2010

Augmented breast tumor classification by perfusion analysis

Bruce Yu-Sun Lin
Louisiana State University and Agricultural and Mechanical College

AUGMENTED BREAST TUMOR CLASSIFICATION BY PERFUSION ANALYSIS

A Dissertation
Submitted to the Graduate Faculty of
the Louisiana State University and
Agricultural and Mechanical College
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

in

The Department of Computer Science

by
Bruce Lin
B.S., Computer Science, Louisiana State University, 2002
December 2010

Acknowledgments

Even with the challenges of rebuilding the LSU Health Sciences Center following the disaster of Hurricane Katrina, the physicians and researchers found a way to accommodate and even encourage research presented here. To Dr. Leonard Bok and Dr. Luis Serrano: I express my deep appreciation for your critical roles in acquiring the experimental data used in my research. You were also mentors to me, guiding and refining my understanding of the problem domain. Had you not decided to dedicate your time to me, this research simply would not exist.

To my colleagues at the Naval Research Laboratory–Dr. John Sample, William Avera, Kevin Shaw, and Elias Ioup: I have benefited from your skills, good humor, and understanding as I strove to complete this milestone. I am surrounded by excellence.

Lastly, and most profoundly, I would like to thank Dr. John M. Tyler for his endless encouragement, guidance, and expertise as my major professor. In the most difficult times, you managed to keep me engaged and productive. I simply do not deserve the level of effort that you gave me, and I hope to be a worthy addition to your legacy of students.

I am greatly indebted to all your contributions.

Table of Contents

List of Tables

List of Figures

Abstract

Magnetic resonance and computed tomography imaging aid in the diagnosis and analysis of pathologic conditions. Blood flow, or perfusion, through a region of tissue can be computed from a time series of contrast-enhanced images. Perfusion is an important set of physiological parameters that reflect angiogenesis. In cancer, heightened angiogenesis is a key process in the growth and spread of tumorous masses. An automatic classification technique using recovered perfusion may prove to be a highly accurate diagnostic tool. Such a classification system would supplement existing histopathological tests, and help physicians to choose the most optimal treatment protocol. Perfusion is obtained through deconvolution of signal intensity series and a pharmacokinetic model. However, many computational problems complicate the accurate-consistent recovery of perfusion. The high time-resolution acquisition of images decreases signal-to-noise, producing poor deconvolution solutions. The delivery of contrast agent as a function of time must also be determined or sampled before deconvolution can proceed. Some regions of the body, such as the brain, provide a nearby artery to serve as this arterial input function. Poor estimates can lead to an over or under estimation of perfusion. Breast tissue is an example of one tissue region where a clearly defined artery is not present. This proposes a new method of using recovered perfusion and spatial information in an automated classifier. This classifier grades suspected lesions as benign or malignant. This method can be integrated into a computer-aided diagnostic system to enhance the value of medical imagery.

Chapter 1

Introduction

Cancer is disease of unchecked cell division that is a leading cause of death for persons under 85. Exposure to carcinogenic compounds, dietary content, genetic predisposition, and age are all variables that may increase an individual's risk of developing cancer. Over 200 variants affect nearly every bodily system, for which different screening procedures and treatments are employed. Different sectors of the population are more susceptible to certain forms of cancer. Breast cancer, the focus of this study, is the second most common form after lung cancer, and accounts for a third of all new cases of cancer in females [JMW+05].

Cancerous cells invade healthy tissues and interfere with their normal function. Lymphoma, for example, attacks the immune system and hinders the body's ability to ward off infection. A tumor may develop and remain benign, not spreading from its point of origin. However, if a portion of a tumor breaks from the main mass and spreads to other parts of the body, that tumor is said to have metastasized. These metastacized cells may invade different tissues and grow into new tumors.

Early detection and treatment is key to a patient's survivability of the disease. Physicians may treat the patient with a combination of drug therapy, radiation therapy, and/or surgical removal of the tumor. Attacking the cancer with the least amount of damage to healthy tissue requires accurate diagnostic information from histopathological analysis, serological analysis, and other diagnostic tools such as medical imaging.

Though biopsy is ultimately the gold standard for cancer diagnosis, medical imaging is a valuable tool for gathering information *in vivo* for both screening and diagnosis. Radiologists are able interpret this imagery and recognize characteristics associated with cancer morphology. Accurate interpretation of the imagery assists in determining the best actions for further analysis and treatment.

There are a number of tools for breast imaging. X-ray mammography transmits a dose of radiation through the breast tissue, and captures the projected image on film or digital medium. Com-

puted tomography (CT), generates a 3-dimensional volume by aggregating 2-dimensional slices. Measures are taken to mitigate the patient's exposure to radiation. Magnetic resonance imaging (MRI) manipulates the spin and orientation of mostly hydrogen atoms in the tissue, differentiating the soft tissues in the body by this water content. These imaging processes are powerful aids for location and characterization of abnormal conditions.

Beyond morphology, perfusion, the flow of blood through tissue, may reveal a great deal of information about a suspected tumor. Tumors larger than a few millimeters require additional vasculature to deliver oxygen and nutrients to sustain this growth. The new blood vessels exhibit irregularities that allow larger molecules to pass into the extravascular space. By adding indicator particles into the blood stream, this "leakage" behavior can be captured and compared to how cancerous tissue has been observed to respond.

To acquire contrast-enhanced perfusion imagery with MRI, the patient is injected with a paramagnetic contrast agent, such as gadolinium or feroglucose. The tissue is imaged over time as the contrast agent circulates through the body. As the contrast arrives on the tumor periphery, the MRI may register a signal change at each tissue voxel. Signal analysis of this time series can characterize the rate of blood flow, the volume of the leakage space, and how long any individual indicator particle takes to traverse a unit volume. Our study uses the contrast-enhanced imagery produced by the `VIBRANT` protocol (Volume Imaging for BReast AssessmeNT, General Electric).

Computing these perfusion characteristics for known healthy, benign, and malignant test cases generates a great quantity of data to train a judiciously chosen machine learning algorithm. New cases could then be classified in both screening and diagnostic roles. It is our goal to understand this enhancement process, compute perfusion parameters, and construct a classification framework that produces a machine-generated prediction. Such a framework would be a valuable step in building the next generation of computer-aided diagnostic tools in the domain of cancer research. We will present the parts of this framework that have the greatest impact on improving detection and classification accuracy.

This research addresses several challenges in designing a classification framework. Computational perfusion recovery must be able to handle low signal-to-noise, image registration artifacts, as well as physiological variations between patients. We will examine the signal enhancement as the convolution of an arterial input function (AIF) and a tissue response kernel. The perfusion parameters of flow, volume, and mean transit time are obtained from this tissue kernel. The arterial input function is selected from a fitted pharmacokinetic model (e.g., Toft's), or sampled from an adjacent blood vessel in the image series.

Several options for the training and classification are explored that exploit the recovered perfusion information. Cluster analysis and statistical learning methods are applied to a feature space that we developed, and the results are compared to the model-based three time-points (3TP) method for enhancement analysis. As the name suggests, the 3TP method reduces the enhancement curve to information to three time points. We found that incorporating additional information in the classifier training yields better performance in terms of accuracy.

Chapter 2 presents a foundation of current research on this topic examining the physiology of cancer growth, methods of cancer imaging with emphasis on MRI, the perfusion convolution model, and machine classifiers applied to medical diagnosis.

Chapter 3 describes the nature of the experimental data and develops the method for perfusion recovery by deconvolution. The types of pathology found in the experimental data are also presented. Chapter 3 poses the feature vector construction used for classification. The outline for a procedure to train and test machine learning algorithms, such as the k-nearest neighbor and support vector machine learners is presented.

Chapter 4 organizes the results of the various training methods and classifier results in terms of accuracy, specificity, sensitivity, and clinical observations by a radiologist.

In Chapter 5, the performance of the classification framework results is discussed in the larger context of screening and diagnosis.

Chapter 2

Background

2.1 Cancer Detection and Diagnosis

2.1.1 Overview

Early detection of cancer greatly increases the probability of survival for a patient. For cancers of the breast, the American Cancer Society recommends annual mammography screenings for women after the age of 40 [ZACS01]. If a suspicious lesion is identified in a initial screening, a physician may employ other methods to characterize what that lesion may be. Evidence may support further action where part or all of the lesion is removed in a biopsy procedure for pathological analysis. Here, at a microscopic scale, the tissue is scrutinized for tell-tale signs of cancer. Figure (2.1) shows a basic diagnosis and treatment procedure for breast cancer.

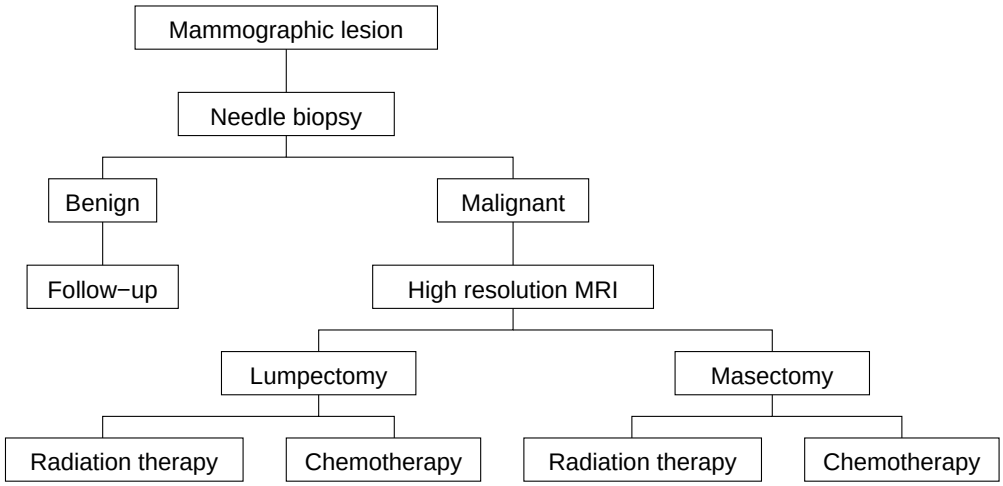

Figure 2.1: Typical breast cancer screening and treatment using MRI (Harms, 1994, Figure 3)

Tissue biopsy followed by pathological analysis remains the gold standard for cancer diagnosis. Pathologists produce a highly accurate diagnosis by analyzing the microvascular structure and cellular characteristics of the tissue. Under certain conditions where invasive cancer is suspected, the physician may recommend that the entire mass is removed in a surgical biopsy. When the classification of a lesion is unsure, fine needle aspiration biopsy is a less invasive option to sur-

gical biopsy that can also produce a highly accurate diagnosis. In this procedure, a small needle is inserted into the mass, extracting cells to be examined. Upon microscopic examination, the determination of benign, malignant, or uncertain for either, is made. The combination of several diagnostic tools, used together, can greatly reduce the risk of missing a positive case for cancer [LAC01].

Preempting an unnecessary surgical procedure not only avoids the usual patient risks and cost of surgery, but may also possibly prevent triggering further tumor growth. Protein factors that promote cell division are released during the wound healing process and may inadvertently trigger a tumor to regrow. To mitigate this risk, drugs that selectively inhibit some cell regrowth might be administered to the patient. One such post-operative drug, Bevacimuzab, binds to VEGF (vascular endothelial growth factor), inhibiting the growth of new blood vessels.

Providing the physician more information from pre-surgical techniques, such as with diagnostic imaging, can only improve the decisions made for patient treatment. The following sections develop the background to understand how diagnostic imaging can be used by a computer-aided method to perform breast cancer detection and classification. Computer-aided diagnostic imaging will not replace the gold standard methods soon, but may reduce the cost and time for analysis of the raw data. A combination of different methods could improve determining the number, location, and classification of any suspicious lesions.

2.1.2 Physiological Characteristics

The predominant form of cancer found in the breast occurs in the lining of the lactiferous ducts that transport milk from the mammary lobes to the nipple. The early form of this condition, having not broken through the endothelial lining of the duct, is referred to as pre-cancerous *ductal carcinoma in-situ*. If the tumor continues to grow and breaks through the endothelial lining, the condition becomes *ductal carcinoma invasive*, and may spread into surrounding tissues. A tumor is said to have *metastasized* if a part of the tumor breaks off and begins growth in another part of the body.

Tumors which grow to larger than several millimeters require vascular support. Angiogenesis, the process of growing new blood vessels, normally occurs during wound healing and healthy tissue growth. New vessels sprout from existing vessels when signaled by growth factors. In a scenario that might involve a tumor, abnormally heightened angiogenesis could indicate invasive behavior or the possibility of metastasis [LMS+05, ELPC+03].

The vasculature in the neighborhood of a tumor is also discernible from normal vasculature, possessing more fenestra, or openings [CP04]. This increased fenestration is likely due to elevated secretion of vascular endothelial growth factor (VEGF) by tumorous tissue [RP95]. As a consequence, larger molecules can pass from the blood into the extravascular space, or leakage space [Pad03, DLRH+03]. This increased exchange of macromolecules is used by certain imaging techniques, described later. The microvascular density (MVD) is another metric that quantifies the blood vessels present in a unit of tissue volume [MSTK99, MCT+02]. Research in serological methods attempt to detect specific antigens in bodily fluids that might signal the presence of cancer [OC98].

Benign growths that do not become invasive are most likely to be *fibroadenoma* [WSC+99]. Fibroadenoma exhibit a lower microvascular density, and are typically composed of fibrous and glandular tissue. Fibroadenoma may be discovered by palpation, and subsequently be examined for any evidence of growth. *Microcalcifications* are also a benign phenomenon that may be imperceptible to the touch, yet may appear as opaque specks on a breast mammogram [PFL05]. *Breast cysts* are another benign, fluid-filled structure that usually do not require treatment after diagnosis.

The physician is tasked with discriminating between these different abnormal conditions. Depending on the imaging protocol used, benign conditions and malignancy could be very difficult to differentiate. In the early phases of cancer, a very small growth might be overlooked in some kinds of imagery.

2.1.3 Diagnostic Imaging

A brief survey of existing diagnostic imaging protocols clarifies the challenges in constructing a computer-aided classification framework. The accuracy of a diagnostic test is a combination of its correctly predicted disease-present and disease-free conditions (true-positives T^+, and true-negatives T^-), as well as the number of mispredicted disease-present and disease-free conditions (false-positives F^+, and false negatives F^-) equation (2.1). The sensitivity of a diagnostic equation (2.2) reflects how likely a test is to miss a disease-present condition. Conversely, specificity measures the rate at which the test is fooled by a disease-free state equation (2.3). Both metrics are needed to evaluate a diagnostic protocol's accuracy.

$$Accuracy = \frac{T^+ + T^-}{T^+ + T^- + F^+ + F^-} \tag{2.1}$$

$$Sensitivity = \frac{T^+}{T^+ + F^-} \tag{2.2}$$

$$Specificity = \frac{T^-}{T^- + F^+} \tag{2.3}$$

In 2-dimensional mammograms, physicians have a standard vocabulary for describing the morphological and visual characteristics of a feature in a mammogram (Figure 2.2). Such a lexicon collects terms for categorizing the shape, number, and density of tumors in a mammogram [BKF96]. The Breast Imaging Reporting and Data System (BI-RADS), for instance, defines terms for a lesion's margin, mass, calcification distribution, and other features that might be observed in a mammogram. For example, "spiculated," "circumscribed," and "lobular" might describe the boundary of a lesion. These qualitative descriptions have been correlated to the actual disease conditions, and establish a standard of communication between radiologists. However, because of its 2-dimensional nature, the structures may overlap and obscure other features.

