if the exact system operator were known, it would be typically very difficult to implement. As a result, practical reconstruction methods use simplified models of imaging systems, leading to less than optimal detection performance. It would thus be of interest to know how errors in the system model used by the reconstructor affect the optimal regularization parameters with respect to detection performance. Furthermore, it is plausible that different system models result in improved detectability for different kinds of signals. For example, one could design the volume elements assumed by the system model to achieve optimal detectability of specific signal profiles.

Our locally shift-invariant analysis has focused on tasks where the mean signal profile $\bar{\boldsymbol{f}}_s$ resembles an impulse. However, it is straightforward to extend it to a $\bar{\boldsymbol{f}}_s$ that is a linear combination of signals resembling impulses. This model encompasses the so called *Rayleigh tasks,* where the observer must discriminate between a signal consisting of one bump and a signal consisting of two bumps. Rayleigh tasks have been analyzed for ideal (GLRT) observers [117]. These tasks, potentially combined with location uncertainty, are intuitively appealing for the purposes of optimizing parameters that control resolution in reconstructed images.

The approximations to the probabilities of detection and false alarm that we used in chapter 4 assume that the observer searches for the maximum local test statistic over a set of possible locations. However, generalizations to this approximation exist that support searching over signal width uncertainty [119] or even rotation uncertainty [116] in addition to location uncertainty. These generalizations could be used to incorporate more flexible models in the optimization of regularization with respect to unknown-location detection performance.

Furthermore, the asymptotic approximations for the distribution tails of $t_{\max}$ that we used in chapter 4 hold for the maximum of Gaussian random fields. The justification for the approximate Gaussianity of the local test statistics was based either on the linearity of the t_ℓ's and the central limit theorem, or on the approximate Gaussianity of images reconstructed by penalized-likelihood methods. However, it is plausible that such assumptions of approximate Gaussianity might not be accurate enough when one is dealing with distribution tails. For such cases, it should be noted that the literature provides asymptotic approximations for the distribution of the maximum of random fields with different statistics, including smoothed Poisson [110].

The analytical performance approximations that we have proposed can be applied to tune user-specified parameters of image reconstruction methods, or even design parameters of tomographic imaging systems, with respect to detection performance in known- or unknown-location tasks. In particular, we would like to establish methods for designing the penalty functions of penalized-likelihood image reconstruction methods with respect to detection performance. The resulting penalties would most likely be space-variant, as is the case when adjusting the penalty to achieve uniform and isotropic spatial resolution [124]. However, the detectability and spatial-resolution criteria would not necessarily yield the same optimal penalties, as suggested by preliminary results of Qi [105], where penalties that achieve isotropic resolution are shown to degrade detection performance in a simple SKE/BKE task. Extending the analysis to penalized-likelihood reconstruction with non-quadratic, edge-preserving penalties is a topic for future investigation.

Developing analytical methods for optimal penalty design with respect to signal detectability would eliminate the need for a subjective and time-consuming trial-and-error search and thus increase the usability of penalized-likelihood image reconstruction algorithms. Similar methods can be developed to tune the resolution of other reconstructor families, such as the popular post-filtered maximum-likelihood. Ultimately, the effectiveness of such methods in improving detection performance would have to be validated by experiments with human observers.

APPENDICES

APPENDIX A

A Comparison of Two System Models for 3-D SPECT

A.1 Introduction

In this appendix we compare system models for 3-D Single-Photon Emission Computed Tomography (SPECT) that compensate for non-uniform photon attenuation as well as the depth-dependent response inherent in SPECT systems. In particular, we consider here rotation- and blob-based implementations of the system model.

Several rotation-based models have been found to achieve superior accuracy when compared to the simple line-integral model in 2-D reconstruction [30]. Furthermore, models that represent images in terms of spherically symmetric blobs have been applied to 3-D Positron-Emission Tomography (PET) and have been found to possess better bias-variance properties than models using cubic voxel representations [91]. The blob-based methods are computationally tractable for shift-variant PET system models when the "footprint" (2-D projection) of the 3-D blob is pre-computed and saved. However, it is not clear if the same holds true for SPECT, where the depth-dependent blur causes these footprints to expand considerably with increasing distance from the detector. Also, for the purpose of SPECT reconstructions, it is of interest to know how these blob-based models compare in terms of bias-variance properties and speed to the rotation-based models.

We have implemented both types of system models for fully 3-D SPECT. We present an overview of the two implementations in Section A.2 and we analyze them in terms of computational load and memory usage in Sections A.3 and A.4 respectively. We present results from reconstructions of simulated phantom data in Section A.5 and further discuss those results in Section A.6.

A.2 System Models

For both types of system models we have implemented the back-projector as the exact adjoint of the respective projector. Using a simpler back-projector can speed up reconstruction but it does not preserve the convergence properties of the iterative reconstruction algorithm and it introduces artifacts that may accumulate with every iteration [155]. To ensure a comparison untarnished of such artifacts, we utilized adjoint pairs.

Both types of system models compensate for non-uniform attenuation. For the sake of speed, the attenuation factors are pre-computed for every volume element (cubic voxel or blob) at every projection angle. Utilizing the central ray approximation, an attenuation factor is computed by summing over the attenuation map along the perpendicular line from the center of the volume element to the detector.

A.2.1 Rotation-Based Projector

The rotation-based projector we use follows the processing steps listed in [154] and performs the rotation step using the fast and accurate three-pass separable rotation method described in [130]. The three-pass method decomposes the rotation of a 2-D image into three sets of 1-D interpolations. These interpolations are equivalent to applying appropriate non-integer shifts first to each row, then to each column, and then again to each row of the image.

Projectors utilizing the three-pass method with linear and cubic interpolation have been compared in [30] and the latter outperformed the line-integral projector in projection accuracy. However, cubic interpolation introduces negative values in the rotated image, which is unnatural in emission tomography, where image values represent photon counts. Thus, we use linear interpolation hereafter.

The rotation-based 3-D SPECT projector implemented for this comparison compensates for depth-dependent detector response. For every projection angle, it first rotates the input emission volume around its z-axis by applying the three-pass method to every slice of the volume separately. It then applies to every voxel the appropriate attenuation factor. Subsequently, every image plane that corresponds to a different distance from the detector face is blurred by a 2-D convolution kernel that models the detector response at the corresponding distance. Finally, the voxels are summed along the direction that is perpendicular to the detector to produce a 2-D projection of the emission volume.

A.2.2 Blob-Based Projector

Rotationally symmetric volume elements can provide a basis that is more natural than cubic voxels for representing smooth radioactivity distributions. Moreover, their symmetry makes them appealing for some image reconstruction applications since the line integrals through them are independent of projection angle and only depend on the distance of the line of integration from the center of the element [82, 83].

Desired properties of the rotationally symmetric basis functions are effective band-limitedness (to match the band-limitedness of tomographic data) and finite spatial support (to facilitate computation). By these criteria, a choice that is nearly optimal yet easy to compute is the Kaiser-Bessel blob:

$$\text{(A.1)}\quad b_{R,\alpha,m}(r) = \begin{cases} [\sqrt{1-(\frac{r}{R})^2}]^m \; I_m[\alpha\sqrt{1-(\frac{r}{R})^2}] \;/\; I_m(\alpha), & 0 \le r \le R \\ 0, & \text{otherwise,} \end{cases}$$

where I_m denotes the modified Bessel function of the first kind of order m, R controls the blob support, α controls the blob taper, and m controls the number of continuous derivatives at $r = R$.

We have implemented a 3-D SPECT projector utilizing Kaiser-Bessel blobs, similar to the one described for PET in [91]. The additional step in our projector is that the blobs at each different distance from the detector now need to be blurred by the corresponding point-spread function (PSF) of the SPECT detector.

Of course, different points within a 3-D blob lie at different distances from the detector and therefore correspond to different PSFs. However, when the Full Width at Half-Maximum (FWHM) of the depth-dependent system blur changes slowly with distance relatively to the blob size, as illustrated in Fig. A.1, we can assume that the blur FWHM is approximately the same throughout a blob as it is at the blob's center. This approximation simplifies the implementation significantly, since it makes projecting the blurred blob equivalent to blurring the projections of the blob. Thus, we can pre-compute and save a set of finely sampled 2-D footprints of the blob, blurred by the detector PSF at each of a set of distances from the detector.

The spacing of the stored samples for a single footprint, as well as the number of different distances for which footprints are stored, depends on the required accuracy of the implementation. Clearly, although the continuous projections of the rotationally symmetric volume elements are angle-independent, the discrete projections are not necessarily so. In general, the same footprint is sampled at slightly different points (by a small shift) at different angles. However, we do not have to store a different footprint for each angle, as we may have done if applying this method to, say, cubic voxels. Storing a finely sampled version of the footprint and picking out samples based on the nearest-neighbor criterion will yield adequate accuracy for SPECT. A similar nearest-neighbor approach is taken with respect to distance from the detector when choosing which of the stored footprints to

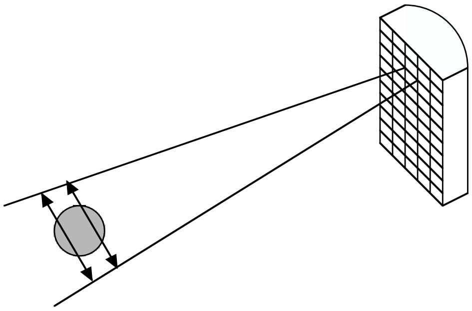

Figure A.1: Approximating the FWHM of the system blur as constant throughout a single blob. If the FWHM of the system blur does not increase very fast with distance (relatively to the size of the blob), we assume that it is approximately constant throughout the blob.

use. In some sense, this approximation is used implicitly by the rotation-based model as well, since after rotation by a non-trivial angle, some edges of the cubic voxel will be at different distances from the detector than others.

We also model the attenuation as being constant throughout a single blob, as we did for cubic voxels in the rotation-based implementation. We arrange the blobs over a uniform 3-D Cartesian grid for simplicity, although more efficient blob arrangements are also possible [90].

Taking all of the above into account, a projection operation with the blob-based model is performed as follows. At each projection angle and for each blob in the volume, we select the footprint that corresponds to the distance of the blob from the detector in a nearest-neighbor sense. For every sample on this footprint that corresponds to a projection bin in a nearest-neighbor sense, we multiply the sample by the appropriate attenuation factor and the coefficient c_j of the blob and add the product onto the projection plane.

In general, we assume that the object $x(\vec{r})$ can be represented as a linear combination of shifted copies of some volume element $b(\vec{r})$:

$$x(\vec{r}) = \sum_j c_j b(\vec{r} - \vec{r}_j)$$

and we attempt to reconstruct the coefficients c_j from noisy projections of the object. The rotation-based model assumes $b(\vec{r})$ to be a rect function, whereas the blob-based model assumes $b(\vec{r})$ to have the form (A.1), where $r = ||\vec{r}||$. Therefore, an extra step of convolving the reconstructed coefficients with the blob function is needed when using the blob-based model, if the images are to be displayed or compared to images represented in voxel space.

A.3 Computation

We provide here an assessment of the number of floating-point operations (additions and multiplications) required by each implementation to compute one projection of an $n_x \times n_y \times n_z$ image.

The operations performed by the rotation-based projector are as follows.

- *Rotation:* Let n_i be the length of the 1-D interpolation filter used by the separable three-pass rotation method. To rotate one $n_x \times n_y$ slice of the input volume, each of the three passes convolves the interpolation filter with all rows or all columns of the slice. This amounts to a total of $\frac{3\pi}{4}(2n_i - 1)n_x n_y n_z$ operations to rotate all slices of the input volume. The $\frac{\pi}{4}$ factor stems from the fact that the rotation algorithm is not applied to all $n_x n_y$ voxels in a slice but only to those that lie within the circle inscribed in the slice. More computation is required when the interpolation filter must be preceded by an IIR filter, *e.g.*, for cubic interpolation [129]. For linear interpolation, where $n_i = 2$ and no IIR pre-processing is needed, the rotation requires $\frac{9\pi}{4}n_x n_y n_z \approx 7.1 n_x n_y n_z$ operations.
- *Attenuation:* Applying the pre-computed attenuation factors to all voxels of the rotated volume requires $n_x n_y n_z$ multiplications.
- *Blurring:* In general, convolving all n_y planes in the volume with 2-D PSFs of average size $n_h \times n_h$, requires $(2n_h^2 - 1)n_x n_y n_z$ operations. However, this step can be significantly speeded up for certain convenient PSF shapes. If, for example, the system PSFs are separable, each 2-D convolution is reduced to two 1-D convolutions, in which case this step can be implemented with as few as $2(2n_h - 1)n_x n_y n_z$ operations.

- *Projection:* Summing the n_y blurred planes of size $n_x \times n_z$ requires $n_x n_y n_z$ additions.

Therefore, with linear interpolation in the rotation step, the rotation-based projector performs a total of $(2n_h^2 + 8.1)n_x n_y n_z$ operations in the general case of non-separable PSFs or $(4n_h + 7.1)n_x n_y n_z$ operations in the case of separable PSFs.

On the other hand, the blob-based projector has to perform the following operations for each blob in the object support.

- *Attenuation:* 1 multiplication of the blob coefficient by a pre-computed attenuation factor.
- *Projection:* n_{bb}^2 multiplications to scale the blob footprint by the blob coefficient and n_{bb}^2 additions to add it onto the projection view, where $n_{bb} \times n_{bb}$ is the average size of the blob footprint after it has been blurred by the appropriate PSF.

On paper, the $2n_{bb}^2$ operations needed for the projection step above could be reduced to $\frac{\pi}{2}n_{bb}^2$, if we only projected the footprint samples that lay within the circular support of the footprint, rather than the entire $n_{bb} \times n_{bb}$ square. However, we do not consider this reduction further, as it is not trivial to implement. In the rotation-based projector, the $\frac{\pi}{4}$ improvement corresponds to a unique circle, the one inscribed in the image matrix, and it is simple to pre-compute its limits and reuse them throughout the computations. In the blob-based case, the $\frac{\pi}{4}$ factor would correspond to a different circle for each blob, whose limits change with blob position from angle to angle. Pre-computing and reusing those limits would require a large amount of memory and the overhead due to memory access would undermine the computational savings.

Therefore the operations add up to $2n_{bb}^2 + 1$ per blob. Assuming that the blob grid is the same as the voxel grid used by the rotation-based projector above, the blob-based projector performs a total of $(2n_{bb}^2 + 1)n_x n_y n_z$ operations for the entire projection view. If n_b is the width of the basic blob before blurring, then $n_{bb} = n_b + n_h - 1$. For any blob radius between 1 and 2 times the grid spacing, $n_b = 3$ and therefore the blob-based projector performs a total of $(2n_h^2 + 8n_h + 9)n_x n_y n_z$ operations.

Based on the calculations above, the rotation-based implementation would always be faster than the blob-based implementation with a regular cubic grid of the same size. However, when the blobs are arranged on a body-centered cubic (BCC) grid instead of the Cartesian grid, the operations required by the blob-based method can be reduced by a factor of $\sqrt{2}$ without loss in image quality, as described in [90].

If the system PSFs are separable (*e.g.*, Gaussian), the rotation-based model with separable 1-D convolutions would be faster than the blob-based model even with the BCC arrangement. In the general case of non-separable PSFs, the blob-based model with the BCC grid will be faster only when $2n_h^2 + 8n_h + 9 < \sqrt{2}(2n_h^2 + 8.1)$, which is true for $n_h \geq 10$ pixels. Whether this is a realistic PSF size will depend on the properties of the imaging system, the radius of the detector orbit and the size of the reconstruction. For example, in the case of the two SPECT reconstruction problems tackled in the simulations of Section A.5 below (technetium and iodine studies), where the diameter of the system PSF at the center of the field of view was about 31 and 52mm respectively, this would mean that the blobs would be faster for a pixel size smaller than 3.1 and 5.2mm respectively, which would be true for matrix sizes of 128×128 or larger.

Furthermore, because the blob-based implementation processes the blobs in the image grid sequentially, it can very naturally take advantage of the object support, which is readily available in SPECT from the attenuation map. If $\lambda \leq 1$ is the ratio of blobs inside the object support, the projector or back-projector only needs to process $\lambda n_x n_y n_z$ blobs (or approximately $\frac{1}{\sqrt{2}}\lambda n_x n_y n_z$ blobs when using the BCC grid). The rotation-based projector could also be modified to take advantage of the object support to some extent. However, it would be less natural than it would be in the blob case and it would significantly increase the complexity of the rotation-based implementation and its underlying data structure.

A.4 Memory Usage

Both system models pre-compute and store attenuation factors corresponding to all volume elements at all angles, to reuse them over iterations. This amounts to an array of type float and size $n_\theta n_x n_y n_z$, where n_θ is the number of projection views, and it is the dominant term in the total memory usage of both models. For a total of $128 \times 128 \times 50$ volume elements and 120 projection views, this requires roughly 400Mbytes of memory, which is reasonable for modern workstations.

A.5 Simulation Results

The CPU time required to calculate a full projection set with each system model on a 3 GHz Xeon processor with 2 GB RAM, gcc and Linux is shown in Table A.1. Times are reported for a system with a pixel size of 4mm and 2mm in the $64 \times 64 \times 25$ and $128 \times 128 \times 50$ case respectively, and PSF widths identical to the ones described for the simulations below. We assume the object support to be the circle that is inscribed in the $n_x \times n_y$ image matrix, therefore the blob-based model benefits by only processing $\frac{\pi}{4} n_x n_y n_z$ grid positions. Since our blob-based implementation uses a Cartesian grid, we also provide a prediction of the speed that could be attained with the BCC grid by dividing the practical CPU times by a factor of $\sqrt{2}$. CPU times for the rotation-based model are reported both for separable PSFs (implemented with 1-D separable convolutions as in the analysis above) and for non-separable PSFs (implemented here with 2-D FFTs). For the $64 \times 64 \times 25$ matrix the computational savings afforded by the BCC grid and the circular object support make the blob-based model faster than the rotation-based model with 2-D FFTs, whereas this is not the case for the $128 \times 128 \times 50$ matrix. The blob-based

Table A.1: Projection speeds (sec) with different system models on a 3GHz Xeon processor. *Column 3:* CPU times for blob-based model with $R = 2$ times the grid spacing, $\alpha = 10.4$, $m = 2$, and the same Cartesian grid as the rotation-based model. *Column 4:* CPU times from column 3 divided by $\sqrt{2}$ (a prediction of CPU times for blobs with the BCC grid). *Column 5:* CPU times for rotation-based model with 2-D FFTs *Column 6:* CPU times for rotation-based model with separable 1-D convolutions

$n_x \times n_y \times n_z$	n_θ	Blob	Blob BCC	Rot 2-D	Rot 1-D
$64 \times 64 \times 25$	60	3.3	2.3	3.2	1.6
$128 \times 128 \times 50$	120	113.	80.	69.	31.

model with a Cartesian grid is consistently the slowest and the rotation-based model with separable 1-D convolution is consistently the fastest.

To evaluate the two models in terms of the bias-variance trade-off, we reconstructed $64 \times 64 \times 23$ images of an anthropomorphic chest phantom (shown in Fig. A.2) from projections taken at 60 uniformly spaced angles over a $[0, 2\pi)$ range. The heart, liver, lungs, spine, and surrounding tissue in the phantom have uniform radioactivity densities with ratio 4:3:1:0:2 and uniform attenuation coefficients equal to 0.15, 0.15, 0.027, 0.191 and 0.15cm^{-1} respectively, values appropriate for the 140KeV photons emitted by Tc-99m.

The pixel size of the system was set to 4mm. The depth-dependent PSFs were assumed to be 2-D Gaussians with a FWHM increasing from 9.4 at the plane of the emission volume that is nearest to the detector to 22.6mm at the one that is furthest. The same set of PSFs were used at all projection angles, as in a circular orbit.

The noiseless 256×256 projections were generated from a high-resolution version of the phantom image using the rotation-based projector (therefore the central ray approximation used by both system models to apply attenuation is here consistent with the data) and were then downsampled to 64×64 by averaging.[1] We normalized this set of down-

---

[1] Since the 2-D projections are downsampled along two dimensions, if the reconstructed images are to be compared to the true 3-D object, a scaling factor equal to the downsampling ratio must be applied. The bias in the reconstructed images as they appear in [147] is due to lack of this scaling.