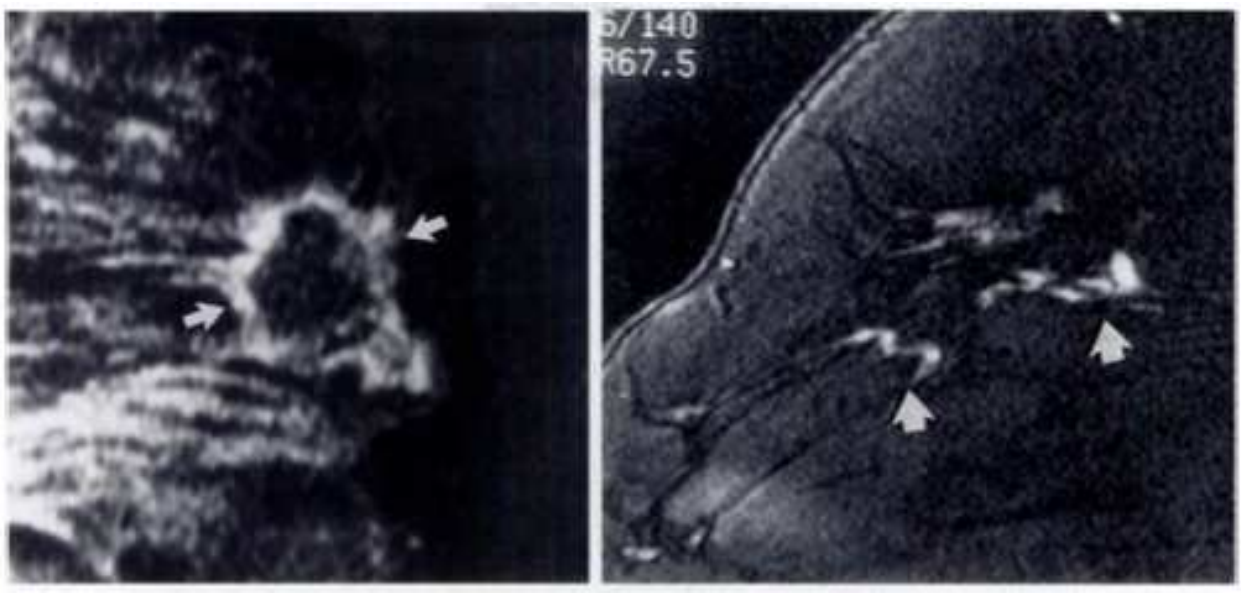

Figure 2.2: Morphologic features of a malignant tumor **(left)** and benign features **(right)** (Orel, 1998, Figure 3)

The screening and diagnostic performance of mammography alone is largely influenced by the radiologist's interpretation. In one study, 10 radiologists read 150 mammographic studies each and rated these visual characteristics. For some characteristics, such as calcification distribution, the radiologists' assessments had substantial agreement. There was much more variance in categorizing lesion shape and density. The results also showed higher agreement between radiologists who had received the same training. While there is a proven value for manual interpretations of mammography, there is notable variability in specificity and sensitivity depending on a radiologist's training [EWL$^+$94].

Radiologists employ other means of imaging for screening and diagnosis: computed tomography (CT), ultrasound, positron-emission tomography (PET), and magnetic resonance (MR) [BBH$^+$99, SAS00]. Each method is subject to conditions that might introduce artifacts. Poor registration from movement or respiration is more likely to affect an imaging protocol with a longer acquisition time. When a contrast agent is administered prior to a longer-term imaging process, there are effects that arise from recirculation. Both spatial and temporal resolution have a pronounced affect on the level of signal-to-noise in CT and MRI [Mil03].

Characterization of choline metabolism with MR spectroscopy (MRS) has also been shown to differentiate benign and malignant lesions [Yeu01, KBSRS$^+$02]. Choline is an essential nutrient produced during lactation. But in the presence of malignant cancer, increased levels of choline and choline metabolites might be used as a marker. The levels of choline metabolites detected by MR spectroscopy can be mapped on to a 3-dimensional volume as a visualization and diagnosis aid. This method can improve specificity of breast cancer detection when coupled with contrast enhanced MRI.

Combining information from a variety of imaging methods may improve accuracy. [MSSS$^+$96] has explored the integration of conventional mammography, sonography, and contrast-enhanced MRI for cancer diagnosis. Investigators evaluated each test case in each imaging method independently, classifying each study as being disease-free, having a benign lesion, or having a malignant lesion. Specificity was the most improved metric by using this multi-method diagnosis when compared with pathological findings.

No non-invasive diagnostic imaging technique should be used in isolation when considering risk of a false-negative prediction. When presented with uncertainty, treatment still tends towards using the gold-standard method of biopsy and pathological analysis. Regardless, imaging provides a great deal of information. The imaging of contrast-enhanced MRI is especially useful for characterizing perfusion, which is sought to be a highly accurate screening and diagnostic method.

2.1.4 Perfusion Analysis

Perfusion estimates computed from contrast-enhanced imagery is an effective, non-invasively way to gauge vascularity [CP04, POM$^+$09]. Ultimately, perfusion can be used to determine how likely a tumor is to metastasize [Pad99]. In addition to detection and diagnosis, perfusion can be used to monitor the effectiveness of anti-angiogenic cancer treatments. Bevacimuzab, mentioned in Section 2.1.1, inhibits the growth of new vasculature, though not directly attacking the tumor [GPM$^+$03]. Depriving the tumor of oxygen and nutrients slow maturation and may encourage

regression. Perfusion imaging, in this case, would show the drug affect over successive imaging sessions.

Each volume element, or voxel, is the smallest unit of resolution in an MR image. By recovering the "amount" of perfusion, in relative terms, the blood flow through the tissue can be color-mapped on the image. These maps illustrate areas of elevated perfusion, found near tumors, and depressed perfusion, such as in necrotic tissue. Absolute perfusion, in units of fluid per unit tissue per unit time, is difficult to quantify accurately. Relative perfusion can show comparative differences in adjacent tissues, which may be as valuable in a diagnostic sense.

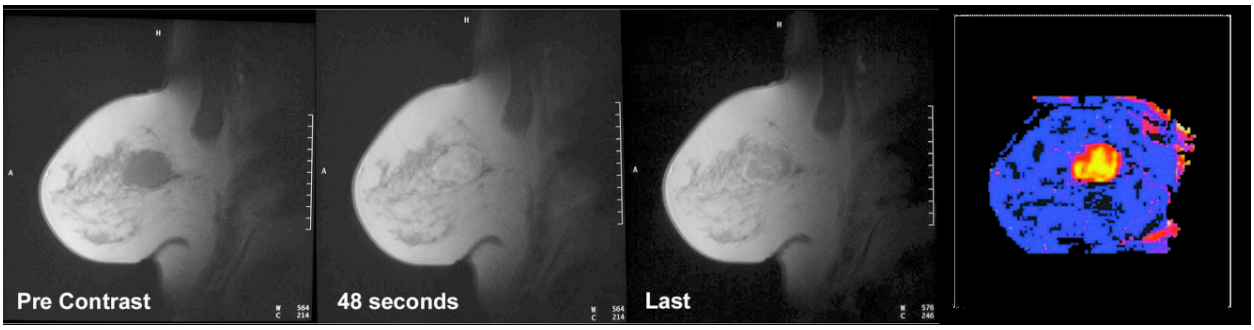

Figure 2.3: Pre-Contrast breast (left) tumor imagery followed by post contrast. Relative blood flow (right) is colored from blue (low) to yellow (high). (Padhani, 2003, Figure 4, 5)

With basic understanding of breast cancer physiology, particularly tumor vasculature characteristics, the methods of relative perfusion recovery can be explained. The process of contrast administration, circulation and imaging are described then related to the indicator-diluter model, central volume theorem, and convection-diffusion model. Here, a formulation of perfusion as a convolution equation is presented and solved analytically.

2.1.5 Detection in MRI

The exploitation of magnetic resonance physics for imaging is a complex phenomenon and a rudimentary understanding of its mechanics is important to how the experimental data will be addressed. MRI exploits the fact that the majority of body mass is water, and therefore has a large proportion of hydrogen. MRI for medical imagery uses a specific radio frequency pulse that ex-

cites the hydrogen atoms, then detects the signal strength of energy released that is released as the hydrogen atoms "relax." Because different tissues have differing quantities of water, an image can be reconstructed that shows contrast between these tissues [Hor97].

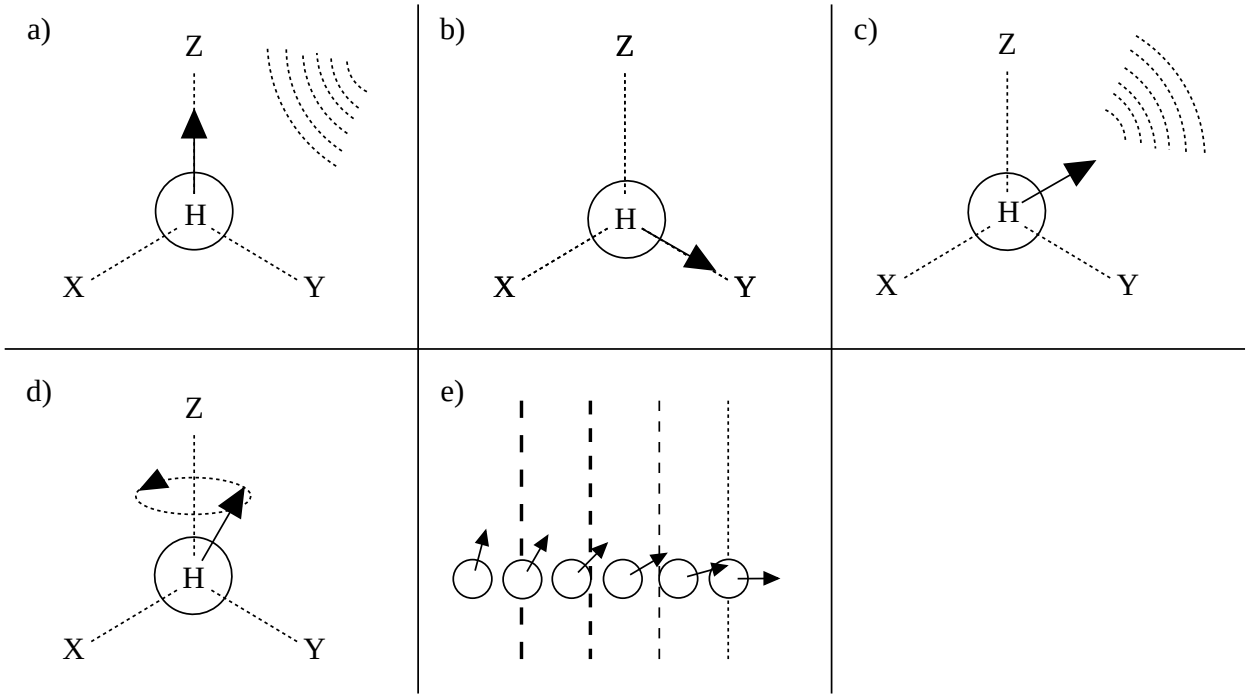

Figure 2.4: a. Hydrogen nuclei at equilibrium is excited by RF pulse, b. Magnetic moment vector is deflected from equilibrium, c. Exponential decay of energy as spin returns to equilibrium, d. Precession occurs about the Z-axis, used in spin encoding e. Magnetic field linear gradient used in phase encoding

At the quantum level, each nuclear particle possesses a spin direction, and a magnetic moment vector associated with this spin. When the body is placed in the magnetic field of the MRI, the magnetic moments of the hydrogen nuclei will align with this field in a quantity proportional to the field strength and temperature. A group of hydrogen spins experiencing the same magnetic field is referred to as a spin packet, which has a net spin vector. When aligned with the magnetic field, a spin packet is typically in its low-energy, equilibrium state. A radio pulse with a frequency specific to hydrogen will be absorbed and cause it to deflect to a higher energy orientation. This frequency is the Larmor frequency, and for hydrogen this value is $\nu = 42.85 MHz/T$.

After this excitation, the energy will decay exponentially and realign with the magnetic field. By convention, the coordinate system is aligned such that the positive Z-axis is in the direction of the magnetic field. The time needed to return the magnetic moment vector in the direction of the magnetic field by a factor of e is the spin-lattice relaxation time, T_1. This is described in equation (2.4) where M_z the magnetic moment in the direction of the magnetic field, M_0 is the initial magnetic moment, t is elapsed time, and T_1 is a characteristic of the tissue being examined.

$$M_z = M_0(1 - e^{-t/T_1}) \tag{2.4}$$

The relaxation time of the magnetic moment in the transverse plane (XY) is similarly represented, and referred to as the T_2 spin-spin relaxation time. Because of the orientation of the coordinate system, the transverse magnetic moment vector is at maximum after a 90° excitation, then decays to 0 as the spin realigns with the magnetic field. It holds that $T_2 \leq T_1$. equation (2.5) describes this as a function with respect to time.

$$M_{xy} = M_{xy_0} e^{-t/T_2} \tag{2.5}$$

Due to slight variations in the magnetic field, or inhomogeneities, there is a T_2^* relaxation time that is made up of the "pure" T_2 time and $T_{2_{inhom}}$ related by equation (2.6).

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_{2_{inhom}}} \tag{2.6}$$

Other mechanisms are used to spatially localize each pixel in 2-dimensions in an MR image slice. As a spin packet is excited from equilibrium, the magnetic moment vector will also rotate, or precess, about the Z-axis at a frequency corresponding to the frequency of excitation. By manipulating the magnetic field as a linear gradient along either the X or Y axis, the spin packets can be differentiated by their rate of spin. If the magnetic field is turned off, these spin packets will halt and continue to decay energy as an RF signal. As these signals are detected in the RF receiver, each

pixel location has a characteristic phase and frequency encoding that can be separated by Fourier methods.

The RF pulses and magnetic field activations are executed in specific imaging sequences that vary frequency and timing. Generally, there are three timing subintervals that can be manipulated: repetition time (TR), inversion recovery time (TI), and spin-echo time (TE). There is a marked difference in how tissues will contrast with each other when these parameters, TE, TR, and TI are varied. Different kinds of tissue, like fat, bone, and blood vessels, will have different characteristic relaxation curves values for T_1, T_2, and T_2^*. Specific combinations of TR, TI, and TE will cause an intensity difference between tissues with different relaxation times.

A sequence that maximizes contrast between tissues with differing T_1 values is referred to as a T_1-weighted image. Tissues with a short T_1 will appear brighter on a T1-weighted image than a tissue with a longer T_1. However, there may be certain kinds of tissue that can be better represented with a T_2-weighted or T_2^*-weighted imagery. Even greater image contrast can be achieved when a weakly magnetic material is added to some tissues. This is the basis for contrast-enhanced MRI, and a contrast agent specially selected for breast tumors would show a unique enhancement pattern.

Prior to MRI acquisition, a contrast agent is often administered to the patient by mechanical injection. The central concept behind tumor perfusion imaging is that the fenestrae, or openings, present in the tumor vasculature are large enough to accommodate these macromolecules to travel into the extravascular space. Experiments have determined that large molecular-weight contrast agents are more effective at measuring endothelial permeability [LLBT$^+$05]. For MRI, a paramagnetic contrast agent such as gadolinium chelates is typically chosen. Each gadolinium atom is sequestered in a ligand structure to render it safe.

A pre-contrast image is acquired before contrast administration. Then, the contrast is injected and circulates from the injection point to the heart and lungs, then out into the body. After some elapsed time, the contrast arrives at the breast tissue and to any lesions that may be present. The expectation is that the gadolinium will accumulate in the leakage space in the vicinity of a tumor. A higher degree of permeability to the contrast may correlate with an invasive or metastasized

tumor. Over time, as blood continues to circulate, the accumulated contrast will be taken back into the blood stream.

The extravascular, extracellular space (EES) in each unit of tissue is filled with fluid that affects the rate of contrast accumulation and uptake [YLL+05]. This proportion of EES to non-EES space is called the extracellular volume fraction, or EVF. There is a single-compartment and a two compartment fluid model that attempt to model this behavior.

2.1.6 Three Time-Points (3TP) Method

The standard method of correlating a signal enhancement curve from contrast to values of EVF and tissue permeability is the 3-Time Points method (3TP). 3TP is employed in many MRI CAD systems, such as the GE CadStream product. 3TP examines the ramp-up behavior as the contrast agent arrives at a tissue voxel (Figure 2.5) [Kuh99, Ore98]. By fitting a pharmacokinetic model with three specifically chosen time points, the tissue permeability and extracellular volume fraction coefficients can be mapped with a 3-color scheme that can indicate a suspicious condition [KRV+00]. Tissue segmentation using 3TP has been explored using basic morphological operations, such as erosion and dilation [LKP+08a].

The creation of the 3TP enhancement map can be generalized by the following method. Three time points are chosen, typically the pre-contrast, contrast arrival, and post-contrast. For each voxel based on these three time points, it is binned into three curve types based on basic curve shape (Figure 2.5).