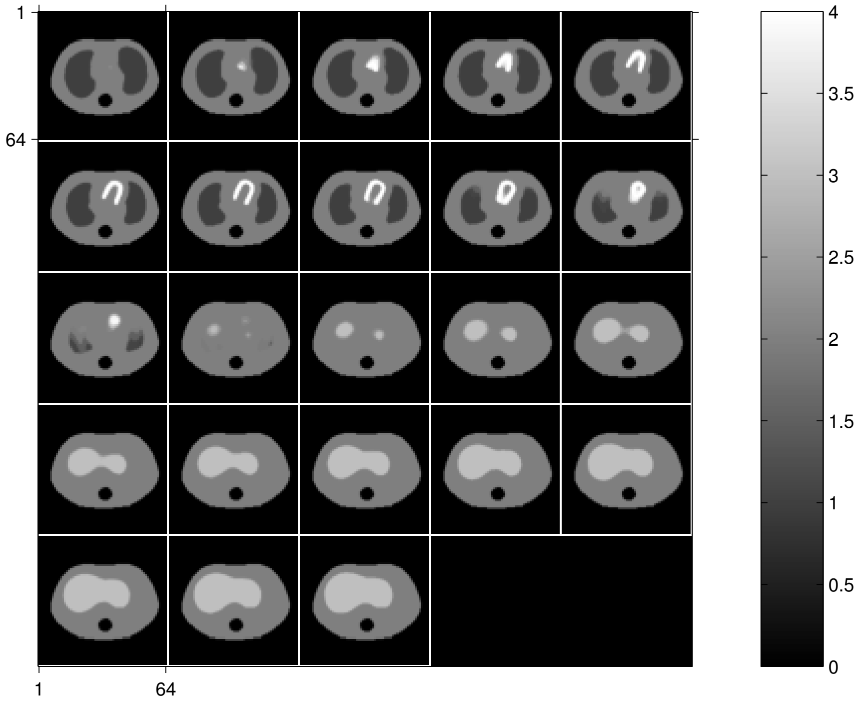

Figure A.2: Anthropomorphic chest phantom used for comparison of system models. The figure shows the phantom downsampled to $64 \times 64 \times 23$.

sampled projections to a total of 5×10^6 counts. To crudely simulate scatter, we uniformly added to the projection set a number of scattered counts equal to 10% of the total true counts. We used this new projection set to generate Poisson-distributed measurements.

We reconstructed images from 200 realizations of the noisy Poisson data using each system model and the Ordered Subsets Expectation Maximization (OSEM) algorithm with 6 subsets. The system models that we used in the OSEM reconstructions assumed the same PSFs as the system that we used to generate the noiseless projections. Although it would be of interest to compare the performance of the system models under PSF mismatch, this was beyond the scope of this paper. Our intention here was to isolate the effects of the choice of volume elements from the effects of such a mismatch. The OSEM reconstructions also assumed knowledge of the average scattered counts in the measurements [28].

We produced two different sets of bias vs. standard deviation plots by varying two different parameters that are commonly used to control the bias-variance trade-off, namely the number of iterations or the FWHM of a post-filter. In the first case we initialized the OSEM reconstructions with a uniform image and plotted the bias and standard deviation of the total uptake in the heart area per iteration, up to 50 iterations. In the second case we initialized the OSEM reconstructions with a smooth FBP image (whose non-positive values had been set to a very small positive value) and allowed the algorithm to run for 1000 iterations. We then smoothed the reconstructed images (in the blob-based case after synthesizing them from the reconstructed coefficients) using 3-D Gaussian post-filters of various FWHMs, up to 7 pixels. Similarly, we recorded the bias and standard deviation of the total uptake in the heart area for each post-filter FWHM.

Fig. A.3 shows plots of the the standard deviation vs. bias normalized as a percentage of the true total heart activity. The FBP curve in Fig. A.3 was acquired from plain FBP reconstructions with a Hanning window of variable width and Chang attenuation correction but no attempt to compensate for the depth-dependent system blur. These curves are just provided for reference, as the focus here was on comparing the blob- and rotation-based system models rather than iterative reconstruction and FBP.

Fig. A.4 shows the seventh slice from an image reconstructed with OSEM using the rotation-based model and two different blob-based models. The images were chosen at different iterations to have a similar bias level (about 16.5%) on the per iteration bias-variance plot. A profile of these slices through the heart area is also shown, superimposed on the corresponding profile of the noiseless high-resolution phantom. Similarly, Fig. A.5 shows the seventh slice from a set of the post-smoothed 1000-iteration OSEM reconstructions with the three different system models, chosen to have a similar bias level (about 19%) on the per FWHM bias-variance plot. All images in Fig. A.4 and Fig. A.5 were reconstructed from the same realization of the noisy data.

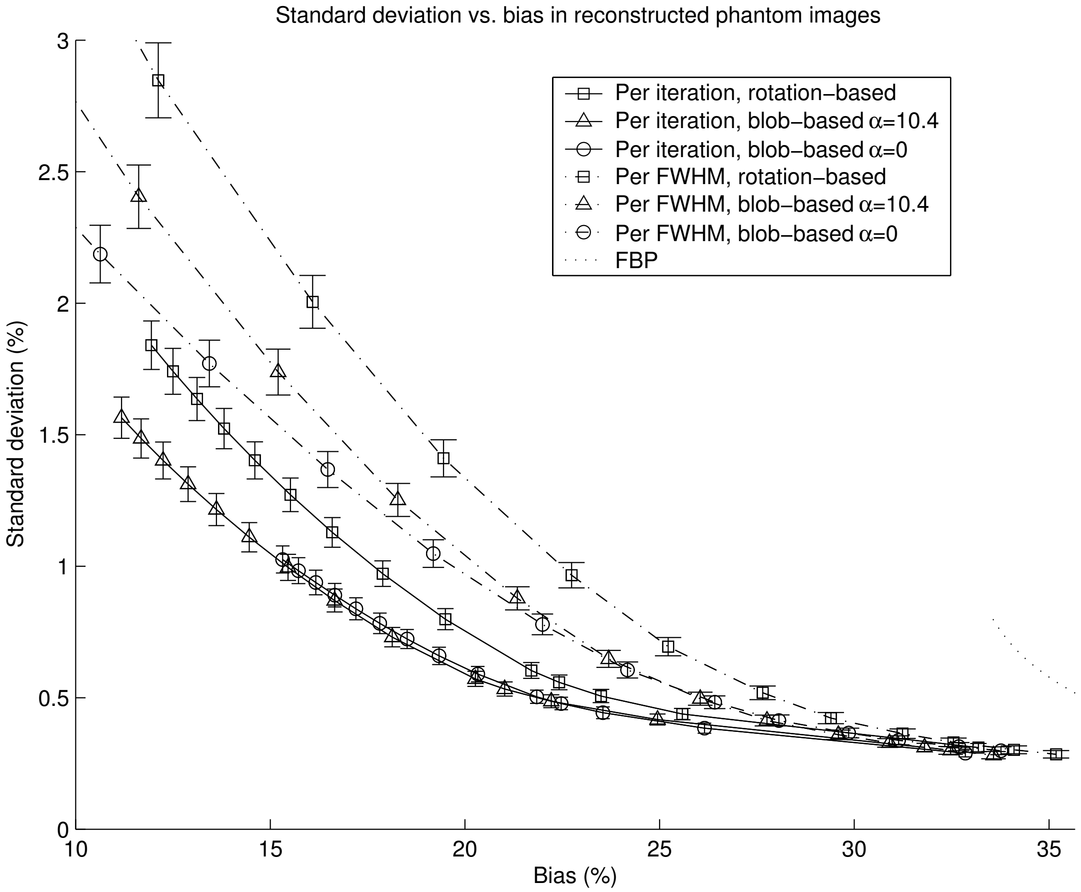

Figure A.3: Bias-variance trade-off per OSEM iteration and per post-smoothing FWHM.

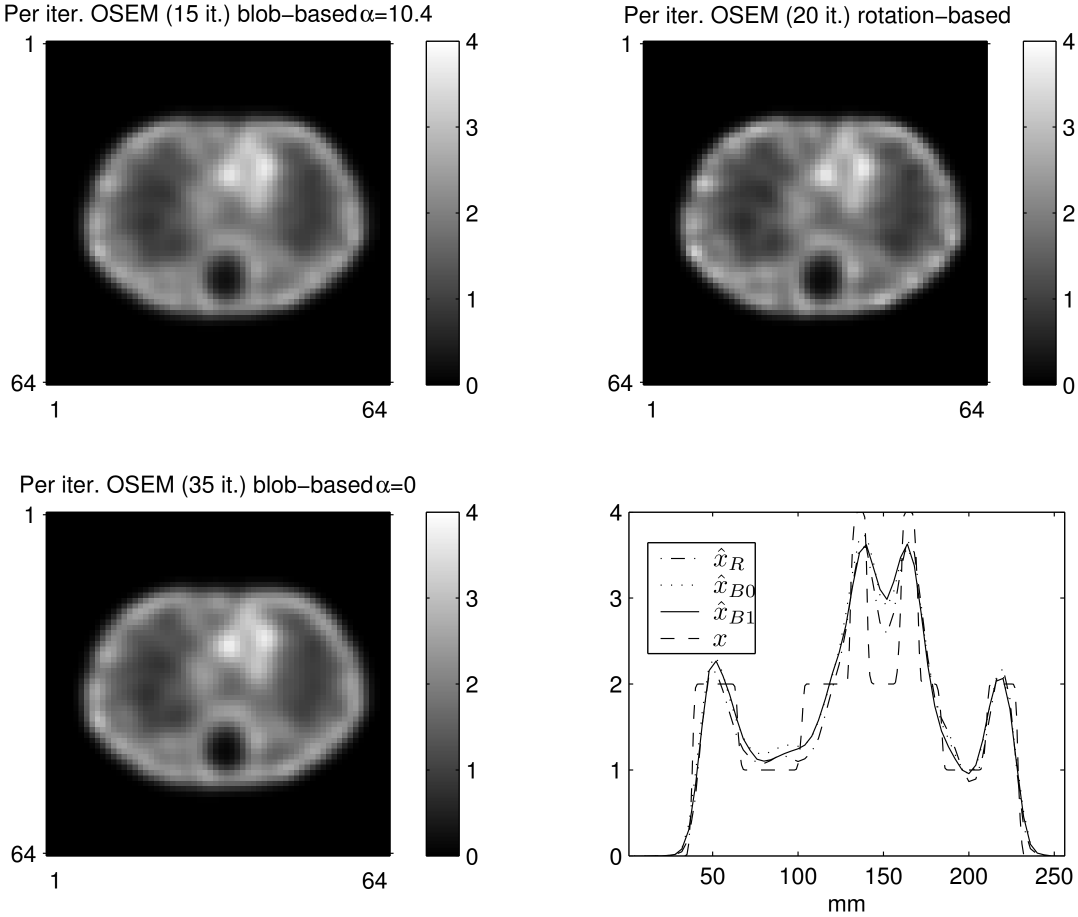

Figure A.4: Images reconstructed with OSEM and different system models. The figure shows a slice and profiles through OSEM-reconstructed images of a single realization of the noisy data. Superimposed on the profile of the true high-resolution phantom (x) are those of the images reconstructed with the rotation-based model ($\hat{x}_R$, NMSE = 3.26%), the blob-based model with $\alpha = 0$ ($\hat{x}_{B0}$, NMSE = 2.76%), and the blob-based model with $\alpha = 10.4$ ($\hat{x}_{B1}$, NMSE = 2.87%).

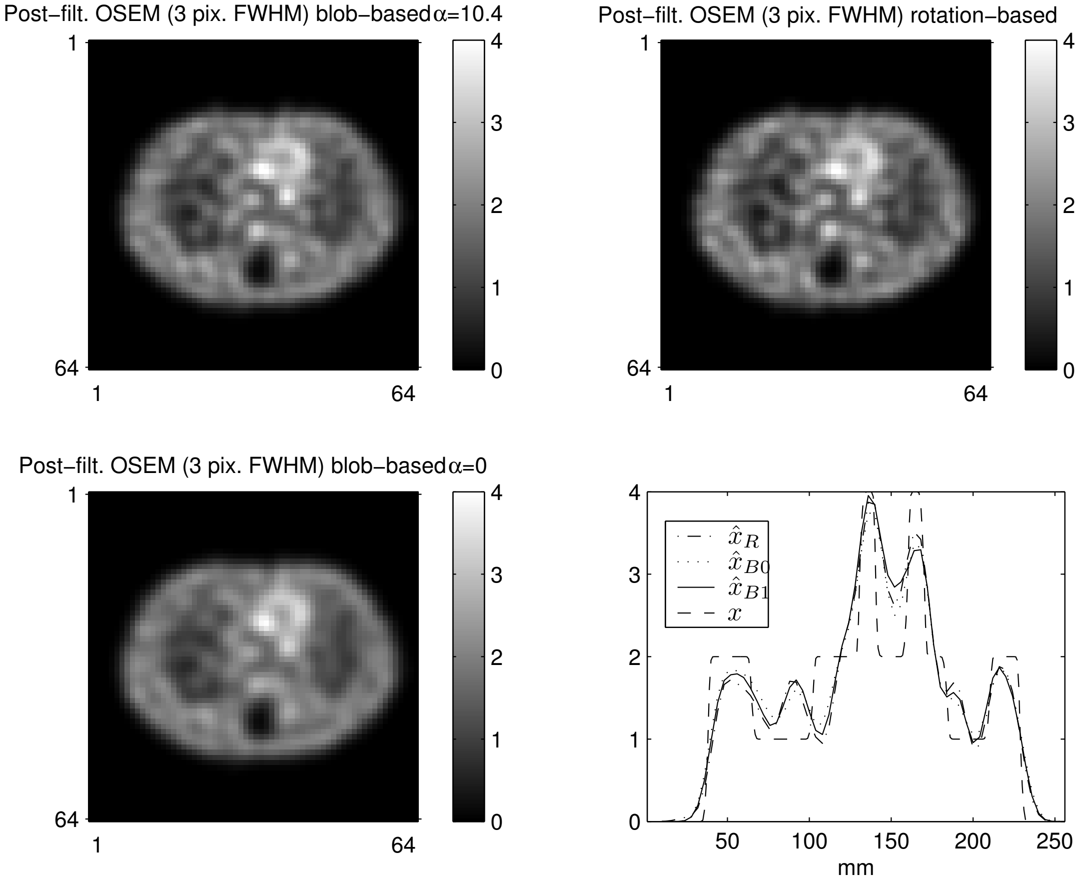

Figure A.5: Images reconstructed with post-smoothed OSEM and different system models. The figure shows a slice and profiles through over-iterated and post-smoothed OSEM-reconstructed images of a single realization of the noisy data. Superimposed on the profile of the true high-resolution phantom (x) are those of the images reconstructed with the rotation-based model ($\hat{x}_R$, NMSE = 4.12%), the blob-based model with $\alpha = 0$ ($\hat{x}_{B0}$, NMSE = 2.99%), and the blob-based model with $\alpha = 10.4$ ($\hat{x}_{B1}$, NMSE = 3.60%).

A.6 Discussion

For the blob-based reconstructions above, we tried two different blob shapes, both of them with radius $R = 2$ times the grid increment and $m = 2$ (continuous and with continuous first derivative at $r = 2$) but with different taper parameters and therefore FWHMs. The blob with $\alpha = 10.4$ has a FWHM of 1.33 times the grid increment and has been shown in [91] to possess good properties with respect to the approximation of the constant 3-D function. The blob with $\alpha = 0$ does not possess such properties and in fact is a uniform sphere with a FWHM equal to its diameter.

Although the shape of these two blobs is quite different, their bias-variance performance was not characterized by statistically significant differences at the commonly acceptable range of bias values, as shown by the error bars in Fig. A.3. In addition to these two choices of blob shape, we also considered the other choices listed in Table I of [91] and they resulted in similar performance, although these results are not shown here to avoid excessive clutter in the plots. No further effort was made here to optimize the blob parameters for SPECT reconstruction. It appears that the wide PSFs of the SPECT detector dominated the bias-variance properties of the reconstruction and varying the blob shape did not make a significant difference in the simulations presented here.

The plots do show, however, that the bias-variance performance of the blob-based projectors was somewhat better than that of the rotation-based one, similarly to what investigators have reported for PET [89, 91]. It is possible that spherically symmetric volume elements were a better match for the contours of the shapes in the high-resolution phantom used in our simulations. However, as shown by the bias-variance plots and confirmed by the reconstructed images, the improvement is modest for SPECT.

The spacing of the pre-computed blob footprint samples that we used in all the simulations above was equal to $\frac{1}{10}$ of the system pixel size. Decreasing this sample spacing did not affect the quality of the reconstructed images.

Similar results (not shown here) were obtained when simulating an I-131 study with the same phantom and system parameters appropriate for an ultra-high energy collimator (PSF FWHM from 3 to 43mm, pixel size 7.2mm).

In general, our simulations indicated only a slightly superior performance of the blob-based implementations of the 3-D SPECT system model in terms of the bias-variance trade-off. On the other hand, the differences in terms of computational load were significant. When the system PSFs can be considered separable, such as the Gaussians used above, the rotation-based model with 1-D separable convolutions will always be faster. If we wanted to model collimator properties in greater detail with non-symmetric PSFs, we would resort to 2-D convolutions, and in that case the blob-based model with an efficient grid could be faster for large PSF sizes. However, FFTs can be used instead of 2-D convolutions to speed up the rotation-based model with those larger PSFs.

The fact that use of the object support can be very naturally built into the blob-based system model makes it very suitable for cases where we only need to reconstruct a few or even a single volume element from the object, such as in a recently proposed method for fast estimation of resolution and covariance in SPECT [125] and in algorithms that perform element-by-element processing, such as the Space-Alternating Generalized EM (SAGE) algorithm [39].

Finally, our simulations did not include realistic Compton scatter effects. We speculate that since Compton scatter is relatively spatially smooth, and since scatter computations are usually performed with relatively coarse sampling [14, 103, 139], including a realistic scatter model would not significantly alter the outcome of our comparisons in favor of any choice of basis function.

APPENDIX B

Locally Shift-Invariant Approximation of $\boldsymbol{\mathcal{A}}'\boldsymbol{\Pi}^{-1}\boldsymbol{\mathcal{A}}$

We treat our model of the tomographic imaging system as a linear continuous-to-discrete operator $\boldsymbol{\mathcal{A}}$ mapping the object space into the measurement space. We assume that the object can be written as a function $f : \mathbb{R}^2 \to \mathbb{R}$, and the measurement can be written as a vector $\boldsymbol{p} \in \mathbb{R}^{n_\mathrm{d}}$. We focus on the 2-D case where $n_\mathrm{d} = n_\mathrm{b} n_\varphi$, but the analysis below can be easily generalized to higher-dimensional models. The measurement $\boldsymbol{p}$ is acquired by the lexicographic ordering of elements $p_{ij},\ i = 1, \ldots, n_\mathrm{b},\ j = 1, \ldots, n_\varphi$.

In its most general form, the system operator $\boldsymbol{\mathcal{A}}$ is defined by a superposition integral, possibly accounting for position-dependent characteristics of the imaging system:

$$\boldsymbol{p} = \boldsymbol{\mathcal{A}} \boldsymbol{f} \Leftrightarrow p_{ij} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} a(r_i, \varphi_j; x, y)\, f(x, y)\, \mathrm{d}x\, \mathrm{d}y, \quad i = 1, \ldots, n_\mathrm{b}, \quad j = 1, \ldots, n_\varphi,$$

where r_i and φ_j are the radial position and angle corresponding to the (i, j) measurement bin respectively. The adjoint operator $\boldsymbol{\mathcal{A}}'$ is then defined as follows:

$$\boldsymbol{f} = \boldsymbol{\mathcal{A}}' \boldsymbol{p} \Leftrightarrow f(x, y) = \sum_{j=1}^{n_\varphi} \sum_{i=1}^{n_\mathrm{b}} a^*(r_i, \varphi_j; x, y) p_{ij}, \quad x, y \in \mathbb{R}.$$

Consider a pixel of interest (x_0, y_0) and let $\boldsymbol{e}_0$ denote a Dirac impulse centered at this point. We define the following impulse responses locally at (x_0, y_0):

$$\begin{aligned} \boldsymbol{h}_0 &\triangleq \boldsymbol{\mathcal{A}}' \boldsymbol{\Pi}^{-1} \boldsymbol{\mathcal{A}} \boldsymbol{e}_0 \\ \boldsymbol{a}_0 &\triangleq \boldsymbol{\mathcal{A}} \boldsymbol{e}_0, \end{aligned}$$

where $\boldsymbol{h}_0$ corresponds to a function $h_0 : \mathbb{R}^2 \to \mathbb{R}$, and $\boldsymbol{a}_0 \in \mathbb{R}^{n_\mathrm{d}}$. We write $h(x, y; x_0, y_0) \triangleq h_0(x, y)$. Then the value of h_0 at some point (x, y) is

$$h(x, y; x_0, y_0) = \sum_{j=1}^{n_\varphi} \sum_{i=1}^{n_\mathrm{b}} \frac{a^*(r_i, \varphi_j; x, y) a(r_i, \varphi_j; x_0, y_0)}{\sigma_{ij}^2}, \tag{B.1}$$

where the variances σ_{ij}^2, $i = 1, \ldots, n_\mathrm{b}$, $j = 1, \ldots, n_\varphi$, when lexicographically ordered, yield the diagonal of $\boldsymbol{\Pi}$.