3TP assigns a color and intensity that indicates membership with three curve families. Type I curves exhibit a monotonic increase throughout the time series. Type II curves exhibit signal increase from the first to second time point, then remains relatively constant. Type III curves exhibit a signal increase between the first two time points, then a "wash-out" phase between second and third time points where there is an intensity decrease. Generally, voxels with a Type II curve are associated with tissue in a benign mass, whereas Type III curves may be suspect for an invasive condition. Algorithm 1 outlines the 3TP algorithm. The next sections develop the basis for perfu-

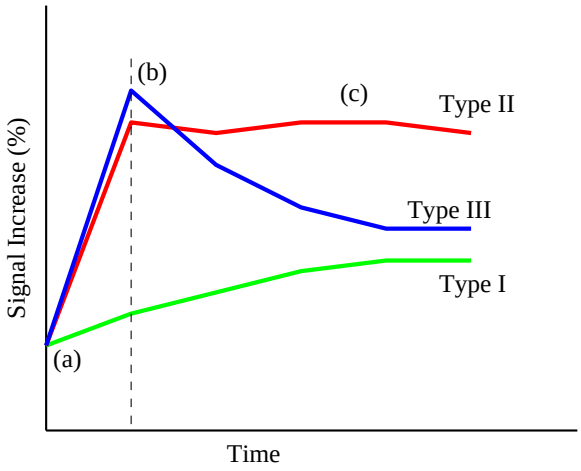

Figure 2.5: 3TP Time-intensity curve types. Type I shows a gradual signal increase. Type II is described as a plateau. Type III shows a sharp washout. Time points (a), (b), and (c) are pre-contrast, contrast arrival, and post-contrast time points, respectively.

sion behavior, and perfusion parameter recovery, with a goal of being able to exploit the remainder of the signal series for use in a classifier.

Algorithm 1 The 3TP Algorithm

for all Voxel $v \in V$ **do**
  Let the first intensity change $\alpha = t_1 - t_0$
  if $\alpha > 0$ **then**
    Let the second intensity percentage change

$$\beta = \frac{t_2 - t_1}{\alpha}$$

    if $\beta > 0.1$ (increases by more than 10%) **then**
      Let blue intensity $B = \alpha$
    else if $\beta < 0.1$ (decreases by more than 10% **then**
      Let red intensity $R = \alpha$
    else
      Let green intensity $G = \alpha$
    end if
  end if
end for
The voxel color $C \leftarrow (R, G, B)$

2.1.7 Indicator-Diluter Model

In the simplest of perfusion models, each measurable unit of tissue is represented as a fixed volume with a single constant inflow and single constant outflow of fluid. In the imagery, this is the smallest unit of tissue that can be measured in one volumetric pixel. The fluid in this model is blood mixed with the injected contrast agent, or indicator that can be described as the concentration of indicator per fluid unit. Zierler developed the indicator-diluter theory that is the basis for current models of perfusion [Zie62]. As described, Zierler represents a single fluid compartment fed by a input artery and drained by a output vein (Figure 2.6).

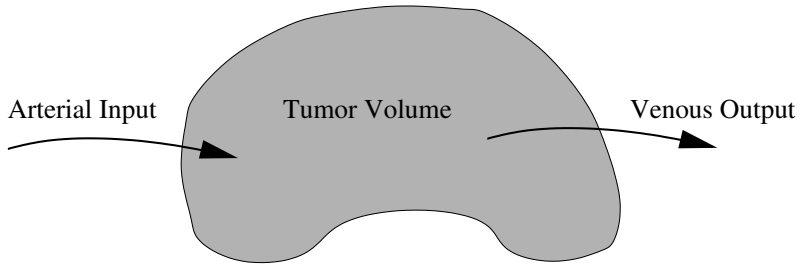

Figure 2.6: Simple single-compartment model for tumor blood flow

In Zierler's model, this fluid compartment has constant rates of blood flow at entrance and exit, and also has constant total volume. At time $t = 0$, a finite quantity of indicator is impulsively introduced at the input artery. Let m be the amount of indicator introduced into the compartment, and let F be the constant rate of blood flow. The model assumes that all of the indicator will drain from the compartment in a finite amount of time. Let $c(t)$ be the concentration of indicator leaving the compartment at time t. With these assumptions, the amount of blood leaving the compartment in an infinitesimal amount of time is $F \cdot dt$. The amount of indicator that leaves with this volume at time t is $c(t) \cdot F \cdot dt$. This quantity "summed" over all time is the total amount of indicator, or m. (2.7).

$$m = F \int_0^\infty c(t)dt \tag{2.7}$$

It is useful to describe the fractional amount of indicator that leaves the compartment at time t as a function, $h(t)$. Because all the indicator leaves the compartment, the sum of $h(t)$ over all time is unity.

$$\int_0^\infty h(t)dt = 1$$

To solve for the total volume of fluid V passing through the compartment, the model defines a infinitesimal unit of volume, dV, that contains the concentration of indicator exiting the compartment at time $t + dt$. The rate at which dV exits is described by $F \cdot h(t)dt$. dV is equal to this fractional rate times t (2.8). The total volume of blood is obtained by summing over all time (2.9).

$$dV = F \cdot t \cdot h(t)dt \tag{2.8}$$

$$V = F \int_0^\infty t \cdot h(t)dt \tag{2.9}$$

Observe that $\int_0^\infty t \cdot h(t)dt$ is the mean value of t. This is better understood as the mean transit time, $\bar{t}$, or the average time that a small amount of indicator takes to drain from the compartment. The relationship between V, F, and $\bar{t}$ is the central volume theorem:

$$\bar{t} = \frac{V}{F} \tag{2.10}$$

This model of fluid flow is based on several critical assumptions. In addition to constant blood flow and volume of the fluid compartment, every indicator particle enters and exits the volume only once. In a physiological context, there can be no recirculation of the fluid and indicator. This model also assumes that $h(t)$ is stationary: the distribution of indicator transit times remains time invariant. Furthermore, $h(t)$ is assumed to match the distribution of transit times of the blood. This implies a complete homogeneous mixing of indicator and blood.

In reality, these assumptions are rarely satisfied. Effects from recirculation do appear in the image following the initial arrival indicator. Also, the manner of injection and vessel structure

affect the arrival of indicator [Zie62]. As the bolus of contrast agent flows through the body, it passes into the heart and through the lungs before being arriving at the region of interest. Only a fraction of the initial injection will circulate through the tumor. Instead of an impulsive arrival, the concentration will gradually increase over a small span of time. However, the indicator-diluter model provides the basis for more sophisticated models of perfusion.

Other tissue models use multiple compartments to describe the exchange of macromolecular indicator particles. A two-compartment model, for example, may represent the vascular space and leakage space, with the exchange of indicator concentrations driven by a simple two-way transfer function [Dal98]. The vascular integrity is characterized by constants that correlate with permeability parameters within the tumor.

2.1.8 Convection-Diffusion Model

The convection-diffusion model of perfusion extends the ideas of Zierler's indicator-diluter model with a more realistic physiological representation. Here, indicator particles move through the tissue region with two transfer methods. Convection transfers indicator particles by fluid motion through the vascular bed, and diffusion transfers indicator particles in both directions across blood vessel membranes. The vasculature is then represented as a series of parallel convective tubes that exchange the indicator with the extravascular space via diffusion. These tubes are so numerous that from a macroscopic scale, they form an anisotropic continuum [KBKS04].

Consider the region Ω, consisting of N discrete convective tubes. Like the indicator-diluter model, the flow in and out of Ω is temporally independent. The blood flow is also spatially constant with respect to the boundaries of Ω. Let C_i be the concentration of indicator and V_i be the volume of i-th compartment. The instantaneous rate of change of indicator quantity, $V_i C_i'$, is the sum of contributions from other the remaining $N-1$ compartments. The $(i-1)$ compartment contributes indicator concentration to the i compartment via convection. Any other of the $j \neq i$ compartments may exchange indicator concentration via diffusion. The constants k_i and k_{ij} describe the rate of the convection and diffusion, respectively. The rate of concentration change for a

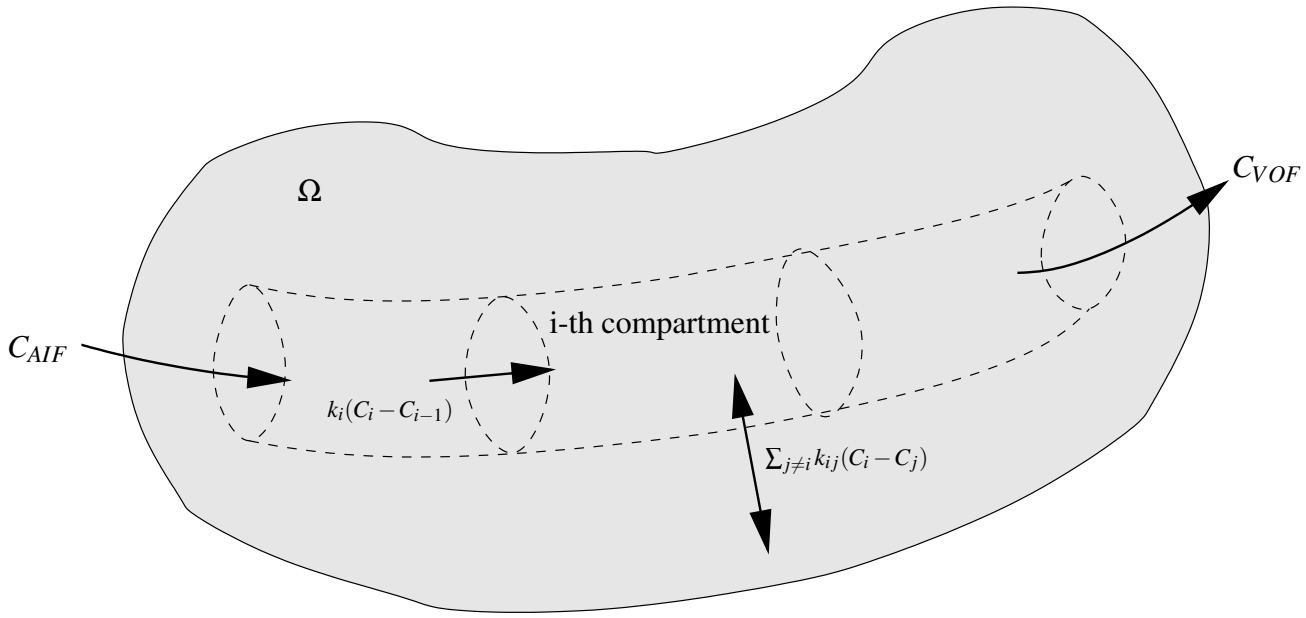

Figure 2.7: Convection-diffusion model of perfusion. Transfer constants k_i describe convective movement and coefficients k_{ij} describe diffusive transfer of indicator.

single compartment is the sum of the convective contribution from the upstream compartment and the sum of all diffusive exchange:

$$V_i C_i' = k_i(C_{i-1} - C_i) + \sum_{j \neq i} k_{ij}(C_j - C_i) \tag{2.11}$$

By expanding terms, equation (2.11) can be written as a matrix equation (2.12). V is the diagonal matrix of each compartment's volume, $diag\{V_i\}$. $\vec{C}$ is the column vector of indicator concentrations in each compartment, $\langle C_1, \ldots, C_N \rangle^T$ The matrix b is the diagonal matrix of convective transfer constants, $b = diag\{k_i\}$. Finally, the matrix A represents the remaining arithmetic operations on the diffusive transfer constants, k_{ij}. Under the assumptions of this model, the convective exchange between each compartment in Ω is constant. Therefore, the quantity $(C_i - C_{i-1})$ is simply the arterial input function, C_{AIF} for all i.

$$VC' = AC + bC_{AIF} \tag{2.12}$$

The general solution of this equation gives an expression for the concentration of indicator as a function of time (2.13). The initial concentration of indicator in the region Ω is zero at $t = 0$. This creates an initial condition problem where substituting $C(0) = 0$ yields a particular solution (2.14). The quantity $e^{V^{-1}A(t-s)}V^{-1}b$ is the tissue kernel, written as $K(t-s)$ (2.15).

$$C(t) = C(0)e^{V^{-1}At} + \int_0^t C_{AIF}(s)e^{V^{-1}A(t-s)}V^{-1}bds \tag{2.13}$$

$$C(t) = \int_0^t C_{AIF}(s)e^{V^{-1}A(t-s)}V^{-1}bds \tag{2.14}$$

$$C(t) = \int_0^t C_{AIF}(s)K(t-s)ds \tag{2.15}$$

This form relates the concentration over time as a convolution of the tissue kernel and the arterial input function [KBKS04].

$$C(t) = C_{AIF}(t) * K(t) \equiv \int_0^t C_{AIF}(t) \cdot K(t-s)ds \tag{2.16}$$

The divergence theorem is applied in this form to relate the properties of $K(t)$ to volume, flow, and mean time parameters. According to this theorem, the net flow through the region Ω is the sum of all sources and sinks across its boundary surface:

$$vol(\Omega)C_T(t) = \int_\Omega C(t)dx = \int_{\nu^T\vec{n}<0} F|\nu^T\vec{n}|dx \int_0^t C_{AIF}(s)ds - \int_{\nu^T\vec{n}>0} F|\nu^T\vec{n}|dx \int_0^t C_{VOF}(s)ds \tag{2.17}$$

The quantity $|\nu^T\vec{n}|$ denotes the magnitude of the fluid flow vector ν in the direction of normal vector $\vec{n}$ at a point on the boundary of Ω. At initial time, the concentration of indicator in the volume is zero. Therefore, the venous output of concentration is also zero. For $t = 0$, let $C_{VOF}(t) = 0$. To illustrate the significance of the initial value of $K(0)$, let $C_{AIF}(t) = \delta(t)$, the Dirac delta function. This step function represents the instantaneous arrival of the entire quantity of indicator.

When $C_{AIF}(t) = \delta(t)$ is substituted into equation (2.16), $C(t) = K(t)$. Making these substitutions to equation (2.17) and taking the limit $t \to 0$ yields an expression for the net volumetric flow:

$$F = \frac{\int_{\nu^t n<0} F|\nu^T n|}{vol(\Omega)} = K(0) \tag{2.18}$$

This model assumes $vol(\Omega)$ remains constant and there is no diffusive exchange across the boundary of Ω. Furthermore, the total volume of fluid per unit of tissue is the area-under-curve of $K(t)$:

$$V = \int_0^\infty K(t)dt \tag{2.19}$$

Applying the central volume theorem (2.10), the mean-transit time, $\bar{t}$ is expressed as the ratio of volume to flow:

$$\bar{t} = \int_0^\infty \frac{K(t)}{K(0)}dt \tag{2.20}$$

[KBKS04] provides a complete derivation of these parameters. Beyond the convection-diffusion model, n-diffusive compartment distributed parameter models may improve estimates [KCS03]. In addition to a series of convective tubes, multiple diffusive compartments are identified with individual transfer coefficients. The number of compartments is chosen to account for the degree of heterogeneity of tissues in the volume Ω.

2.1.9 Perfusion Recovery

However, given the nature of our data acquisition, we refrained from trying to obtain absolute quantification of the flow, and volume, due to the complex set of influencing variables of patient physiological parameters, and concentrated on developing the machine classification to perform diagnosis in the relative quantities, and operating primarily on the attributes of the residual function. However, as an alternative to removing tissue by biopsy, medical imaging can reveal information about the amount of new vasculature surrounding a tumor. A diagnostic protocol based on medical

imagery has yet to provide results that are as accurate as biopsy. Some existing techniques which use MRI are significantly less accurate in comparison [HFE+94]. Clearly, signal enhancement-time curves alone are unable to reliably differentiate levels of a tumor severity [FHSK+05].

Absolute perfusion in real units of mL/g/s is difficult to quantify for a number of reasons. A known quantity of contrast agent injected into the bloodstream passes through multiple body parts: the lungs and heart, for instance. These compartments "spread out" the concentration of contrast agent before arriving at the tumor locus. Without knowing precisely how much contrast agent is being delivered to the region-of-interest per unit time, the signal change in the image cannot be related to an absolute quantity [Kis01]. In some cases, with model approximations and specific imaging hardware, an approximate quantity of blood flow can be determined. The FAIRER (Flow-sensitive Alternating Inversion Recovery with an Extra RF Pulse) scanning technique, for example, can approximate the perfusion in lung tissue. Results of this experiment were slightly overestimated due to a ROI selection that may have included pulmonary artery input [MCE00].

Measurement of the concentration of contrast agent as it enters the tumor volume, $C_{AIF}(t)$, is not an easy task. In brain tissue, the superior saggital sinus (SSS) provides a convenient artery to measure $C_{AIF}(t)$ as the bolus of contrast agent arrives. In breast tissue, however, there is not such a distinct structure to measure this function. In these cases, the $C_{AIF}(t)$ is typically approximated by applying a exemplary pharmacokinetic model for the particular contrast agent used [YLL+05].