We now define the following angle-dependent *local certainties*, *cf.* equation (19) in [124]:

$$\kappa_{\varphi_j}(x, y) \triangleq \sqrt{\frac{\sum_{i=1}^{n_\mathrm{b}} |a(r_i, \varphi_j; x, y)|^2 \frac{1}{\sigma_{ij}^2}}{\sum_{i=1}^{n_\mathrm{b}} |a(r_i, \varphi_j; x, y)|^2}}, \quad j = 1, \ldots, n_\varphi. \tag{B.2}$$

This reduces to $\kappa_{\varphi_j}(x, y) = \frac{1}{\sigma_j}$ when the measurement variance σ_j^2 depends only on the angle φ_j and not on the radial position r_i. The local certainties in (B.2), used in [36, 149], are a refinement of the single angle-independent local certainty,

$$\kappa_0(x, y) \triangleq \sqrt{\frac{\sum_{j=1}^{n_\varphi} \sum_{i=1}^{n_\mathrm{b}} |a(r_i, \varphi_j; x, y)|^2 \frac{1}{\sigma_{ij}^2}}{\sum_{j=1}^{n_\varphi} \sum_{i=1}^{n_\mathrm{b}} |a(r_i, \varphi_j; x, y)|^2}}, \tag{B.3}$$

that other researchers have used in their locally shift-invariant approximations for lesion detectability. Using (B.2) instead of (B.3) in the analysis that follows leads to more accurate approximations by better capturing aspects of the system response that are angle-dependent, *e.g.*, non-uniform attenuation.

Along the "diagonal" of $\boldsymbol{\mathcal{A}}'\boldsymbol{\Pi}^{-1}\boldsymbol{\mathcal{A}}$ we have:

$$h(x_0, y_0; x_0, y_0) = \sum_{j=1}^{n_\varphi} \sum_{i=1}^{n_\mathrm{b}} \kappa_{\varphi_j}^2(x_0, y_0) |a(r_i, \varphi_j; x_0, y_0)|^2.$$

We now adopt the following approximation (*cf.* [40]), which is exact along the diagonal:

$$h(x, y; x_0, y_0) \approx \sum_{j=1}^{n_\varphi} \kappa_{\varphi_j}(x, y) \kappa_{\varphi_j}(x_0, y_0) \sum_{i=1}^{n_\mathrm{b}} a^*(r_i, \varphi_j; x, y) a(r_i, \varphi_j; x_0, y_0).$$

Since the local impulse response $\boldsymbol{h}_0$ takes its largest values around the pixel (x_0, y_0), we are mostly interested in (x, y) in the neighborhood of (x_0, y_0). In that case, using the fact that typically the local certainties vary slowly, we can further approximate the above as follows:

$$h(x,y;x_0,y_0) \approx \sum_{j=1}^{n_\varphi} \kappa^2_{\varphi_j}(x_0,y_0) \sum_{i=1}^{n_\mathrm{b}} a^*(r_i,\varphi_j;x,y)a(r_i,\varphi_j;x_0,y_0).$$

To further simplify, define the following "local detector response" function for the object point (x_0, y_0):

$$b_0(\tau,\varphi) \triangleq a(\tau + [x_0 \cos\varphi + y_0 \sin\varphi], \varphi; x_0, y_0),$$

so that

$$a(r,\varphi;x_0,y_0) = b_0(r - [x_0 \cos\varphi + y_0 \sin\varphi], \varphi)\,.$$

Now assume that the detector response functions $a(\cdot)$ are *locally shift-invariant* in the following sense

(B.4) $$a(r,\varphi;x,y) \approx b_0(r - [x \cos\varphi + y \sin\varphi], \varphi),$$

for $(x, y) \approx (x_0, y_0)$. This assumption leads to the following approximation for the inner sum:

$$\begin{aligned}
&\sum_{i=1}^{n_\mathrm{b}} a(r_i,\varphi_j;x,y)^* a(r_i,\varphi_j;x_0,y_0) \\
&\quad\approx \sum_{i=1}^{n_\mathrm{b}} b_0^*(r_i - [x\cos\varphi_j + y\sin\varphi_j], \varphi_j)\, b_0(r_i - [x_0\cos\varphi_j + y_0\sin\varphi_j], \varphi_j) \\
&\quad= \sum_{i=1}^{n_\mathrm{b}} b_0^*(\tau_i, \varphi_j)\, b_0(\tau_i + (x - x_0)\cos\varphi_j + (y - y_0)\sin\varphi_j, \varphi_j) \\
&\quad= (b_0 \star b_0)[(x - x_0)\cos\varphi_j + (y - y_0)\sin\varphi_j],
\end{aligned}$$

where $\tau_i = r_i - [x\cos\varphi_j + y\sin\varphi_j]$ and $\star$ denotes discrete-space 1-D autocorrelation (with respect to τ_i).

This leads to our final approximation for the local impulse response:

$$\begin{aligned} &h(x, y; x_0, y_0) \\ &\approx \sum_{j=1}^{n_\varphi} \kappa^2_{\varphi_j}(x_0, y_0)\,(b_0 \star b_0)[(x - x_0)\cos\varphi_j + (y - y_0)\sin\varphi_j] \\ &= \int_0^\pi \sum_{j=1}^{n_\varphi} \delta(\varphi - \varphi_j)\kappa^2_\varphi(x_0, y_0)\,(b_0 \star b_0)[(x - x_0)\cos\varphi + (y - y_0)\sin\varphi]\,\mathrm{d}\varphi\,. \end{aligned} \tag{B.5}$$

This final *shift-invariant* approximation agrees with the exact expression (B.1) “along the diagonal” of $\boldsymbol{\mathcal{A}}'\boldsymbol{\Pi}^{-1}\boldsymbol{\mathcal{A}}$, *i.e.*, when $(x, y) = (x_0, y_0)$. And it agrees everywhere if σ_{ij} is independent of i and if the detector response functions $a(\cdot)$ are shift-invariant in the sense that (B.4) holds exactly.

If we are to apply the operator $\boldsymbol{\mathcal{A}}'\boldsymbol{\Pi}^{-1}\boldsymbol{\mathcal{A}}$ to signals highly localized around (x_0, y_0), then it is reasonable to approximate the operator as a convolution by the local impulse response $\boldsymbol{h}_0$, to which we can apply Fourier analysis. So we approximate $\boldsymbol{\mathcal{A}}'\boldsymbol{\Pi}^{-1}\boldsymbol{\mathcal{A}}$ locally around (x_0, y_0) as

$$\boldsymbol{\mathcal{A}}'\boldsymbol{\Pi}^{-1}\boldsymbol{\mathcal{A}} \approx \boldsymbol{\mathcal{U}}^{-1}\boldsymbol{\mathcal{D}}\{\lambda(\rho, \Phi)\}\boldsymbol{\mathcal{U}},$$

where $\boldsymbol{\mathcal{D}}\{\cdot\}$ denotes a “diagonal” operator and $\boldsymbol{\mathcal{U}}$ a 2-D Fourier operator. The local frequency response is given by

$$\lambda(\rho, \Phi) \triangleq \sum_{j=1}^{n_\varphi} \kappa^2_{\varphi_j}(x_0, y_0)\frac{|B_0(\rho, \varphi_j)|^2}{|\rho|}\delta(\Phi - \varphi_j), \tag{B.6}$$

and $B_0(\cdot, \varphi)$ denotes the 1-D FT of $b_0(\cdot, \varphi)$.

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] C. K. Abbey and H. H. Barrett. Human- and model-observer performance in ramp-spectrum noise: effects of regularization and object variability. *J. Opt. Soc. Am. A*, 18(3):473–88, March 2001.

[2] C. K. Abbey, H. H. Barrett, and M. P. Eckstein. Practical issues and methodology in assessment of image quality using model observers. In *Proc. SPIE 3032, Med. Im. 1997: Phys. of Med. Im.*, pages 182–94, 1997.

[3] C. K. Abbey and M. P. Eckstein. Maximum-likelihood and maximum-a-posteriori estimates of human-observer templates. In *Proc. SPIE 4324, Medical Imaging 2001: Image Perc. and Perf.*, pages 114–22, 2001.

[4] C. K. Abbey and M. P. Eckstein. Optimal shifted estimates of human-observer templates in two-alternative forced-choice experiments. *IEEE Trans. Med. Imag.*, 21(5):429–40, May 2002.

[5] R. J. Adler. *The geometry of random fields*. Wiley, New York, 1981.

[6] R. J. Adler. On excursion sets, tube formulas and maxima of random fields. *Ann. Appl. Probab.*, 10(1):1–74, February 2000.

[7] H. H. Barrett. Objective assessment of image quality: effects of quantum noise and object variability. *J. Opt. Soc. Am. A*, 7(7):1266–1278, July 1990.

[8] H. H. Barrett, C. K. Abbey, and E. Clarkson. Objective assessment of image quality III: ROC metrics, ideal observers and likelihood-generating functions. *J. Opt. Soc. Am. A*, 15(6):1520–35, June 1998.

[9] H. H. Barrett, J. L. Denny, R. F. Wagner, and K. J. Myers. Objective assessment of image quality. II. Fisher information, Fourier crosstalk, and figures of merit for task performance. *J. Opt. Soc. Am. A*, 12(5):834–52, May 1995.

[10] H. H. Barrett, T. Gooley, K. Girodias, J. Rolland, T. White, and J. Yao. Linear discriminants and image quality. In A. C. F. Colchester and D. J. Hawkes, editors, *Information Processing in Medical Im.*, pages 458–73. Springer-Verlag, Berlin, 1991.

[11] H. H. Barrett and K. J. Myers. *Foundations of image science*. Wiley, New York, 2003.

[12] H. H. Barrett, D. W. Wilson, and B. M. W. Tsui. Noise properties of the EM algorithm: I. Theory. *Phys. Med. Biol.*, 39(5):833–46, May 1994.

[13] H. H. Barrett, J. Yao, J. P. Rolland, and K. J. Myers. Model observers for assessment of image quality. *Proc. Natl. Acad. Sci.*, 90:9758–65, November 1993.