Delayed arrival of the contrast agent due to dispersal introduces error in the perfusion estimate. In these models, the arrival of contrast agent is represented with a step function, $\delta(t)$. Several strategies attempt to reduce this effect [JHHMS04, WsK+03]. One approach to correct for this is to introduce a time shift, t_D into the convolution (2.21). A linear regression line is computed for data points occurring before t_D: the pre-contrast signal. A second regression line is fit to points following t_D, but before maximal signal enhancement. An optimal choice for t_D maximizes the difference between the slopes of these lines.

$$C_T(t-t_D) = \int_0^t K(t-s)C_{AIF}(t-t_D)ds \tag{2.21}$$

Perfusion estimation error may also result from stenosis, or narrowing of vessels from which the arterial input function is measured. This produces additional delay and dispersion of the contrast agent, such that the measured $C_{AIF}(t)$ is not the true input to the tumor voxels. Block-circulant deconvolution may make the perfusion recovery process independent of dispersal and delay [WsB$^+$03].

Injection rate of the contrast agent also impacts perfusion results. For most image acquisition protocols, a powered injector delivers the bolus at a rapid-constant rate to reduce dispersal effects. Delivering the contrast agent at slower rates may improve perfusion estimates by increasing the signal-to-noise ratio with a coarser temporal resolution [NDDT$^+$03]. Simulation of a slow infusion contrast injection suggests a SNR nearly 5 times greater than fast injection followed by high resolution imaging.

Recirculation effects break the assumptions of the convection-diffusion model. The contrast agent which has left the volume may be reintroduced and registered if the acquisition time is sufficiently long. Instead of using a sampled value for $C_{AIF}(t)$, a gamma-variate curve is often used to represent tracer dilution with recirculation effects [PAK$^+$02].

Once the system identifies the tumor region, we can compute relative levels of perfusion in the tumor classification. The signal enhancement is a convolution of the arterial input function with the tissue response kernel. Commonly, experimental procedure calls for the application of a pharmacokinetic model as the arterial input function. To further improve our perfusion estimates, we will adjust the phase, amplitude, and dilation of the stock pharmacokinetic equation for use in our deconvolution. We may also experiment with other pharmacokinetic equations if required.

The discussion of perfusion analysis results in the convolution equations that are solved computationally and presented in Chapter 3. The goal of this research, however, is to develop a classification framework that can be trained to use this recovered perfusion information to generate a diagnosis of disease-free, benign, or malignant ROIs given a signal enhancement series and a controlled image acquisition protocol. Tools for this classification are discussed next.

2.1.10 Methods for Deconvolution

The "undoing" of a convolution operation is deconvolution. In the general context of signal processing, deconvolution may be used to remove the effects of noise to recover the original function. In the context of perfusion analysis, it is the properties of the kernel more than the properties of the input function that describe the exchange behavior between the vasculature and the extravascular space. There are numerous methods to perform deconvolution. Some methods are more computationally efficient, while other may provide more stable solutions in low signal-to-noise conditions. What is represented in $K(t)$ also includes noise effects. Increased spatial resolution or temporal resolution worsens these effects. Even though the data acquisition in this research focuses on T1-weighted processes, where temporal resolution is low, the contribution of noise in the experimental method deserves some inspection and treatment.

Fourier transform of a signal reveals its spectral components. A signal can be represented as a summation of trigonometric basis functions. In a stationary signal, the spectral components remain constant over time. In a transient signal, these spectral components change over time.

$$\hat{f}(\omega) = \int_{R^n} f(x) e^{-i\omega \cdot x} dx \tag{2.22}$$

$$\hat{f}(x) = \frac{1}{(2\pi)^n} \int_{R^n} \hat{f}(\omega) e^{i\omega \cdot x} d\omega \tag{2.23}$$

The Fourier transform is described by equation (2.22), where ω is in radians. The inverse Fourier transformation is equation (2.23). When applied to a convolution operation, the Fourier transform changes it to multiplication. Let the Fourier transformation of a function f be $\mathcal{F}(f)$. Then:

$$\mathcal{F}(f * g) = \mathcal{F}(f) \cdot \mathcal{F}(g) \tag{2.24}$$

equation (2.24) is the convolution theorem. Deconvolution can then be accomplished by dividing both sides of the equation by the transformed input function. However, if the coefficients are near zero, then the solutions may be very unstable.

Another technique is deconvolution using the Moore-Penrose pseudoinverse. When the continuous convolution equation is discretized, it can be written as a matrix equation, $Ax = b$. Here, x represents the original function being recovered in the deconvolution. A unique solution for x may not exist, but the pseudoinverse of A can be used to determine a unique, least-squares fit. Mechanically, the convolution matrix cannot be inverted, due to the row rank of A being greater than is column rank. A is not a square matrix. Algebraically, there are more linear equations in this system than variables. The Moore-Penrose pseudoinverse can be derived as follows.

$$\begin{aligned} A\hat{x} &= b \\ \Rightarrow A^T A\hat{x} &= A^T b \\ \Rightarrow (A^T A)^{-1} A^T A\hat{x} &= (A^T A)^{-1} A^T b \\ \Rightarrow \hat{x} &= (A^T A)^{-1} A^T b \end{aligned} \tag{2.25}$$

Singular value decomposition can be used to compute the pseudoinverse as well. The singular value decomposition of an ill-conditioned matrix can be useful in computing better solutions. The singular value decomposition of a matrix A is written as the as a matrix product $U\Sigma V^T$. Σ is the $n \times n$ diagonal matrix of the eigenvectors of A ordered in descending magnitude. U and V^T are each orthogonal matrices containing the left-singular and right-singular vectors of A, respectively.

$$A = U\Sigma V^T = \begin{bmatrix} | & & | \\ u_0 & \dots & u_n \\ | & & | \end{bmatrix} \begin{bmatrix} \sigma_0 & & & & \\ & \sigma_1 & & & \\ & & \ddots & & \\ & & & \sigma_{n-1} & \\ & & & & \sigma_n \end{bmatrix} \begin{bmatrix} | & & | \\ v_0 & \dots & v_n \\ | & & | \end{bmatrix}^T \tag{2.26}$$

Observe that the rows and columns of U, Σ, and V form an orthonormal inverse of a product of orthonormal matrices and is the reversed product of their transposes [Lea04]. Substituting the matrix decomposition into (2.25), the least-squares solution is (2.27).

$$\begin{aligned} A\hat{x} &= b \\ \Rightarrow (U\Sigma V^T)\hat{x} &= b \\ \Rightarrow (U\Sigma V^T)^T(U\Sigma V^T)\hat{x} &= (U\Sigma V^T)^T b \\ \Rightarrow \hat{x} &= \left((U\Sigma V^T)^T(U\Sigma V^T)\right)^{-1}(U\Sigma V^T)^T b \\ \Rightarrow \hat{x} &= V\Sigma^{-1}U^T b \end{aligned} \tag{2.27}$$

Written in this form, the system can be easily adjusted to mitigate noise effects to improve the stability of solutions [KWCL04]. The eigenvectors of A, $\sigma_0, \dots, \sigma_n$ are ordered in decreasing magnitude. The magnitude of the condition number can be controlled by choosing some k of the eigenvectors, such that σ_0/σ_k is acceptably small. This method of truncated singular value decomposition (TSVD) eliminates the smallest magnitude eigenvectors, which likely correspond to noise effects.

Choosing the optimal number of eigenvalues to retain can remove noise and preserve signal information. Koh, et.al. developed strategies to choose an optimal truncation value by observing properties of the solution kernel [KWCL04]. Koh uses a Picard plot of the solution norm, $\|x\|$

versus the residual, $\|x - Ab\|$. Using a piecewise regression approach, the optimal truncation value is the point of greatest curvature in the plot.

Maximum likelihood (ML) deconvolution is another approach used to recover perfusion parameters. Fitzgerald develops an iterative method for a non-linear least squares algorithm [FON08]. The algorithm uses a grid search to optimize parameters for the reconstruction of the residue function. In cases of low signal-to-noise, this algorithm demonstrates better perfusion parameter estimation for simulated MRI data.

In the extreme case where only the output function is available, and neither function f or g is known, blind deconvolution techniques may be able produce a result. In blind deconvolution, a general constraint on either f or g is used to iteratively generate estimates of both operands [AD88]. For example, in $c = f * g$ it may be known *a priori* that both f and g are non-negative. Blind deconvolution begins with an initial estimate for f. The function g is obtained by a conventional deconvolution method, such as Fourier division. The known constraints for g are then applied to this estimate. Any negative values are replaced with zero. A new estimation for f is computed, and the process is repeated until two non-negative functions with a convolution of c are determined. As the iterations proceed, computational problems may arise as matrices formed by f and g may become ill-posed for deconvolution.

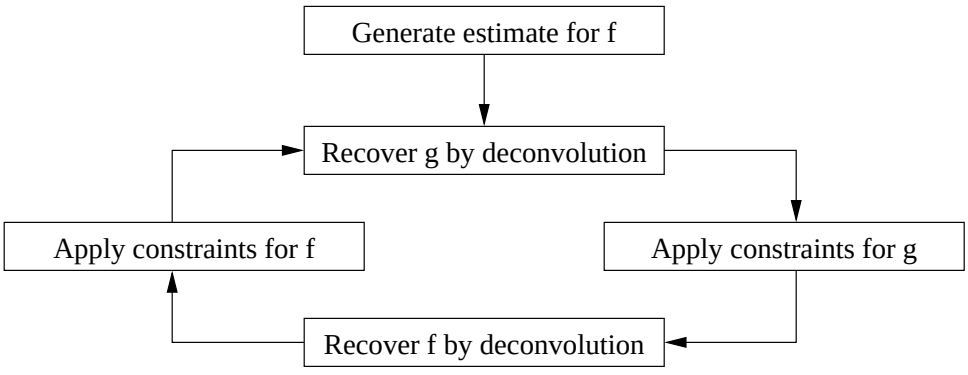

Figure 2.8: Generalized blind deconvolution

2.2 Machine Learning Algorithms in Diagnosis

Having reviewed the relevant perfusion models, the next step in constructing a computer-aided diagnosis framework is selecting a set of machine learning algorithms. Machine learning algorithms, given a training data set, can make objective predictions for a set of new test cases. First, the medical data must be conditioned into a representation that the learning algorithm can use. Time series data, an image raster, or a combination of both can be formulated into a feature vector of arbitrary length. There are several examples of medical data that can be used for learning. With electrocardiogram data, a time series can be "learned" for detecting ischaemic heart disease. [KKG+99] applied several different algorithms including back-propagation neural networks, a naive Bayes classifier, and a K-nearest neighbor algorithm. Breast ultrasound represented in digital imagery has been studied for the purpose of characterizing microcalcifications [CWMC03]. In this study, windowed subimages where generated from the original ultrasound image, and trained in a support vector machine as a form of texture analysis. A similar approach for digital mammograms was studied in [ENYW+02].

There is no shortage of machine learning algorithms. Artificial neural networks (ANN) are a popular choice for concept learning and are based on their biological counterpart: connected neurons in the human brain. The network is a set of interconnected neurons, with an input layer, output layer, and hidden layer. A pattern is applied to the input layer, and an output pattern is generated through the weights and interconnections of the hidden layer(s).

Twellmann, et. al., have developed a back-propagation neural network for signal enhancement data as input that generates a tissue classification [TLN05]. Their classification framework is designed as a model-free tissue classification model versus model-based methods, such as the 3TP method. The classifier developed in their paper is called the Adaptive Tissue Classification Neural Network (ATCNN), which compared its accuracy to that of the 3TP method. The methodology developed in Chapter 3 subsumes this by considering both the signal enhancement data and the characteristics of perfusion recovery, to improve classification accuracy. Hoyt, et. al., also use a neural network approach for analyzing contrast-enhanced ultrasound imagery [HWU+10]. The

machine learning method was used to evaluate the effect of an anti-angiogenic drug, Bevacimuzab, on breast cancer in mice.

2.2.1 Support Vector Machine Classification

The support vector machine approach to data classification has demonstrated excellent performance in the realm of two-class pattern recognition in a feature space of arbitrary dimensions. The support vector machine is based on statistical learning theory, which has found application in a number of application domains: optical character recognition, face matching, and text classification. In the medical realm, it has been applied largely to ultrasound imagery and texture-based classification [CWMC03, CWM$^+$03, KP03, WSM93, LKP$^+$08b, KT10]. In this section, we present a basic overview of the Support Vector Machine approach, where thorough treatments of the statistical learning foundations can be found in [Bur98, MM08].

A set of vectors $x_i \in \mathbb{R}^n, i = 0 \ldots l$ each have a classification, $y_i \in -1, 1$, where $-1, 1$ represent the two classes in consideration. As with any learning algorithm, a subset of these vectors and their respective classes (x_i, y_i) are chosen as a training set. The support vector machine determines a mapping $x_i \rightarrow f(x_i, \alpha)$, where α are the learned parameters of the mapping. The analog for α in a back propagation neural-network classifier would be the node weights at each perceptron. The SVM can then be tested using the remaining examples.

For a given set of parameters α, the average error rate on the training set, or the *empirical risk* is defined by (2.28), where l is the size of the training data set.

$$R_{emp}(\alpha) = \frac{1}{2l} \sum_{i=0}^{l} |y_i - f(x_i, \alpha)| \tag{2.28}$$

The "loss" of each classified training example, $\frac{1}{2}|y_i - f(x_i, \alpha)|$ can be proved to be bounded by (2.29) [Bur98].

$$R(\alpha) = R_{emp}(\alpha) + \sqrt{\frac{(h(log(2l/h) + 1) - log(\eta/4))}{h}} \tag{2.29}$$

The variable η denotes the probability of the loss function taking the value 0 correctly trained, or $1-\eta$ incorrectly trained. The variable h is a non-negative integer called the Vapnik-Chervonenkis (VC) dimension. The VC dimension represents the "capacity" of the learned machine, $f(x,\alpha)$. The capacity of a trained classifier is a measure of how it generalized beyond the training data set. A classifier with too high a capacity will over fit the training data, and fail to generalize to new testing vectors. Conversely, a classifier with too low a capacity fails to properly discriminate new testing vectors.

Assuming the training set is independently drawn, and individually separated, the SVM algorithm proceeds to find a set of partitioning hyperplanes that maximize the margin of separation between the classes being sought (Figure 2.9).

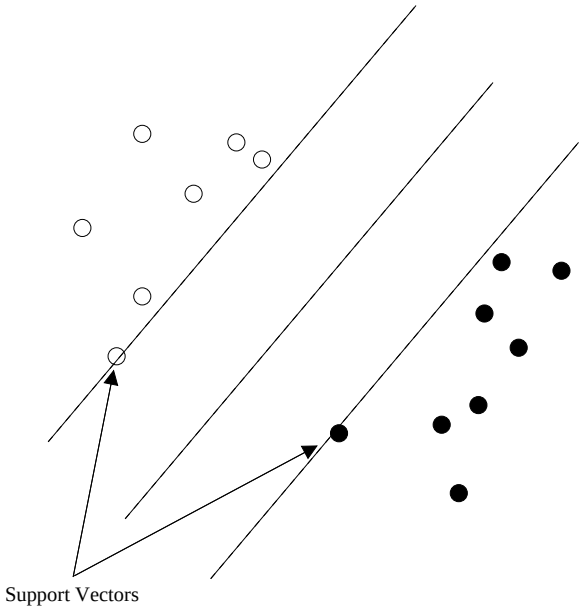

Figure 2.9: Support Vector Machine maximizes the margin of separation between feature classes

The determination of the separating function(s), are obtained by numerical optimization of these error bounds, even though faster greedy algorithms have been developed to converge faster and more stably [VM03]. In this research, we have selected the sequential minimal optimization for quadratic decomposition implemented in `libsvm` [CL01] and utilized by the PyML code [BH07].

The general applicability of the SVM method has popularized its use in the medical domain for screening and diagnosis. The most common application is for texture or morphology of suspicious image features. In [ENYW$^+$02], SVM is applied to ultrasound imagery to detect microcalcifications. The feature vectors, x_i, in this are $M \times M$ windows surrounding each input pixel.