[14] F. J. Beekman, H. W. A. M. de Jong, and S. van Geloven. Efficient fully 3D iterative SPECT reconstruction with Monte Carlo based scatter compensation. *IEEE Trans. Med. Imag.*, 21(8):867–77, August 2002.

[15] F. J. Beekman, E. T. P. Slijpen, H. W. A. M. . Jong, and M. A. Viergever. Estimation of the depth-dependent component of the point spread function of SPECT. *Med. Phys.*, 26(11):2311–22, November 1999.

[16] F. J. Beekman, E. T. P. Slijpen, and W. J. Niessen. Selection of task-dependent diffusion filters for the post-processing of SPECT images. *Phys. Med. Biol.*, 43(6):1713–30, June 1998.

[17] M. Belge, M. E. Kilmer, and E. L. Miller. Efficient determination of multiple regularization parameters in a generalized L-curve framework. *Inverse Prob.*, 18(4):1161–83, August 2002.

[18] F. O. Bochud, C. K. Abbey, and M. Eckstein. Statistical texture synthesis of mammographic images with clustered lumpy backgrounds. *Optics Express*, 4(1):33–42, January 1999.

[19] P. Bonetto, J. Qi, and R. M. Leahy. Covariance approximation for fast and accurate computation of channelized Hotelling observer statistics. *IEEE Trans. Nuc. Sci.*, 47(4):1567–72, August 2000.

[20] A. B. Brill and R. N. Beck. Evolution of clinical emission tomography. In M N Wernick and J N Aarsvold, editors, *Emission tomography*, chapter 3. Academic Press, San Diego, 2004.

[21] V. A. Brookerman and T. J. Bauer. Collimator performance for scintillation camera systems. *J. Nuc. Med.*, 14(1):21–5, January 1973.

[22] A. E. Burgess and B. Colborne. Visual signal detection. IV. Observer inconsistency. *J. Opt. Soc. Am. A*, 5(4):617–, April 1988.

[23] A. E. Burgess, R. F. Wagner, R. J. Jennings, and H. B. Barlow. Efficiency of human visual signal discrimination. *Science*, 214(4516):93–94, October 1981.

[24] Y. Censor. Row-action methods for huge and sparse systems and their applications. *SIAM Review*, 23(4):444–66, October 1981.

[25] L. T. Chang. A method for attenuation correction in radionuclide computed tomography. *IEEE Trans. Nuc. Sci.*, 25(1):638–643, February 1978.

[26] D. A. Chesler and S. J. Riederer. Ripple suppression during reconstruction in transverse tomography. *Phys. Med. Biol.*, 20(4):632–6, July 1975.

[27] M. S. Chesters. Human visual perception and ROC methodology in medical imaging. *Phys. Med. Biol.*, 37(7):1433–84, July 1992.

[28] M. E. Daube-Witherspoon, R. E. Carson, Y. Yan, and T. K. Yap. Scatter correction in maximum likelihood reconstruction of PET data. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pages 945–7, 1992.

[29] J. D. Daugman. Two-dimensional spectral analysis of cortical receptive field profiles. *Vision Research*, 20(10):847–56, 1980.

[30] E. V. R. Di Bella, A. B. Barclay, R. L. Eisner, and R. W. Schafer. Comparison of rotation-based methods for iterative reconstruction algorithms. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pages 1146–50, 1995.

[31] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern classification*. Wiley, New York, 2001.

[32] M. P. Eckstein, C. K. Abbey, and F. O. Bochud. A practical guide to model observers for visual detection in synthetic and natural noisy images. In J. Beutel, H. L. Kundel, and R. L. Van Metter, editors, *Handbook of Medical Imaging, Volume 1. Physics and Psychophysics*, chapter 10. SPIE, Bellingham, 2000.

[33] J. A. Fessler. Penalized weighted least-squares image reconstruction for positron emission tomography. *IEEE Trans. Med. Imag.*, 13(2):290–300, June 1994.

[34] J. A. Fessler. ASPIRE 3.0 user's guide: A sparse iterative reconstruction library. Technical Report 293, Comm. and Sign. Proc. Lab., Dept. of EECS, Univ. of Michigan, Ann Arbor, MI, 48109-2122, July 1995. Available from `http://www.eecs.umich.edu/~fessler`.

[35] J. A. Fessler. Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): Applications to tomography. *IEEE Trans. Im. Proc.*, 5(3):493–506, March 1996.

[36] J. A. Fessler. Analytical approach to regularization design for isotropic spatial resolution. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 3, pages 2022–6, 2003.

[37] J. A. Fessler. *Image reconstruction: Algorithms and analysis*. 2005. To be published.

[38] J. A. Fessler and S. D. Booth. Conjugate-gradient preconditioning methods for shift-variant PET image reconstruction. *IEEE Trans. Im. Proc.*, 8(5):688–99, May 1999.

[39] J. A. Fessler and A. O. Hero. Penalized maximum-likelihood image reconstruction using space-alternating generalized EM algorithms. *IEEE Trans. Im. Proc.*, 4(10):1417–29, October 1995.

[40] J. A. Fessler and W. L. Rogers. Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs. *IEEE Trans. Im. Proc.*, 5(9):1346–58, September 1996.

[41] J. A. Fessler and A. Yendiki. Channelized Hotelling observer performance for penalized-likelihood image reconstruction. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pages 1040–4, 2002.

[42] R. D. Fiete, H. H. Barrett, W. E. Smith, and K. J. Myers. The Hotelling trace criterion and its correlation with human observer performance. *J. Opt. Soc. Am. A*, 4:945–953, 1987.

[43] V. Fragkou. *A generalization of fundamental concepts of metric spaces*. PhD thesis, Aristotle Univ. of Thessaloniki, Thessaloniki, Greece, 1969.

[44] K. Fukunaga. *Introduction to statistical pattern recognition*. Academic, Boston, 1990.

[45] B. D. Gallas and H. H. Barrett. Validating the use of channels to estimate the ideal linear observer,. *J. Opt. Soc. Am. A*, 20(9):1725–38, September 2003.

[46] S. Geman and D. E. McClure. Bayesian image analysis: an application to single photon emission tomography. In *Proc. of Stat. Comp. Sect. of Amer. Stat. Assoc.*, pages 12–8, 1985.

[47] H. C. Gifford, P. E. Kinahan, C. Lartizien, and M. A. King. Evaluation of multiclass model observers in PET LROC studies. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 7, pages 4068–71, 2004.

[48] H. C. Gifford, M. A. King, D. J. de Vries, and E. J. Soares. Channelized Hotelling and human observer correlation for lesion detection in hepatic SPECT imaging. *J. Nuc. Med.*, 41(3):514–21, March 2000.

[49] H. C. Gifford, M. A. King, P. H. Pretorius, and R. G. Wells. A comparison of human and model observers in multislice LROC studies. *IEEE Trans. Med. Imag.*, 24(2):160–9, February 2005.

[50] H. C. Gifford, M. A. King, R. G. Wells, W. G. Hawkins, M. V. Narayanan, and P. H. Pretorius. LROC analysis of detector-response compensation in SPECT. *IEEE Trans. Med. Imag.*, 19(5):463–73, May 2000.

[51] H. C. Gifford, P. H. Pretorius, and M. A. King. Comparison of human- and model-observer LROC studies. In *Proc. SPIE 5034: Medical Imaging 2003: Image Perc., Obs. Perf. and Tech. Assess.*, pages 112–22, 2003.

[52] H. C. Gifford, R. G. Wells, and M. A. King. A comparison of human observer LROC and numerical observer ROC for tumor detection in SPECT images. *IEEE Trans. Nuc. Sci.*, 46(4):1032–7, August 1999.

[53] F. A. Graybill. *Matrices with applications in statistics*. Wadsworth, Bellmont, CA, 2 edition, 1983.

[54] R. Gupta. *Quantization strategies for low-power communications*. PhD thesis, Univ. of Michigan, Ann Arbor, MI, 2001.

[55] P. C. Hansen. Analysis of discrete ill-posed problems by means of the L-curve. *SIAM Review*, 34(4):561–580, December 1992.

[56] T. Hebert and R. Leahy. A generalized EM algorithm for 3-D Bayesian reconstruction from Poisson data using Gibbs priors. *IEEE Trans. Med. Imag.*, 8(2):194–202, June 1989.

[57] T. J. Hebert. Statistical stopping criteria for iterative maximum likelihood reconstruction of emission images. *Phys. Med. Biol.*, 35(9):1221–32, September 1990.

[58] H. V. Henderson and S. R. Searle. On deriving the inverse of a sum of matrices. *SIAM Review*, 23(1):53–60, January 1981.

[59] A. O. Hero, J. A. Fessler, and M. Usman. Exploring estimator bias-variance tradeoffs using the uniform CR bound. *IEEE Trans. Sig. Proc.*, 44(8):2026–41, August 1996.

[60] D. M. Higdon, J. E. Bowsher, V. E. Johnson, T. G. Turkington, D. R. Gilland, and R. J. Jaszczak. Fully Bayesian estimation of Gibbs hyperparameters for emission computed tomography data. *IEEE Trans. Med. Imag.*, 16(5):516–26, October 1997.

[61] H. Hotelling. The generalization of Student's ratio. *Ann. Math. Stat.*, 2(3):360–78, 1931.

[62] P. J. Huber. *Robust statistics*. Wiley, New York, 1981.

[63] H. M. Hudson and R. S. Larkin. Accelerated image reconstruction using ordered subsets of projection data. *IEEE Trans. Med. Imag.*, 13(4):601–9, December 1994.

[64] D. A. Hutton and R. N. Strickland. Channelized detection filters for detecting tumors in nuclear medical images. In *Proc. SPIE 3034: Medical Imaging 1997: Image Proc.*, pages 457–66, 1997.

[65] M. Jacobson, R. Levkovitz, A. Ben-Tal, K. Thielemans, T. Spinks, D. B. E. Pagani, V. Bettinardi, M. C. Gilardi, A. Zverovich, and G. Mitra. Enhanced 3D PET OSEM reconstruction using inter-update Metz filtering. *Phys. Med. Biol.*, 45(8):2417–39, August 2000.

[66] M. C. Jones, J. S. Marron, and S. J. Sheather. A brief survey of bandwidth selection for density estimation. *J. Am. Stat. Ass.*, 91(433):401–7, March 1996.

[67] C-M. Kao and X. Pan. Non-iterative methods incorporating a priori source distribution and data information for suppression of image noise and artefacts in 3D SPECT. *Phys. Med. Biol.*, 45(10):2801–19, October 2000.

[68] P. K. Khurd and G. Gindi. Decision strategies maximizing the area under the LROC curve. In *spie-5749*, pages 150–61, 2005.

[69] P. K. Khurd and G. R. Gindi. Rapid computation of LROC figures of merit using numerical observers (for SPECT/PET reconstruction). In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 4, pages 2516–20, 2003.

[70] P. K. Khurd and G. R. Gindi. LROC model observers for emission tomographic reconstruction. In *Proc. SPIE 5372, Medical Imaging 2004: Image Perception, Observer Performance, and Technology Assessment*, pages 509–20, 2004.

[71] M. A. King, S. J. Glick, P. H. Pretorius, R. G. Wells, H. C. Gifford, M. V. Narayanan, and T. Farncombe. Attenuation, scatter, and spatial resolution compensation in SPECT. In M N Wernick and J N Aarsvold, editors, *Emission tomography*. Academic Press, San Diego, 2004. Chapter 22.

[72] K. F. Koral, Q. Lin, A. Yendiki, A. Akhtar, Y. K. Dewaraja, and J. A. Fessler. Application of unregularized OSEM with 3D detector response to I-131 SPECT with ultra-high- and high-energy collimation. *J. Nuc. Med. (Abs. Book)*, 43(5):215, May 2002.

[73] K. F. Koral, A. Yendiki, Q. Lin, and Y. K. Dewaraja. Comparison of 3-D OSEM versus 1-D SAGE for focal total-activity quantification in I-131 SPECT with HE collimation. *IEEE Trans. Nuc. Sci.*, 52(1):154–8, February 2005.

[74] K. F. Koral, A. Yendiki, Q. Lin, Y. K. Dewaraja, and J. A. Fessler. Update on HE vs UHE collimation for focal total-activity quantification in I-131 SPECT using 3D OSEM. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 4, pages 2910–13, 2003.

[75] K. F. Koral, A. Yendiki, Q. Lin, Y. K. Dewaraja, and J. A. Fessler. Determining total I-131 activity within a VoI using SPECT, a UHE collimator, OSEM, and a constant conversion factor. *IEEE Trans. Nuc. Sci.*, 51(3):611–8, June 2004.

[76] H. Kunsch. Robust priors for smoothing and image restoration. *Ann. Inst. Stat. Math.*, 46(1):1–19, 1994.

[77] K. J. LaCroix and B. M. W. Tsui. An evaluation of the effect of nonuniform attenuation compensation on defect detection for Tc-99m myocardial SPECT images. *J. Nuc. Med. (Abs. Book)*, 38(5):19, May 1997.

[78] K. Lange. Convergence of EM image reconstruction algorithms with Gibbs smoothing. *IEEE Trans. Med. Imag.*, 9(4):439–46, December 1990. Corrections, T-MI, 10:2(288), June 1991.

[79] R. Lecomte, D. Schmitt, and G. Lamoureux. Geometry study of a high resolution PET detection system using small detectors. *IEEE Trans. Nuc. Sci.*, 31(1):556–61, February 1984.

[80] S-J. Lee, A. Rangarajan, and G. Gindi. Bayesian image reconstruction in SPECT using higher order mechanical models as priors. *IEEE Trans. Med. Imag.*, 14(4):669–80, December 1995.

[81] T. Lewellen and J. Carp. PET systems. In M N Wernick and J N Aarsvold, editors, *Emission tomography*. Academic Press, San Diego, 2004. Chapter 10.

[82] R. M. Lewitt. Multidimensional digital image representations using generalized Kaiser-Bessel window functions. *J. Opt. Soc. Am. A*, 7(10):1834–46, October 1990.

[83] R. M. Lewitt. Alternatives to voxels for image representation in iterative reconstruction algorithms. *Phys. Med. Biol.*, 37(3):705–16, March 1992.

[84] R. M. Lewitt and S. Matej. Overview of methods for image reconstruction from projections in emission computed tomography. *Proc. IEEE*, 91(10):1588–611, October 2003.

[85] L. N. D. Loo, K. Doi, and C. E. Metz. A comparison of physical image comparison indices and observer performance in the radiographic detection of nylon beads. *Phys. Med. Biol.*, 29(7):837–57, July 1984.

[86] L. B. Lusted. Signal detectability and medical decision-making. *Science*, 171(3977):1217–19, March 1971.

[87] A. Macovski. *Medical imaging systems*. Prentice-Hall, New Jersey, 1983.

[88] S. Marčelja. Mathematical description of the responses of simple cortical cells. *J. Opt. Soc. Am.*, 70(11):1297–300, 1980.

[89] S. Matej, G. T. Herman, T. K. Narayan, S. S. Furuie, R. M. Lewitt, and P. E. Kinahan. Evaluation of task-oriented performance of several fully 3D PET reconstruction algorithms. *Phys. Med. Biol.*, 39(3):355–67, March 1994.

[90] S. Matej and R. M. Lewitt. Efficient 3D grids for image reconstruction using spherically-symmetric volume elements. *IEEE Trans. Nuc. Sci.*, 42(4):1361–70, August 1995.

[91] S. Matej and R. M. Lewitt. Practical considerations for 3-D image reconstruction using spherically symmetric volume elements. *IEEE Trans. Med. Imag.*, 15(1):68–78, February 1996.

[92] C. E. Metz. Basic principles of ROC analysis. *Seminars in Nuclear Medicine*, 8(4):283–98, October 1978.

[93] C. E. Metz, F. B. Atkins, and R. N. Beck. The geometric transfer function component for scintillation camera collimators with straight parallel holes. *Phys. Med. Biol.*, 25(6):242–50, November 1980.

[94] K. Millett. *Sexual politics*. Doubleday, New York, 1970.

[95] S. Mustafovic and K. Thielemans. Object dependency of resolution in reconstruction algorithms with inter-iteration filtering applied to PET data. *IEEE Trans. Med. Imag.*, 23(4):433–46, April 2004.

[96] K. J. Myers and H. H. Barrett. Addition of a channel mechanism to the ideal-observer model. *J. Opt. Soc. Am. A*, 4(12):2447–57, December 1987.

[97] K. J. Myers and K. M. Hanson. Comparison of the algebraic reconstruction technique with the maximum entropy reconstruction technique for a variety of detection tasks. In *Proc. SPIE 1231, Med. Im. IV: Image Formation*, pages 176–87, 1990.

[98] M. V. Narayanan, M. A. King, P. H. Pretorius, S. T. Dahlberg, F. Spencer, E. Simon, E. Ewald, E. Healy, K. MacNaught, and J. A. Leppo. Human-observer receiver-operating-characteristic evaluation of attenuation, scatter, and resolution compensation strategies for 99mTc myocardial perfusion imaging. *J. Nuc. Med.*, 44(11):1725–34, November 2003.

[99] F. Natterer. *The mathematics of computerized tomography*. Teubner-Wiley, Stuttgart, 1986.

[100] F. Natterer. Inversion of the attenuated Radon transform. *Inverse Prob.*, 17(1):113–9, February 2001.

[101] R. G. Novikov. On the range characterization for the two-dimensional attenuated x-ray transformation. *Inverse Prob.*, 18(3):677–700, June 2002.

[102] K. Ogawa, Y. Harata, T. Ichihara, A. Kubo, and S. Hashimoto. A practical method for position-dependent Compton-scatter correction in single photon emission CT. *IEEE Trans. Med. Imag.*, 10(3):408–12, September 1991.

[103] J. M. Ollinger. Model-based scatter correction for fully 3D PET. *Phys. Med. Biol.*, 41(1):153–76, January 1996.

[104] J. Qi. Analysis of lesion detectability in Bayesian emission reconstruction with nonstationary object variability. *IEEE Trans. Med. Imag.*, 23(3):321–9, March 2004.

[105] J. Qi. Comparison of statistical reconstructions with isotropic and anisotropic resolution. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 6, pages 3624–8, 2004.

[106] J. Qi and R. H. Huesman. Theoretical study of lesion detectability of MAP reconstruction using computer observers. *IEEE Trans. Med. Imag.*, 20(8):815–22, August 2001.

[107] J. Qi and R. H. Huesman. Fast approach to evaluate MAP reconstruction for lesion detection and localization. In *Proc. SPIE 5372, Medical Imaging 2004: Image Perception, Observer Performance, and Technology Assessment*, pages 273–82, 2004.

[108] J. Qi and R. M. Leahy. A theoretical study of the contrast recovery and variance of MAP reconstructions with applications to the selection of smoothing parameters. *IEEE Trans. Med. Imag.*, 18(4):293–305, April 1999.

[109] J. Qi and R. M. Leahy. Resolution and noise properties of MAP reconstruction for fully 3D PET. *IEEE Trans. Med. Imag.*, 19(5):493–506, May 2000.

[110] D. Rabinowitz and D. Siegmund. The approximate distribution of the maximum of a smoothed Poisson random field. *Statistica Sinica*, 7(1):167–80, January 1997.

[111] A. J. Rockmore and A. Macovski. A maximum likelihood approach to emission image reconstruction from projections. *IEEE Trans. Nuc. Sci.*, 23:1428–32, 1976.

[112] J. G. Rogers. A method for correcting the depth-of-interaction blurring in PET cameras. *IEEE Trans. Med. Imag.*, 14(1):146–50, March 1995.

[113] J. P. Rolland and H. H. Barrett. Effect of random background inhomogeneity on observer detection performance. *J. Opt. Soc. Am. A*, 9(5):649–58, May 1992.

[114] S. Sankaran, E. C. Frey, K. L. Gilland, and B. M.W. Tsui. Optimum compensation method and filter cutoff frequency in myocardial spect: A human observer study. *J. Nuc. Med.*, 43(3):432–8, March 2002.

[115] S. S. Saquib, C. A. Bouman, and K. Sauer. ML parameter estimation for Markov random fields, with applications to Bayesian tomography. *IEEE Trans. Im. Proc.*, 7(7):1029–44, July 1998.

[116] K. Shafie, B. Sigal, D. Siegmund, and K. J. Worsley. Rotation space random fields with an application to fMRI data. *Ann. Stat.*, 31(6):1732–71, December 2003.

[117] M. Shahram and P. Milanfar. Imaging below the diffraction limit: a statistical analysis. *IEEE Trans. Im. Proc.*, 13(5):677–89, May 2004.

[118] L. A. Shepp and Y. Vardi. Maximum likelihood reconstruction for emission tomography. *IEEE Trans. Med. Imag.*, 1(2):113–22, October 1982.