2.2.2 K-Means Cluster Analysis

While not a learning method, the k-means clustering algorithm is often employed in image segmentation. The algorithm divides vectors in a multi-dimensional feature space into k-fixed groups in an attempt to minimize the intercluster variance. When used for image segmentation, each pixel becomes a feature vector with coordinates and the pixel value $I(x,y)$ as the feature vector components: $\langle x, y, I(x,y) \rangle$. The k-clusters are randomly initialized to members in the feature space. The algorithm iteratively reassigns the cluster groups by closest cluster centroid. Each cluster centroid is recomputed and the recombination is run again. The algorithm converges and halts when no cluster-reassignments occur (Algorithm 2).

Algorithm 2 K-Means Algorithm

```
Given a set of feature vectors, V, and k initial clusters,
Assign each vector to a random cluster.
repeat
  Let ClustersChanged = false
  for all v ∈ V do
    for all C_j ≠ memberOf(v) do
      if distance(v, centroid(C_j)) < distance(v, centroid(memberOf(v))) then
        Assign v to C_j
        Let ClustersChanged = true
      end if
    end for
  end for
until ClustersChanged = false
```

The algorithm is implemented using the Euclidean distance measure. Many improvements can be made to the direct K-means method. A more efficient implementation using KD-Trees can be used to improve the execution time [ARS98]. The number of clusters, k, is selected arbitrarily by

the user, and increased or decreased to improve the "fit" of clusters to the vector space. The initial assignment of the clusters to vectors in the set may also affect the segmentation result. Methods have been developed for refining the selection of k and the cluster initialization [BF98]. For image segmentation, the result can be evaluated in terms of differentiation and connectivity. The quality of clustering for higher-dimensional data can be measured by the total error function equation (2.30)

$$E = \sum_{j=1}^{k} \sum_{i_l \in C_j} |i_l - w_j|^2 \tag{2.30}$$

K-means segmentation has been applied to medical image segmentation of the MR and X-ray images of the brain and other tissues [SWDH02, PP99, PXP00]. Chapter 3 describes how K-means clustering is used as a segmentation aid for identifying training data in the MR imagery.

Subspace clustering and projected clustering are other techniques that may be relevant. With higher-dimensional, manually composed feature vectors, the meaning of a distance measure becomes less intuitive. For instance, a feature vector that contains a spatial coordinate, pixel intensities over time, and other meta-information might cluster very strongly spatially, but fail to cluster on intensity. Subspace clustering is one way to discover clusters on arbitrary subspaces with the n dimensions of the feature vector. In this way, a feature vector may be part of 0 or more subspace clusters. This is may discover clusters on seemingly unrelated components of the vector. An extension of subspace clustering is projected clustering, where each vector is assigned to one, and only one cluster which may exist on a subspace of the vector space [AWY$^+$99].

We employ some of these extensions to the classification framework during segmentation for training sets to improve the training data in our research.

In the next chapter, we formulate a classification framework that exploits both the spatial and temporal nature of the medical data, and employ these learning algorithms to evaluate their efficacy for breast tumor detection and diagnosis.

Chapter 3

Materials and Methods

3.1 Data Sets

The 29 MRI breast studies with the Volume Imaging for BReast AssessmeNT (VIBRANT) acquisition protocol were retrospectively examined in this research with all patient identifiers eliminated. The dynamics image series began with a full resolution, pre-contrast T1-weighted acquisition. The contrast agent, 20cc of gadolinium with saline, was administered by powered injection. The time series consisted of five post-contrast imaging sequences spaced 90 seconds apart. Each acquisition produced a three-dimensional volume with a spatial resolution of 512 pixels left-to-right and 512 pixels anterior to posterior. The Z-axis resolution of each data set varied from 84 to 156 pixels from superior-to-inferior. The MRI sequence parameters for the contrast enhanced series were $TR = 5.2 - 5.7$, $TE = 2.1$, and $TI = 17$. Each pixel had 8-bit grayscale intensity values from $0 - 255$. Each data set included other non-contrast enhanced images with MRI parameters set for different tissue contrasts. The pathology for each data set was retrieved and correlated to lesions in the images. Additional X-ray and ultrasound studies also helped correlate the pathology analysis to a ground-truth answer.

During initial examination of the cases, the GE CADStream software package visualized enhancement data in the VIBRANT series. CADStream offered several useful views of the imagery for enhancement analysis. The default 3-panel view in CADStream showed axial, saggital, and coronal projections. A Dynamics overlay could be toggled, showing enhancement areas based on the 3-Time Points method (3TP), though the colorations were non-standard. Yellow, blue, and red were used instead of the red, green, and blue color scheme described in the typical method. When compared to a manually implemented 3TP result, CADStream also appears to be filtering single-points of enhancement, possibly with a median or averaging post-processing filter.

In CADStream, the Curves dialog was used to graph the enhancement series at a mouse-selected point. Here, the first and second enhancements are displayed as percent change. An-

other Surface dialog showed a three dimensional surface generated at a mouse-selected feature. It is unknown what criteria was used to segment and generate this surface, though the segmentation roughly included areas of similar enhancement. The 3D surface view computed the selected feature's dimension and volume, as well as an ellipsoidal buffer region for a surgical biopsy estimation.

Within CADStream, there was no way to access the enhancement information in a computable form. Also, to ensure that there were no post-processing effects from CADStream, the images were exported into DICOM format. The imagery was saved in separate folders for each study. The DICOM format contained imaging parameters in the metadata sections to identify the time, spatial orientation, and dimension. LSU Health Science technicians aided in the anonymization of each study such that only image data and a generic identifier remain prior to analysis (e.g., Data Set 01, Data Set 02, etc.).

OsiriX was a medical imaging application that supports the DICOM format [RSR04]. OsiriX could generate 2 and 3 dimensional views, maximum-intensity projections and surface renderings. Each anonymized DICOM directory was imported into OsiriX as a separate study. OsiriX failed to group the images with respect to time. However, OsiriX did support exporting the DICOM images to an alternate image format.

In order to work directly with the data, the DICOM had to be exported into a binary format that was both lossless and conveniently indexed. Uncompressed TIFF served as an intermediate format to extract the DICOM images. The export was performed with OsiriX on a dual-quad core Apple MacPro 3GHz, with 4GB of system memory. After each axial slice was converted to a 512×512 TIFF image, an optimal storage format for the images was constructed. Each data set was stacked and reordered such that the time series for a single voxel was stored contiguously (Figure 3.1). A Python script collected the time series at each voxel, iterating over each image in row-major order. Each slice was visited in an inferior-to-superior order. The final binary files for computation were referred to as stacked enhancement volumes.

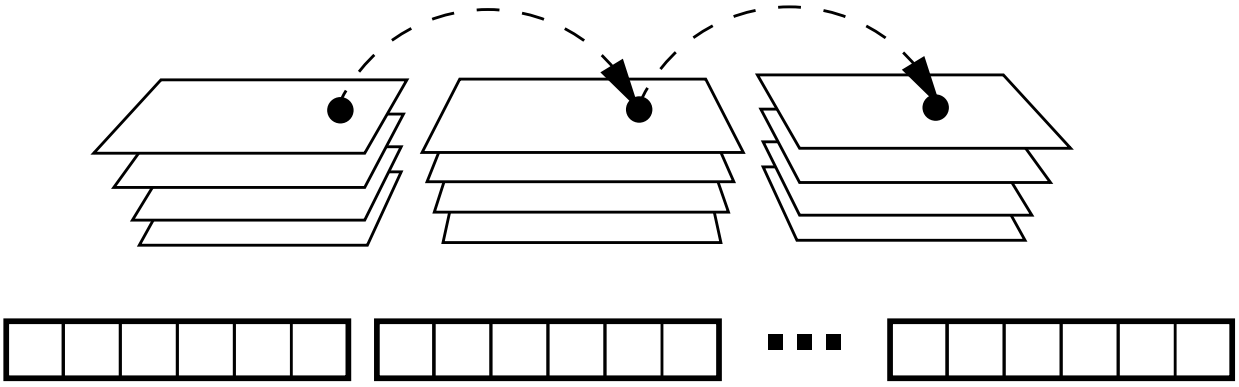

Figure 3.1: Contiguous-time stacked enhancement files

A Python class library was written to interact with these stacked enhancement volumes intuitively and efficient for computation. With the Python NumPy libraries, these volume files could be memory-mapped and treated as a 4-dimensional array, (x, y, z, t). The library access functions ignored any zero-valued time series as a background pixel, to eliminate unnecessary processing operations.

3.2 Computation

3.2.1 3-Time Points Implementation

To produce a baseline result, the 3TP algorithm was implemented in Python as defined in section 2.1.6. Times t_0, t_1, and t_3 were chosen to be the pre-contrast, maximum enhancement, and post-contrast times, respectively. For each time series, the color category was computed, and the 8-bit intensity was assigned based on the algorithm. The algorithm was also written in Python and the Python Imaging Libraries (PIL) provided routines to compose the images in memory and write them to disk.

The output according to the algorithm produced pixels too dark to distinguish visually. As a visual aid, the intensity for each pixel was brightened according to the mapping function, $x' = f(x) = \lfloor 255 \cdot \sqrt{x/255} \rfloor$, where x is the red, green, or blue intensity in the range $[0, 255]$ from the 3TP algorithm. The behavior of this function boosts the lower range of the brightness values

higher, but keeps the relative brightnesses in each color channel consistent. Figure 3.2 shows this brightening function.

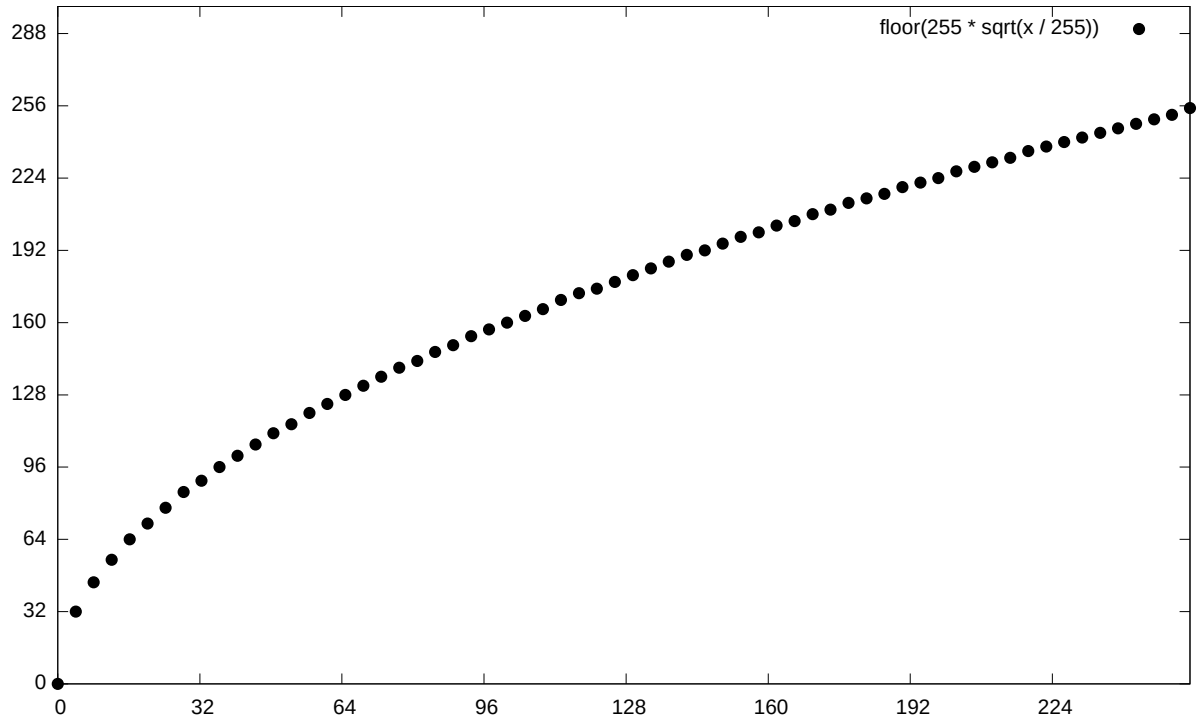

Figure 3.2: Intensity brightening applied to 3TP output

The output was overlaid on the original pre-contrast image for reference, and each slice was output in 32-bit PNG for visualization. Several strategies were employed to reduce the presence of single pixel outliers and more closely match the output of the CADStream Dynamics View. A 3×3 median filter was applied to the output prior to compositing. This filter was applied independently to each channel (red, green, and blue) then to all channels simultaneously. Larger median filters, 5×5 and 7×7, were also attempted on representative data sets.

A second strategy to reduce outliers was to select a threshold for the initial enhancement interval. Below that threshold in the range $[0, 255]$, the 3TP value was set to zero. In this case, only the pixels with a high enough initial enhancement period would remain. Finally, these noise reduction

methods were combined to both remove single-pixel outliers and truncate pixels with low initial enhancement values.

The results of the implemented 3TP were compared against the output from CADStream and qualitative judgments were made. It was not possible to quantify the CADStream 3TP results as the data in the Dynamics view could not be exported. The closest resultant match was used as a baseline result against the classification result described next.

3.2.2 Matrix Construction

In the equation $C(t) = C_{AIF}(t) * K(t)$, the perfusion is characterized by the tissue kernel, $K(t)$. Properties of this function such as maximum value, area-under-curve, and initial value can be related to perfusion quantities as described by equations (2.18), (2.19), and (2.20). To solve for $K(t)$ using the VIBRANT imagery, a discrete convolution equation was constructed and solved with matrix deconvolution methods.

Let $C(t)$, $C_{AIF}(t)$, and $K(t)$ be functions of the discrete time points $t_0, t_1, \cdots, t_{n-1}$. Convolution is performed by first using the elements of one of the operands in descending order. The operands are overlaid such that the first elements overlap. The corresponding elements are multiplied and summed to produce one value of the convolved signal. The reversed operand is shifted to the right by one step, and the corresponding elements are again multiplied and summed. This operation can be represented by a matrix equation $Ax = b$, where A is a Toeplitz matrix created by shifting $C_{AIF}(t)$ to the right for each descending row. The elements of A, x, and b are defined in terms of the values of $C_{AIF}(t)$, $K(t)$, and $C(t)$ in equation (3.1).

$$
\begin{aligned}
a_{ij} &= \begin{cases} C_{AIF}(t_{i-j}), & 0 \le j \le i \\ 0, & i < j \end{cases} \\
x_i &= K(t_i) \\
b_i &= C(t_i) \\
i &= 0, \cdots, n-1
\end{aligned} \tag{3.1}
$$

The tissue concentration function $C(t)$ is the enhancement time series from the imagery at each pixel. Note that the matrix b has $2n-1$ elements, resulting from the dimensions of the matrix multiplication product, where a $2n-1 \times n$ matrix multiplied by a $n \times 1$ matrix yields gives a $2n-1 \times 1$ matrix. To achieve the desired length, we interpolated additional points between actual points as necessary. The justification for increasing the number of the observed points is due to the coarse temporal resolution of the VIBRANT acquisition sequence. As a future improvement, different interpolation methods might be employed to introduce smoothness, such as a cubic spline polynomial fit. Equation (3.2) illustrates this matrix construction.

$$
\begin{bmatrix}
C_{AIF}(t_0) & 0 & \dots & & 0 \\
C_{AIF}(t_1) & C_{AIF}(t_0) & 0 & \dots & \vdots \\
\vdots & & & & \\
C_{AIF}(t_{n-1}) & C_{AIF}(t_{n-2}) & \ddots & C_{AIF}(t_1) & C_{AIF}(t_0) \\
 & & & & \vdots \\
\vdots & \dots & 0 & C_{AIF}(t_{n-1}) & C_{AIF}(t_{n-2}) \\
0 & & \dots & 0 & C_{AIF}(t_{n-1})
\end{bmatrix}
\begin{bmatrix} K(t_0) \\ K(t_1) \\ \\ \vdots \\ \\ K(t_{n-2}) \\ K(t_{n-1}) \end{bmatrix}
=
\begin{bmatrix} C(t_0) \\ \\ \vdots \\ \\ C(t_{n-1}) \end{bmatrix}
\tag{3.2}
$$

3.2.3 Input Function Selection

One very important detail is to define the selection of the input function $C_{AIF}(t)$. Two principle methods for input function selection are used based on physiological appropriateness [STCD04]. In some parts of the body, the tissue pixels are adjacent to a large artery. Here, the enhancement due to the contrast passage is clearly seen. In cerebral perfusion recovery, for example, the superior saggital sinus is located as shown on Figure 3.3. The circled feature in the middle image is at the maximum enhancement point and can be sampled.

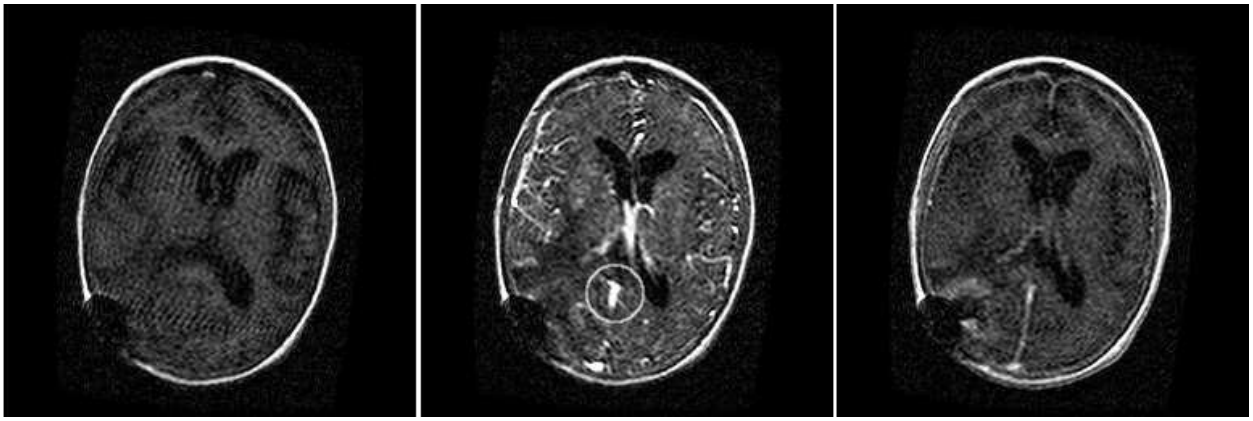

Figure 3.3: Superior saggital sinus in cerebral perfusion

However, the vasculature in the breast tissue does not usually present such a convenient feature. Though blood vessels were visible in the VIBRANT imagery, the spatial resolution of the structure was typically not large enough to guarantee that there was no contribution from surrounding tissue enhancement. In this case, a pharmacokinetic model that approximates the typical bolus arrival pattern was used. A gamma distribution has shown to correlate well with the measured AIFs [WsB$^+$03]. In equation (3.3), $C_{AIF}(t)$ was approximated by the gamma distribution with parameters α and β.

$$C_{AIF}(t) = \begin{cases} 0, & t < t_0 \\ C_0(t-t_0)^{\alpha} e^{(t-t_0)/\beta}, & t > t_0 \end{cases} \tag{3.3}$$

The parameters α and β control the shape and scale of the gamma distribution. Typical values used with a pharmacokinetic system are $\alpha = 3.0$, and $\beta = 1.5$. This equation was sampled at time intervals corresponding to the time series.

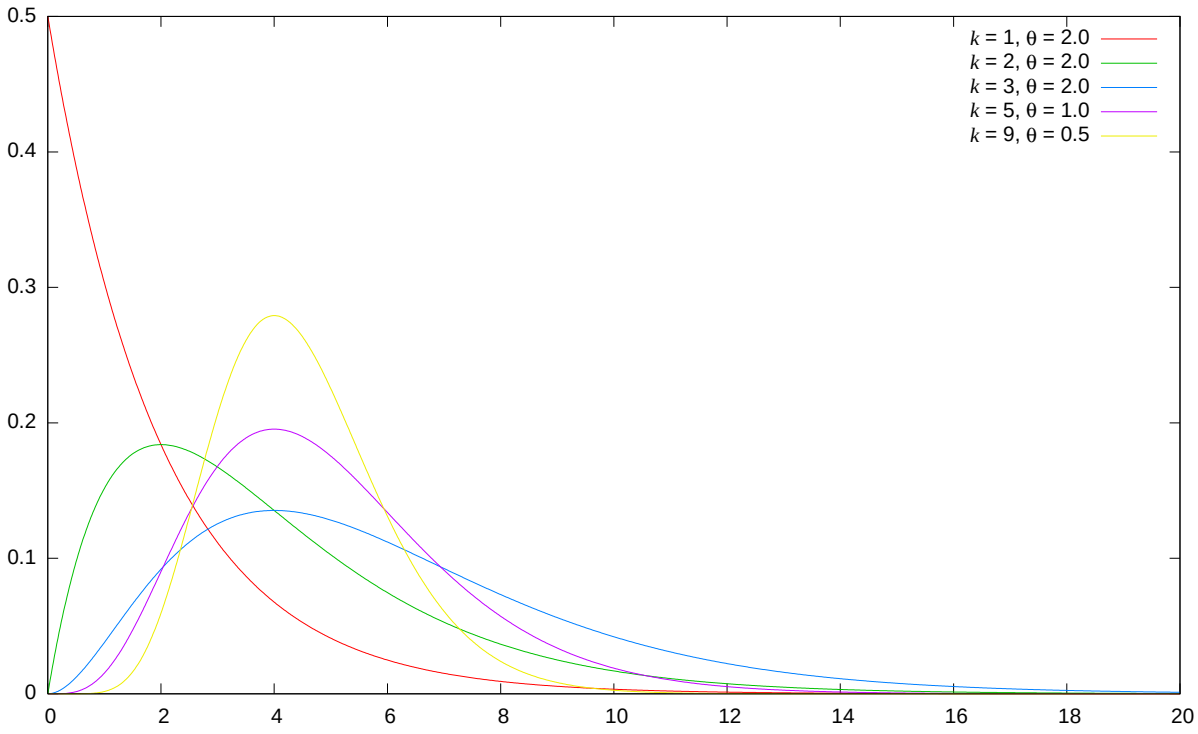

Figure 3.4: Gamma variate function for various values of α, β

Both methods for $C_{AIF}(t)$ selection were attempted. For a sampled $C_{AIF}(t)$ a representative data set was identified. Several blood vessels were located in the breast tissue, sampled, and

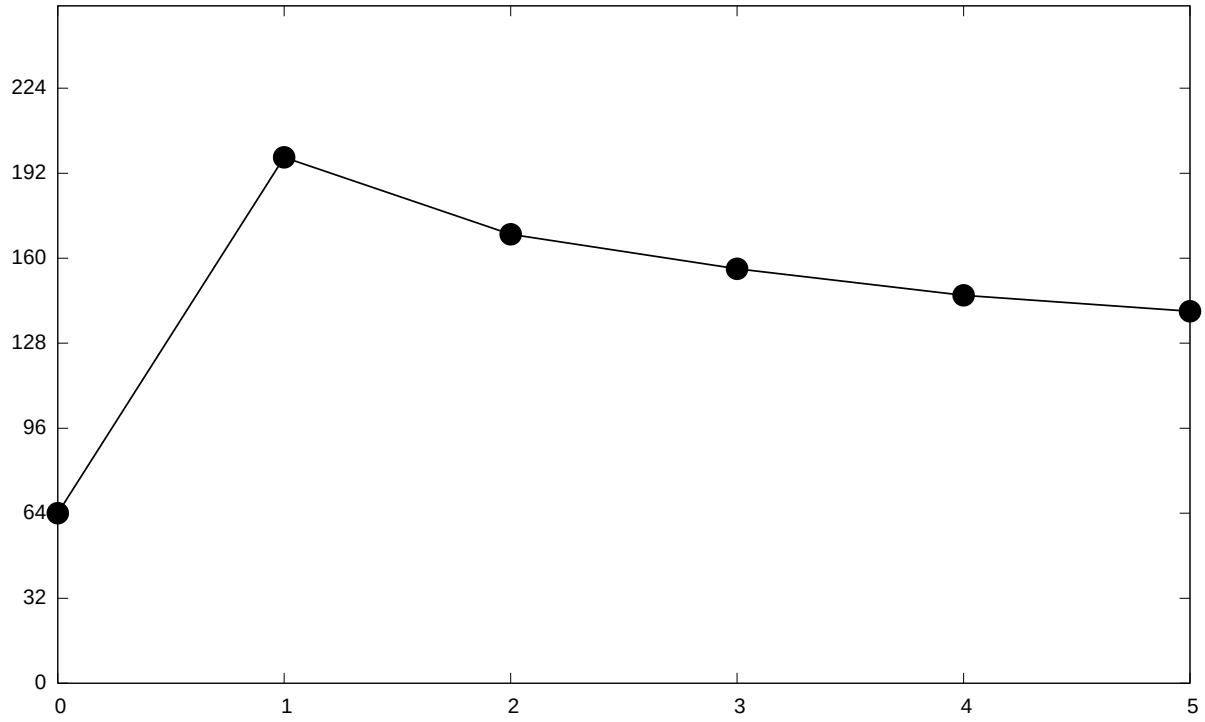

Figure 3.5: $C_{AIF}(t)$ sampled from several blood vessels

averaged. Figure 3.5 shows several enhancement series from manually located blood vessels, and their averaged curve.

3.2.4 Discrete Deconvolution

This system defined $2n \times 1$ equations in n variables. The solution of this system was found by several methods of deconvolution. For both methods of input function selection, three forms of deconvolution were applied: Moore-Penrose pseudoinversion, singular-value decomposition, and blind deconvolution. The kernels from the deconvolution of from each method were compared, as well as how each method affected the classification results.

The Python language and its accompanying scientific library SciPy were selected as the computational platform because:

1. The file-to-memory mapping capabilities allow convenient abstractions to be written to handle the stacked enhancement volumes.

2. The SciPy library provided a set of matrix structures in a "Matlab-like" syntax.

3. Parallel runs were also easily distributable, even though Python does not provide true processor-level parallelism within a single execution instance.

4. General image input, manipulation, and output is an important feature of the language and libraries.

5. Machine learning toolkits such as PyML provided the K-nearest neighbor and SVM for the classification stage.

The perfusion recovery code operated directly on the time series of the stacked enhancement volumes. Each volume was loaded into memory, and mapped into a 4-dimensional array. Iterating over each location, the 8-bit integer series was deconvolved into a 32-bit double-precision representation and stored in another binary file, referred to as a stacked kernel volume. The stacked kernel volume was arranged with the same indexing as the stacked enhancement volume.

Several optimizations improved the computational processing time for each volume. Each time series can be deconvolved independently of any other point in the volume. There were no data dependencies in the processing loop and therefore the operation could be "completely" parallelized. Utilizing up to eight cores on the MacPro workstation, each stacked enhancement volume was divided into a $length/N$ segment for processing. Each segment was sent to a separate core to deconvolve. Then all core results were merged at completion. Due to the global Python interpreter lock, a multi-process implementation was selected versus a multi-threaded implementation. The processing was computationally bound, not communications bound, since the entire stacked volume file could be contained in real memory.

Additionally, the blind deconvolution was applied iteratively to estimate $K(t)$ and $C_{AIF}(t)$. $C_{AIF}(t)$ was initialized as a gamma distribution with $\alpha = 3.0$ and $\beta = 1.5$. $C_{AIF}(t)$ and $K(t)$ were designed to be non-negative functions. Because of this requirement, any negative values encountered in any iteration were replaced by zero.

3.3 Classification

3.3.1 Training Sets

Training data sets were generated for a pre-selected subset of the images using the original MRI sets, their pathology information, and any other information available. From the disease-positive data sets, a single axial slice image was selected if it contained a pathology-confirmed tumor. From the disease-negative cases, a slice containing an enhancing feature, such as a cyst, was selected if present. Otherwise, an axial slice passing through the nipple was chosen.

In the selected disease-positive images, the tumor pixels needed to be distinguished from the non-tumor pixels with a tumor mask. This tumor mask was generated manually at first, and then later with the aid of a segmentation algorithm. A program was created to assist both methods of generating the tumor mask, PerfusionView. PerfusionView was written using C++ and open source Qt libraries its the graphical interface and image handling. The software supported loading the stacked enhancement volume files and displaying axial, saggital, and coronal projections of the three dimensional volume presented on Figure 3.6.

PerfusionView allowed panning through the volume on each axis, selecting a specific time point, or centering on a mouse-selected point in all three views. 3TP results could be toggled as an overlay to the imagery. PerfusionView functioned in three modes: inspection, mark-up, and classification. In inspection mode, the time series for the point underneath the cursor was shown on a graph, along with its 3TP classification. This capability is very similar to the CADStream Curves display. In mark-up mode, the operator could use a variable size drawing tool to paint the tumor mask onto the image. This mask could be saved as a 1-bit image to be used in the classification training phase. Having generated a bitmask for a selected axial slice, the classification mode launched the training and testing phases of the classification.

3.3.2 Segmentation-Aided Marking

The manual marking facility in PerfusionView was one method for indicating disease-positive pixels in the imagery. Correlating the study pathology with the imagery and the radiologist report

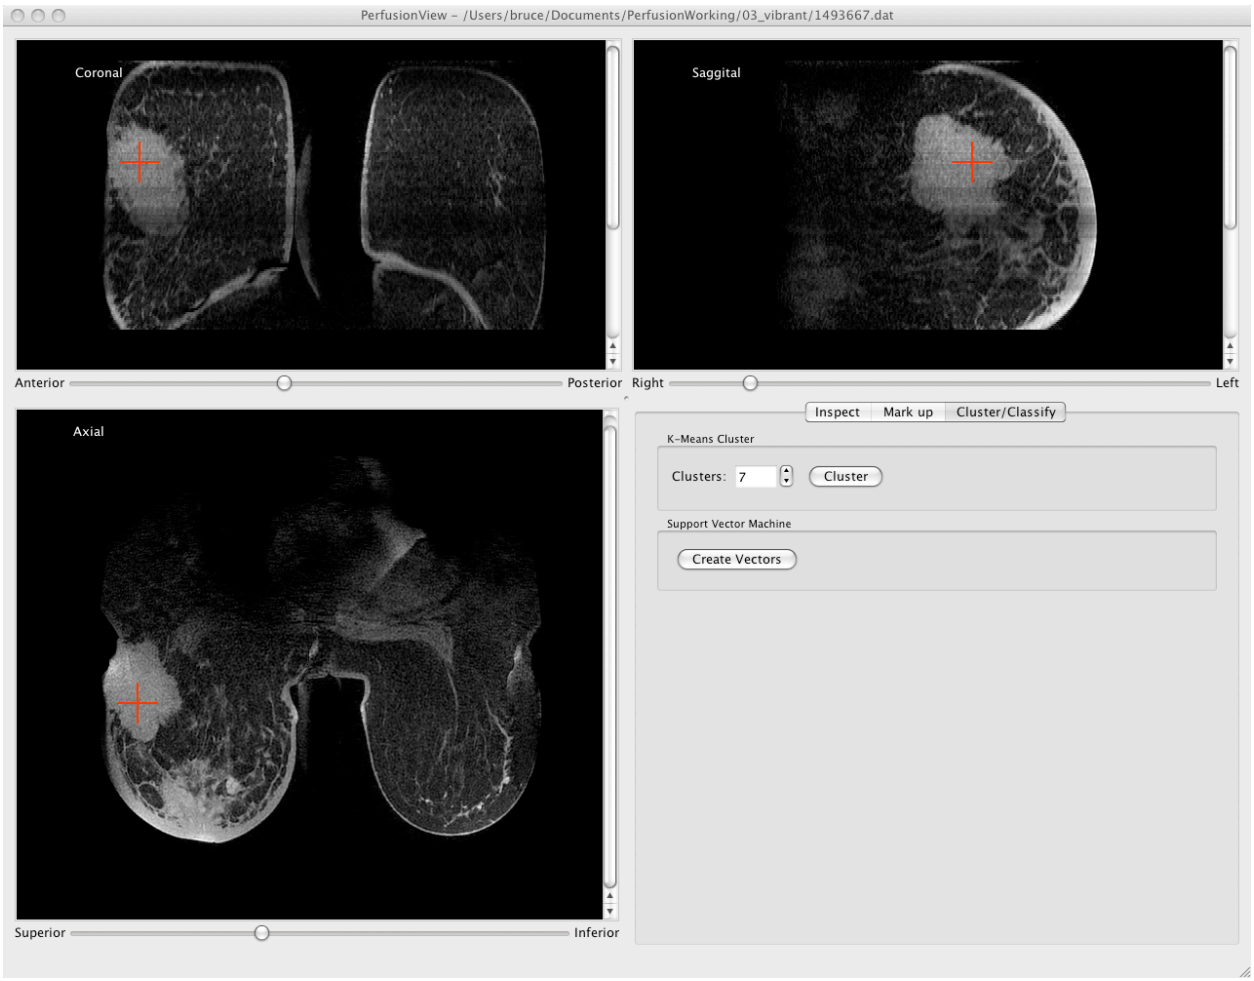

Figure 3.6: PerfusionView with a stacked enhancement volume

provides a rough location and morphological description of a lesion. Radiologist descriptions were defined in "clock face" position. From a coronal view the nipple was the clock center, and the tumor was reported in terms of clock axes, one through twelve o'clock. Tumors located behind the nipple, or retroareolar tumors, may not have had a clock face position. These rough descriptors, aided in locating obvious tumors. However, there was always the possibility of a significant marking error due to poor boundary resolution in the original image. An improved method for assisting in the marking of the training sets was devised using the previously described K-means clustering.