[119] D. O. Siegmund and K. J. Worsley. Testing for a signal with unknown location and scale in a stationary Gaussian random field. *Ann. Stat.*, 23(2):608–39, April 1995.

[120] B. W. Silverman, M. C. Jones, J. D. Wilson, and D. W. Nychka. A smoothed EM approach to indirect estimation problems, with particular reference to stereology and emission tomography. *J. Royal Stat. Soc. Ser. B*, 52(2):271–324, 1990.

[121] E. T. P. Slijpen and F. J. Beekman. Comparison of post-filtering and filtering between iterations for SPECT reconstruction. *IEEE Trans. Nuc. Sci.*, 46(6):2233–8, December 1999.

[122] D. L. Snyder and M. I. Miller. The use of sieves to stabilize images produced with the EM algorithm for emission tomography. *IEEE Trans. Nuc. Sci.*, 32(5):3864–71, October 1985.

[123] E. J. Soares, C. L. Byrne, S. J. Glick, C. R. Appledorn, and M. A. King. Implementation and evaluation of and analytical solution to the photon attenuation and nonstationary resolution reconstruction problem in SPECT. *IEEE Trans. Nuc. Sci.*, 40(4):1231–1237, August 1993.

[124] J. W. Stayman and J. A. Fessler. Regularization for uniform spatial resolution properties in penalized-likelihood image reconstruction. *IEEE Trans. Med. Imag.*, 19(6):601–15, June 2000.

[125] J. W. Stayman and J. A. Fessler. Fast methods for approximation of resolution and covariance for SPECT. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pages 786–8, 2002.

[126] R. G. Swensson. Unified measurement of observer performance in detecting and localizing target objects on images. *Med. Phys.*, 23(10):1709–25, October 1996.

[127] C. Tocharoenchai, B. M. W. Tsui, D. P. Lewis, E. C. Frey, and X. Zhao. Compensation for the response function of medium energy collimator in /sup 67/Ga planar and SPECT imaging. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 3, pages 1405–8, 1998.

[128] G. D. Tourassi, B. S. Carey, C. E. Floyd, and M. T. Munley. Improved lesion detection in SPECT using MLEM reconstruction. *IEEE Trans. Nuc. Sci.*, 30(2):780–3, April 1991.

[129] M. Unser, A. Aldroubi, and M. Eden. B-spline signal processing: Part II—efficient design and applications. *IEEE Trans. Sig. Proc.*, 41(2):834–48, February 1993.

[130] M. Unser, P. Thevenaz, and L. Yaroslavsky. Convolution-based interpolation for fast, high quality rotation of images. *IEEE Trans. Im. Proc.*, 4(10):1371–81, October 1995.

[131] L. van Elmbt and S. Walrand. Simultaneous correction of attenuation and distance-dependent resolution in SPECT: an analytical approach. *Phys. Med. Biol.*, 38(9):1207–17, September 1993.

[132] H. L. Van Trees. *Detection, estimation, and modulation theory*. Wiley, New York, 1968.

[133] E. Veklerov and J. Llacer. Stopping rule for the MLE algorithm based on statistical hypothesis testing. *IEEE Trans. Med. Imag.*, 6(4):313–9, December 1987.

[134] E. Veklerov and J. Llacer. The feasibility of images reconstructed with the method of sieves. *IEEE Trans. Nuc. Sci.*, 37(2):835–41, April 1990.

[135] C. R. Vogel. Non-convergence of the L-curve regularization parameter selection method. *Inverse Prob.*, 12(4):535–47, August 1996.

[136] R. F. Wagner, D. G. Brown, and C. E. Metz. On the multiplex advantage of coded source/aperture photon imaging. In *Proc. SPIE 314, Digital Radiography*, pages 72–6, 1981.

[137] R. F. Wagner, K. J. Myers, K. M. Hanson, et al. Toward optimal observer perfomance of detection and discrimination tasks on reconstructions from sparse data. In K M Hanson and R N Silver, editors, *Maximum Entropy and Bayesian Methods*. Kluwer, New York, 1996.

[138] W. Wang, W. Hawkins, and D. Gagnon. 3D RBI-EM reconstruction with spherically-symmetric basis function for SPECT rotating slat collimator. *Phys. Med. Biol.*, 49(11):2273–92, June 2004.

[139] C. C. Watson, A. Schaefer, W. K. Luk, and C. M. Kirsch. Clinical evaluation of single-photon attenuation correction for 3D whole-body PET. *IEEE Trans. Nuc. Sci.*, 46(4-2):1024–31, August 1999.

[140] R. G. Wells, M. A. King, H. C. Gifford, and P. H. Pretorius. Single-slice versus multi-slice display for human-observer lesion-detection studies. *IEEE Trans. Nuc. Sci.*, 47(3):1037–44, June 2000.

[141] R. G. Wells, M. A. King, P. H. Simkin, P. F. Judy, A. B. Brill, H. C. Gifford, R. Licho, P. H. Pretorius, P. Schneider, and D. W. Seldin. Comparing filtered backprojection and ordered-subsets expectation maximization for small lesion detection and localization in Ga-67 SPECT. *J. Nuc. Med.*, 41(8):1391–9, August 2000.

[142] S. D. Wollenweber, B. M. W. Tsui, D. S. Lalush, E. C. Frey, K. J. LaCroix, and G. T. Gullberg. Comparison of Hotelling observer models and human observers in defect detection from myocardial SPECT imaging. *IEEE Trans. Nuc. Sci.*, 46(6):2098–103, December 1999.

[143] K. J. Worsley. Detecting activation in fMRI data. *Stat. Meth. Med. Res.*, 12(5):401–18, October 2003.

[144] K. J. Worsley, S. Marrett, P. Neelin, A. C. Vandal, K. J. Friston, and A. C. Evans. A unified statistical approach for determining significant signals in images of cerebral activation. *Hum. Brain Map.*, 4(1):58–73, 1996.

[145] Y. Xing, I-T. Hsiao, and G. Gindi. Rapid calculation of detectability in Bayesian single photon emission computed tomography. *Phys. Med. Biol.*, 48(22):3755–74, November 2003.

[146] M. Yavuz and J. A. Fessler. Statistical image reconstruction methods for randoms-precorrected PET scans. *Med. Im. Anal.*, 2(4):369–78, December 1998.

[147] A. Yendiki and J. A. Fessler. A comparison of rotation- and blob-based system models for 3-D SPECT with depth-dependent detector response. In *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med*, 2003.

[148] A. Yendiki and J. A. Fessler. Analysis of observer performance in detecting signals with location uncertainty for regularized tomographic image reconstruction. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 4, pages 2620–4, 2004.

[149] A. Yendiki and J. A. Fessler. Analytical approach to channelized Hotelling observer performance for regularized tomographic image reconstruction. In *Proc. IEEE Intl. Symp. Biomedical Imaging*, pages 360–3, 2004.

[150] A. Yendiki and J. A. Fessler. A comparison of rotation- and blob-based system models for 3D SPECT with depth-dependent detector response. *Phys. Med. Biol.*, 49(11):2157–68, June 2004.

[151] A. Yendiki and J. A. Fessler. Analysis of observer performance in known-location tasks for tomographic image reconstruction. *IEEE Trans. Med. Imag.*, 2005. To appear.

[152] D. F. Yu and J. A. Fessler. Mean and variance of singles photon counting with deadtime. *Phys. Med. Biol.*, 45(7):2043–56, July 2000.

[153] G. L. Zeng, J. R. Galt, M. N. Wernick, R. A. Mintzer, and J. N. Aarsvold. Single-photon emission computed tomography. In M N Wernick and J N Aarsvold, editors, *Emission tomography*, chapter 7. Academic Press, San Diego, 2004.

[154] G. L. Zeng and G. T. Gullberg. Frequency domain implementation of the three-dimensional geometric point response correction in SPECT imaging. *IEEE Trans. Nuc. Sci.*, 39(5-1):1444–53, October 1992.

[155] G. L. Zeng and G. T. Gullberg. Unmatched projector/backprojector pairs in an iterative reconstruction algorithm. *IEEE Trans. Med. Imag.*, 19(5):548–55, May 2000.

[156] Z. Zhou, R. M. Leahy, and J. Qi. Approximate maximum likelihood hyperparameter estimation for Gibbs priors. *IEEE Trans. Im. Proc.*, 6(6):844–61, June 1997.

[157] G. Zubal, G. Gindi, M. Lee, C. Harrell, and E. Smith. High resolution anthropomorphic phantom for Monte Carlo analysis of internal radiation sources. In *IEEE Symposium on Computer-Based Medical Systems*, pages 540–7, 1990.

ABSTRACT

ANALYSIS OF SIGNAL DETECTABILITY IN STATISTICALLY RECONSTRUCTED TOMOGRAPHIC IMAGES

by

Anastasia Yendiki

Chair: Jeffrey A. Fessler

Imaging in general, and emission tomography in particular, has become an important tool in many areas of medical diagnosis. Several common applications of emission tomography, such as the diagnosis of lung tumors or myocardial perfusion defects, involve the detection of a spatially localized target signal in an image reconstructed from noisy data. Such detection tasks are affected by various design parameters of the imaging system and reconstruction algorithm. This thesis is concerned with optimizing regularized image reconstruction methods for emission tomography with respect to the detectability of a spatially localized target signal in the reconstructed images.

We first consider the task of detecting a statistically varying signal of known location on a statistically varying background in a reconstructed tomographic image. We show that a broad family of linear observer models can achieve exactly optimal detection performance in this task if one chooses a suitable reconstruction method. This conclusion encompasses several well-known models from the literature, including those with a frequency-selective

channel mechanism. Interestingly, the "optimal" linear reconstruction methods for many of these observer models are unregularized and in some cases quite unconventional. In the case of channelized models in particular, the observer's ability to prewhiten determines the extent to which its detection performance can benefit from regularization. That is, regularization is more important for channelized observers that have incomplete knowledge of the second-order statistics of the reconstructed images.

Subsequently, we investigate detection tasks where the location of the target signal is unknown to the observer. This location uncertainty complicates the mathematical analysis of observer performance significantly. We consider model observers whose decisions are based on the maximum value of a linear local test statistic over all possible signal locations. Several of our conclusions about the known-location task extend to this case. Previous approaches to this problem have used Monte Carlo simulations to evaluate the localization performance of maximum-statistic observers. We propose an alternative approach, where approximations of tail probabilities for the maximum of correlated Gaussian random fields facilitate analytical evaluation of detection performance. We illustrate how these approximations can be used to optimize the probability of detection (at low probabilities of false alarm) for the observers of interest.