For each training data set, each kernel function in the stacked kernel volume was added as an element in the population to be clustered. The algorithm was initialized to $k = 7$ random clusters

as described in section 2.2.2. The selection of the number of clusters was determined through experimentation. This choice of number of clusters provided the best balance between intercluster variance while not over-fitting the data. The Euclidean distance served as the metric in the K-means algorithm:

$$\begin{aligned} P &= (p_1, p_2, \ldots, p_n) \\ Q &= (q_1, q_2, \ldots, q_n) \\ d(P,Q) &= \sqrt{\sum_{i=1}^{n} (p_i - q_i)^2} \end{aligned} \tag{3.4}$$

The clustering algorithm was implemented in C++, also using open source Qt classes for image input and output. From PerfusionView, the deconvolution routine was first executed for an entire stacked image volume. The clustering function was executed next with a selectable number of clusters. Upon completion, the clustering program generated an output image per axial slice in the same image resolution as the original VIBRANT images. A representative feature vector was generated for each cluster by averaging all the final cluster members. To visualize, each cluster feature vector was ordered by maximum component value, and a color map was applied for each cluster.

This image is reviewed to include or exclude groups of pixels that belonged to the same cluster and were spatially adjacent. This alternate method of marking tumor pixels was expected to have an impact on both training and testing the classifiers used in the final stage shown on Figure 3.9. Both the segmentation-assisted marking method and the completely manual method were used to generate testing data for the classifier training and testing phase. The effects on accuracy in the classification phase were compared with receiver operating characteristic (ROC) analysis .

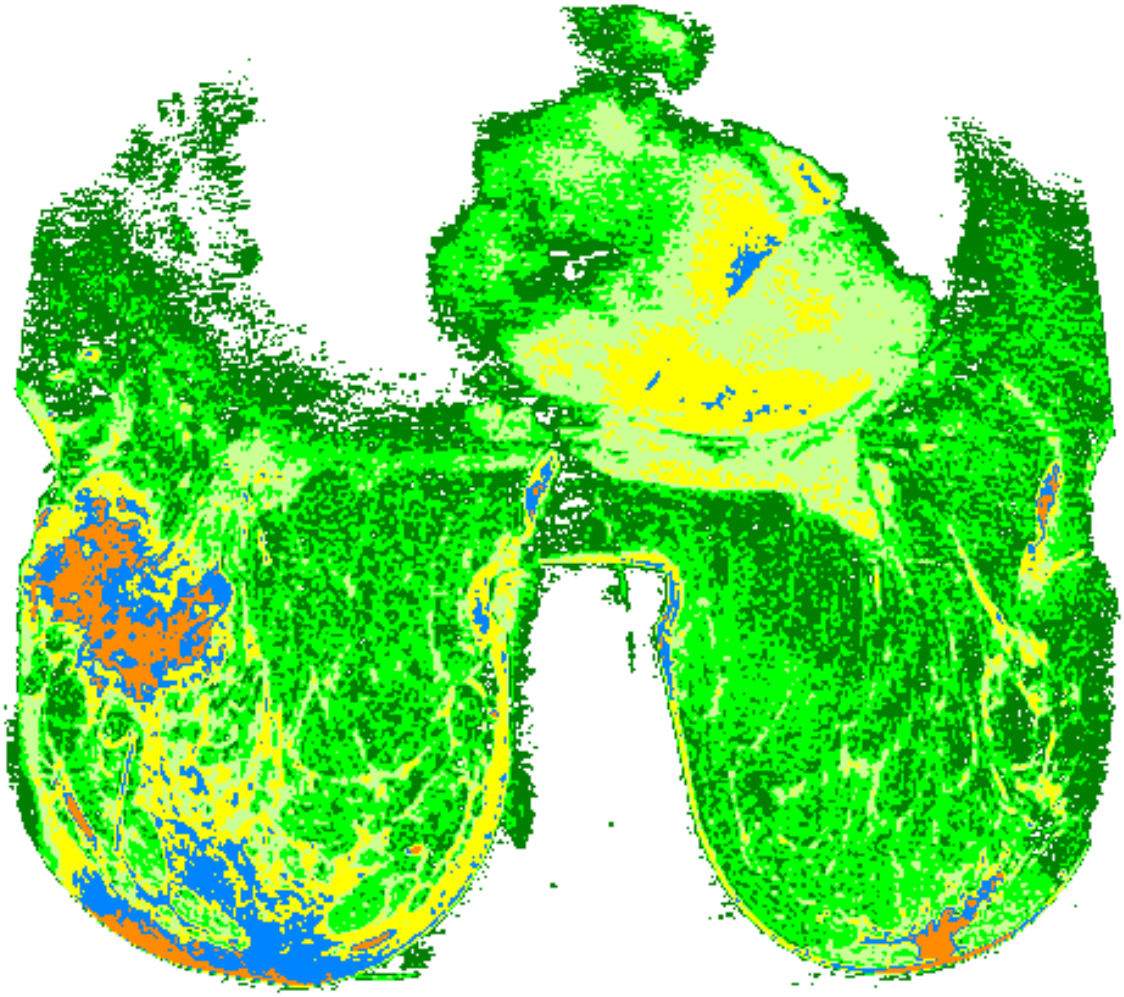

Figure 3.7: K-Means cluster result

3.3.3 Training and Testing

The tumor mask and the corresponding data points generated in both testing and training sets was used for classification. Using PerfusionView, the tumor mask was combined with the stacked kernel volume. An index appended to the end of the tumor mask file name was used to reference the axial slice marked for classification. The 1-bit mask had a value of 1 for a pixel marked

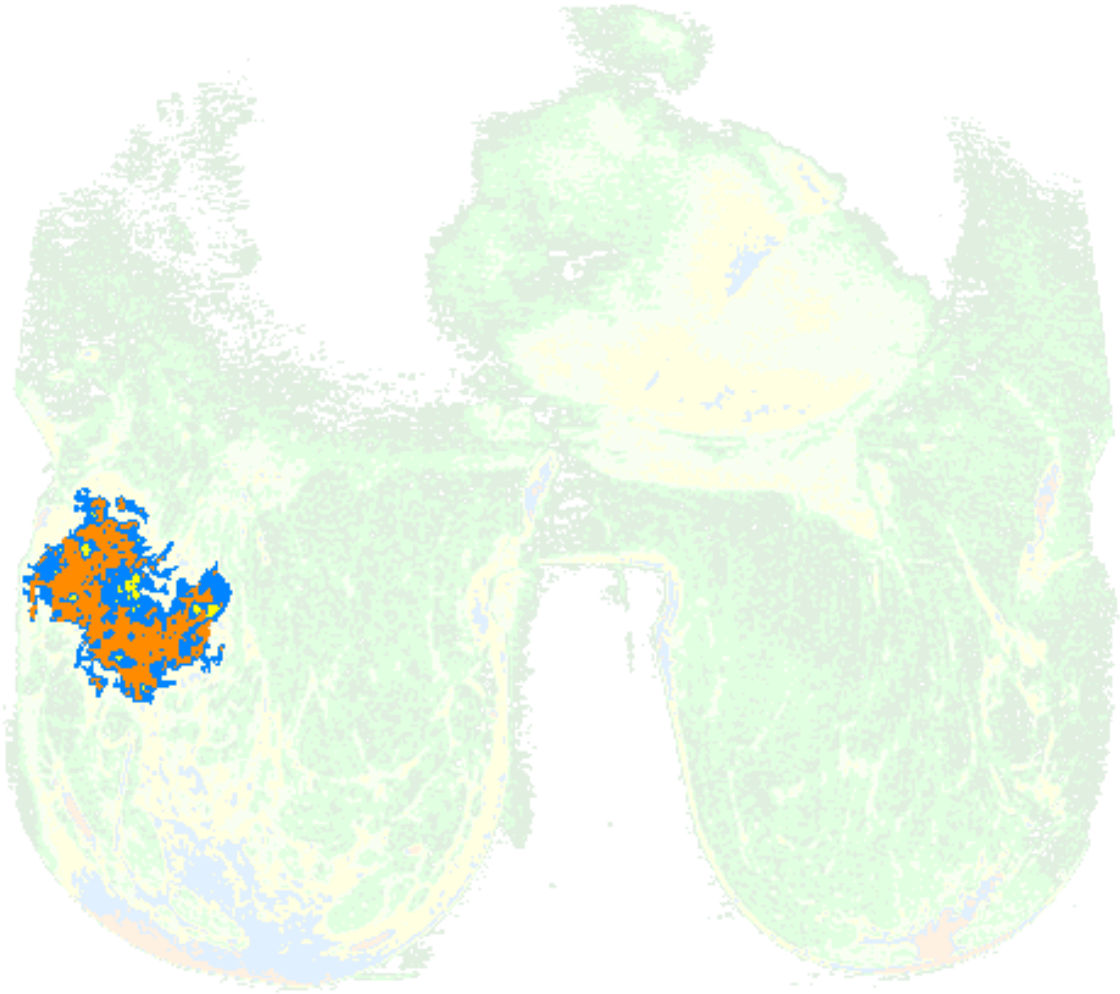

Figure 3.8: Cluster result shown with tumor clusters isolated

disease-positive and 0 for a pixel marked disease-negative. Pixels in the stacked volume that have a zero-valued kernel function were excluded from both the training and testing sets.

PerfusionView generated a feature vector for each point that does not have a zero-valued kernel function. The values of $K(t)$ and the derived quantities from the central volume theorem were added to this feature vector. Because the relative flow was the same as the kernel function at $t = 0$,

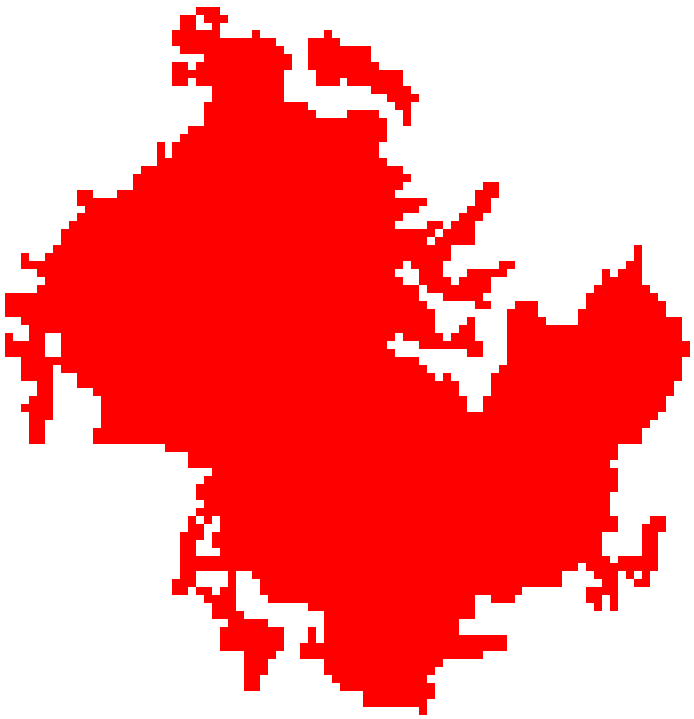

Figure 3.9: Selected clusters added to tumor mask

it was omitted from the feature vector. Total relative blood volume computed from the area under curve of the kernel function was included in the feature vector. The trapezoidal rule was used to compute the area-under the curve. The mean particle transit time was the ratio of volume to flow. The flow is divided by the initial value, $K(0)$, and incorporated in the feature vector.

In addition to the estimated relative perfusion parameters, the feature vector includes other properties of the kernel. The first and second order forward differences characterize the shape of the kernel. The rates of change reflected in these two functions present characteristics that are independent of individual signal or kernel values.

$$k'_i = K(i+1) - K(n) : i = 0, 1, \ldots, max - 1 \tag{3.5}$$

$$k''_i = K(i+1) - 2K(i) + K(i-1) : i = 1, 2, \ldots, max - 1 \tag{3.6}$$

The final formulation concatenates all of these values into one fixed-length feature vector.

$$\left\langle x_0, x_1, x_2, x_3, x_4, k'_0, k'_1, k'_2, k'_3, k'_4, k''_1, k''_2, k''_3, k''_4, \arg\max_t(K(t)), \int K(t) \right\rangle \tag{3.7}$$

Different training set sizes were chosen for the SVM initialization phase. The training set size effects the accuracy, specificity, and sensitivity of all of the classifiers tested. In the trails, percentages of the pixels in axial slices ranged from 10% to 100% in 10% increments. In each trial, with a percentage of training data to be used, the perfusion software randomly selected examples from the population and wrote them to an external file. Each example contained the user-marked classification, pixel coordinate, a data set index, and the fixed-length feature vector. A Python program loaded the training examples file for each trial. The program then instantiated a new classifier using the SVM implementation from the PyML machine learning library. The examples were passed to the SVM for training. The program wrote the trained SVM in an external file.

The classification trials were organized such that examples were included from the same stacked volume as well as from other stacked volumes. The first set of classification trials chose elements from the axial slice where the SVM training examples were selected. The Python code chose classification examples that were mutually exclusive with the original training set. The second set of classification trials used examples from other axial slices in the same volume. The alternate slices

were chosen at regularly-spaced intervals and marked manually and with segmentation assistance. The third set of trials used an SVM model and classified examples from other stacked volumes.

3.3.4 Challenges

Current models for the concentration behavior originated from indicator-diluter theories of Zierler, et.al. [Zie62, Zie63]. These theories made the following assumptions:

1. Every particle of fluid entering a voxel compartment leaves in a finite time.
2. There were no effects due to recirculation.
3. The volume of the fluid compartment remained constant.
4. The flow rate through the fluid compartment was constant.

Most of these assumptions were not realistic in the actual physiology. The circulatory system is closed, so recirculation effects were present [KBKS04]. Also, these tissues were non-rigid, and tumor boundaries were not easily defined. Finally the circulatory system did not maintain constant flow; rather, the heart pumps blood through the tissue in a periodic fashion.

These assumptions about the contrast agent bolus as it travels through the body could have created significant errors in the estimation of flow rate per unit tissue volume. Also, the manner in which the imagery was obtained introduced variable amounts of noise in the observed data. Examining the condition number of the discretized convolution matrix showed the deconvolution solution to be possibly noise-sensitive.

Calamante [CGC02] discussed making assumptions to the arterial input function, $C_{AIF}(t)$ in computing blood flow. Specifically, any deviation of actual bolus arrival behavior from the estimated model affected the amount of perfusion computed. From the injection site into a vein, the bolus of contrast agent had to travel to and from the heart, through the lungs, back into the heart, and then distributed into the breast. Such a path will likely bias the shape of the arterial input function. Adjustments to the parameters of the gamma variate function were made to compensate for these physiological effects.

Chapter 4

Results

4.1 Organization

The experiment consisted of four phases:

1. ground truth specification,
2. creation of controlled results with 3TP,
3. recovery of perfusion parameters by deconvolution,
4. training and of testing of perfusion feature vectors.

Prior to the experiment, the anonymized MRI data was reformatted for efficient access as described in Chapter 3. Both the radiology report and pathology report gave some information for locating the tumor in each study. A custom software tool, PerfusionView, provided visualization and creation tools for manual tracing of each tumor mask. An assisted tracing method was also used to refine the mask creation. The K-means clustering algorithm was run on the prepared data for different numbers of clusters. Images of the clustering results were used to select the optimum number of clusters for the algorithm. The cluster results were used to refine the boundary shape of each tumor mask. Both methods of specifying ground truth were compared in the classification phase of the experiment.

In the second phase of the experiment, the 3-Time Points (3TP) method was run on the MRI data. Different time points were selected in each run until a satisfactory result was produced. Image post-processing methods were tested on the 3TP results to approximate the results of the default radiologist tool, CAD Stream. The performance of the subsequent classification were compared with these results as an experimental control.

The third phase of the experiment investigated different methods for perfusion recovery. This phase of the experiment explored three different ways to select the arterial input function. Not all

methods yielded usable results. Usable perfusion results were carried forward to the fourth phase of the experiment.

In the final phase of the experiment, different feature vector constructions were tested in a support vector machine classifier. The feature vectors contained elements from the unmodified experimental data as well as derived elements from the recovered perfusion data. The performance of each feature vector construction was evaluated using 10-fold stratified cross validation and compared to the 3TP results when used as a classification method.

4.2 Data Preparation

The twenty-nine radiological studies were initially identified for use in this research. Each study was a complete image set of MRI scan types, including contrast-enhanced MRI time series. Two studies were duplicates of patients for pre- and post-mastectomy comparison. The two post-mastectomy studies were discarded. These twenty-nine studies represented twenty-seven individual patients imaged with the VIBRANT protocol. Fifteen of the studies were positive for ductal carcinoma and four of the studies were negative for ductal carcinoma.

In the eight remaining studies, the pathology was unavailable or incomplete, and hence were removed from the pool. Only seven of the fifteen studies that were positive for ductal carcinoma had one main tumor with a well-defined boundary. These studies were selected as training data to reduce manual segmentation error. The remaining eight ductal carcinoma studies either lacked a single, well-defined tumor boundary, or had multiple foci that would require a prohibitive amount of effort to manually mark every tumor. Appendix 5.6 presents a complete list of all studies along with a non-personally identifying data set number.

In preparation for analysis, all of the enhancement data were extracted into the time series-volume format previously described. The following analysis operated on a single axial slice containing the lesion's largest cross-section. For the disease-negative studies, an axial slice through the nipple center line was selected. Each axial slice was visualized with PerfusionView, using the multiple-view windows to locate the enhancement behavior.

In addition to isolating the axial slice in each study, the coronal plane that partitioned the breast tissue from the chest was recorded. Though there are automated algorithms for determining this plane, PerfusionView was used to manually inspect each study and isolate the breast tissue anterior of the sternum.

Figure 4.2 shows a single axial image from each disease-positive study. The white bounding box identifies the location of the tumor in each study; these are drawn manually. The brightness level of these images were increased for improved illustration and do not indicate the intensities of the 16-bit raw signal. Signal intensity within the DICOM image is expressed in Hounsfield units (HU). Average brightness within each lesion was 1213 HU with a standard deviation of $\sigma = 443$ HU. The lesion texture varies between each study. The tumors in studies D15 and D16 have a smoother appearance, while D1, D11, and D12 have a mottled appearance. The number of pixels that belong to each lesion varied from 400 to 5039 pixels. The average lesion pixel count was 1737 pixels. Figure 4.3 shows each tumor area-of-interest in greater detail.

Figure 4.1 shows an axial image from the four disease-negative studies. In some of these images, there are lesions that have been identified as non-cancerous masses, such as cysts. The brightness of these images has also been adjusted for better illustration, and do not reflect the brightness in the raw data. In Study D3, the left breast has a small feature that shows enhancement, but with relatively few other landmarks in the breast. Study D14 has a fundamentally different appearance in both breasts, with a higher initial enhancement value.

4.3 Specification of Ground Truth

The radiologist notes on the tumor morphology and the pathology report aided in establishing the ground truth for each study. The radiologist report only described general location (left or right breast, clock face position) and qualitative morphology. The pathology findings only described presence or absence of cancer by tissue biopsy. Subsequently, the tumor in each study needed to be traced with a manually drawn mask. A pixel was “cancer-positive” if it intersected this mask.

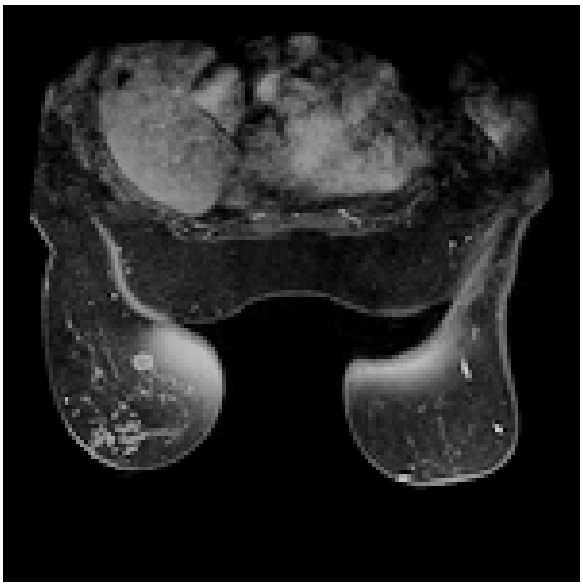

(a) Study D3

(b) Study D4

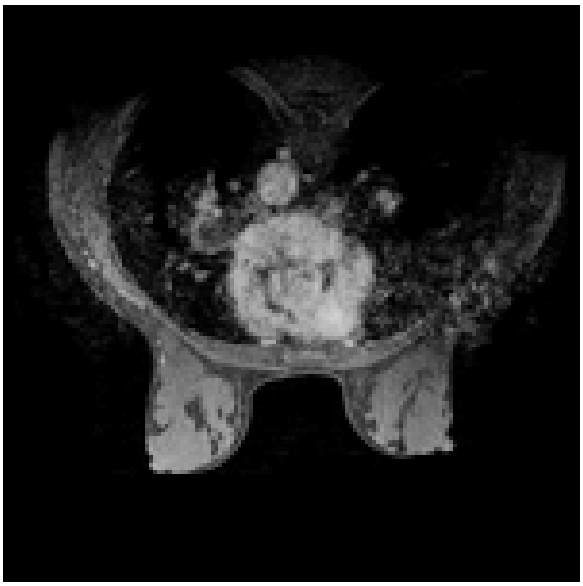

(c) Study D14

(d) Study D27

Figure 4.1: Disease-negative studies, pre-enhancement axial imagery

Conversely, a pixel was "cancer-free" if the mask excluded it. Figure 4.4 shows the drawn tumor masks.

This method of marking "cancer-positive" pixels depended upon the precision of the mask-drawing tools. Therefore, a second method of specifying training data used a cluster-assisted segmentation method.

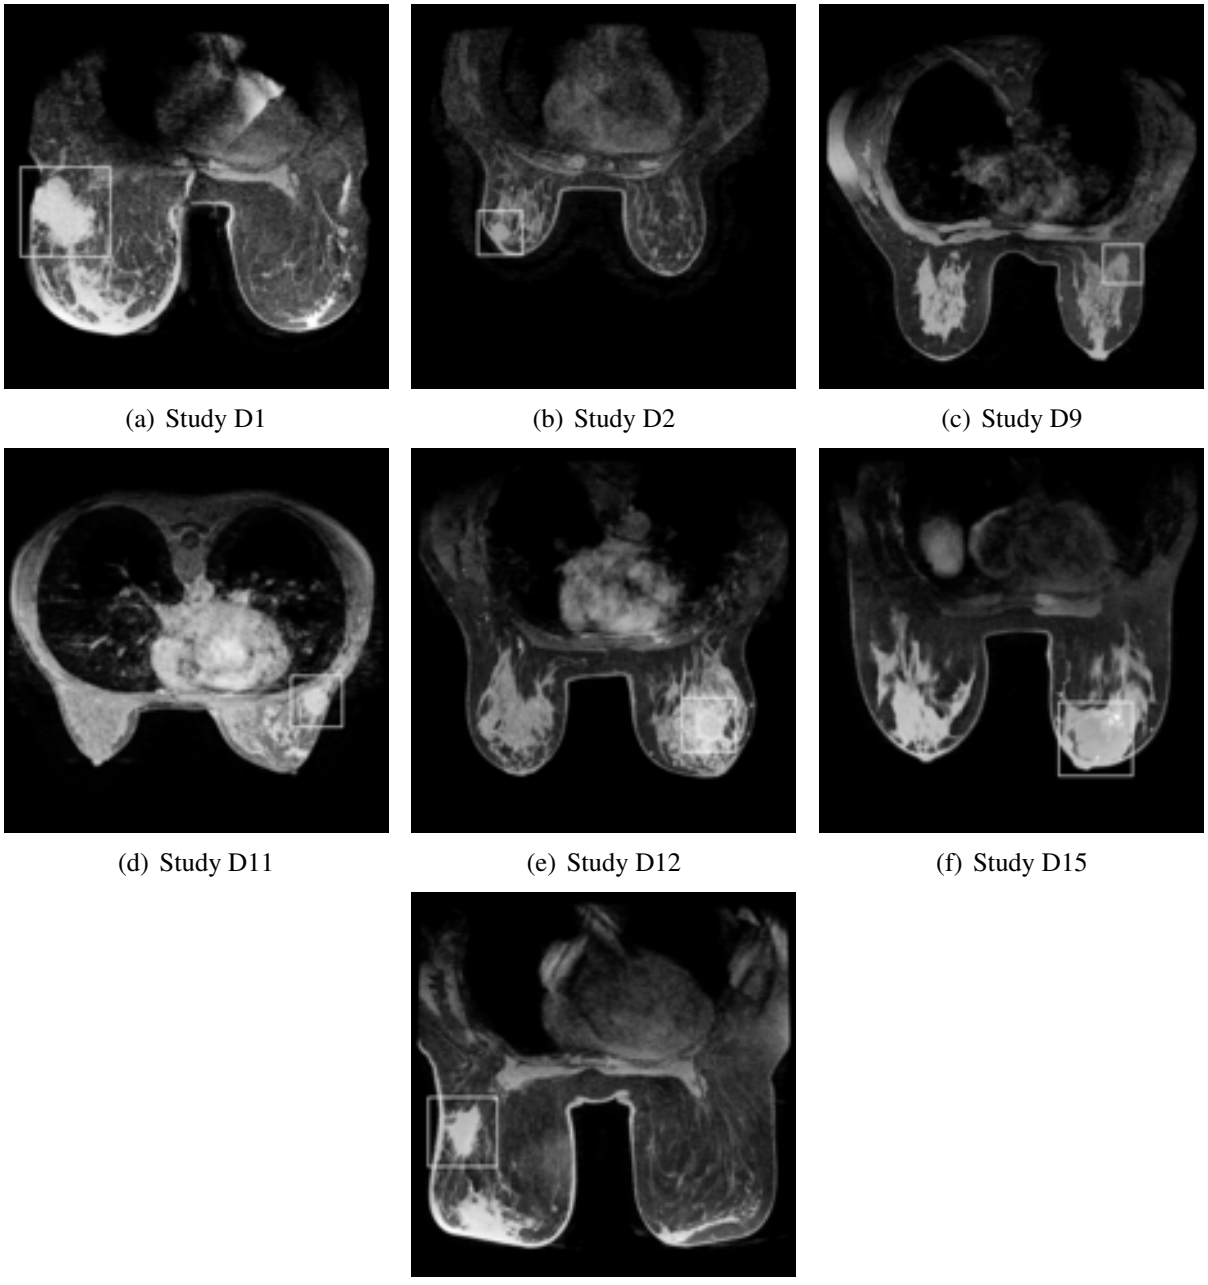

(a) Study D1 (b) Study D2 (c) Study D9

(d) Study D11 (e) Study D12 (f) Study D15

(g) Study D16

Figure 4.2: Disease-positive studies, pre-enhancement axial imagery

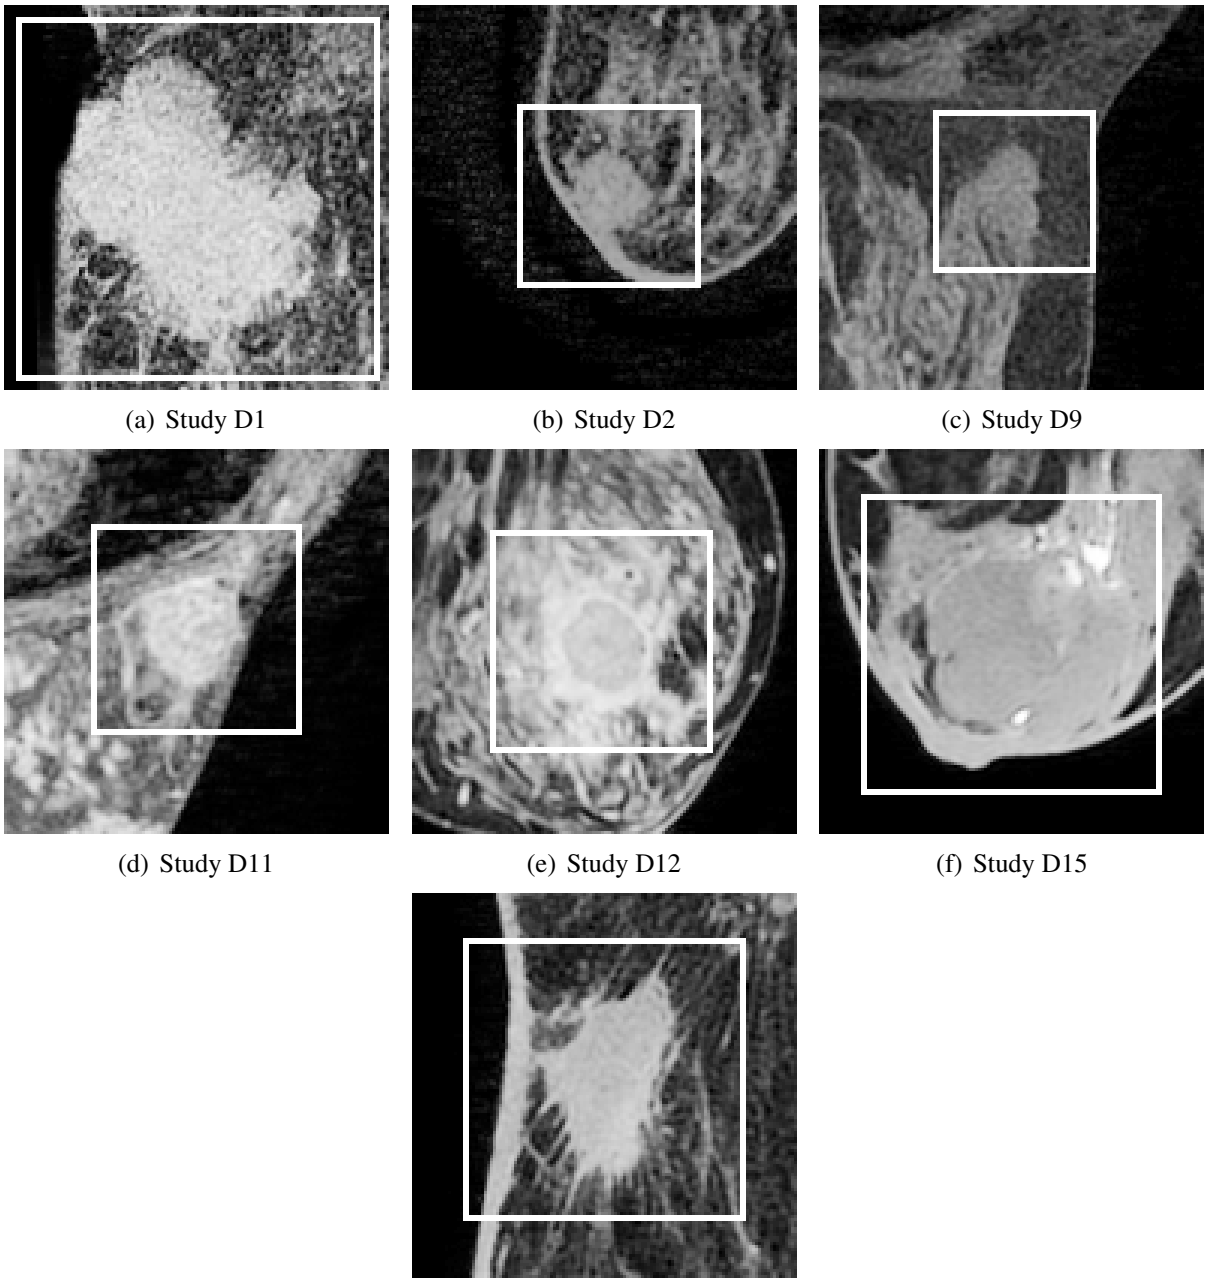

(a) Study D1 (b) Study D2 (c) Study D9

(d) Study D11 (e) Study D12 (f) Study D15

(g) Study D16

Figure 4.3: Disease-positive studies, cropped to AOI

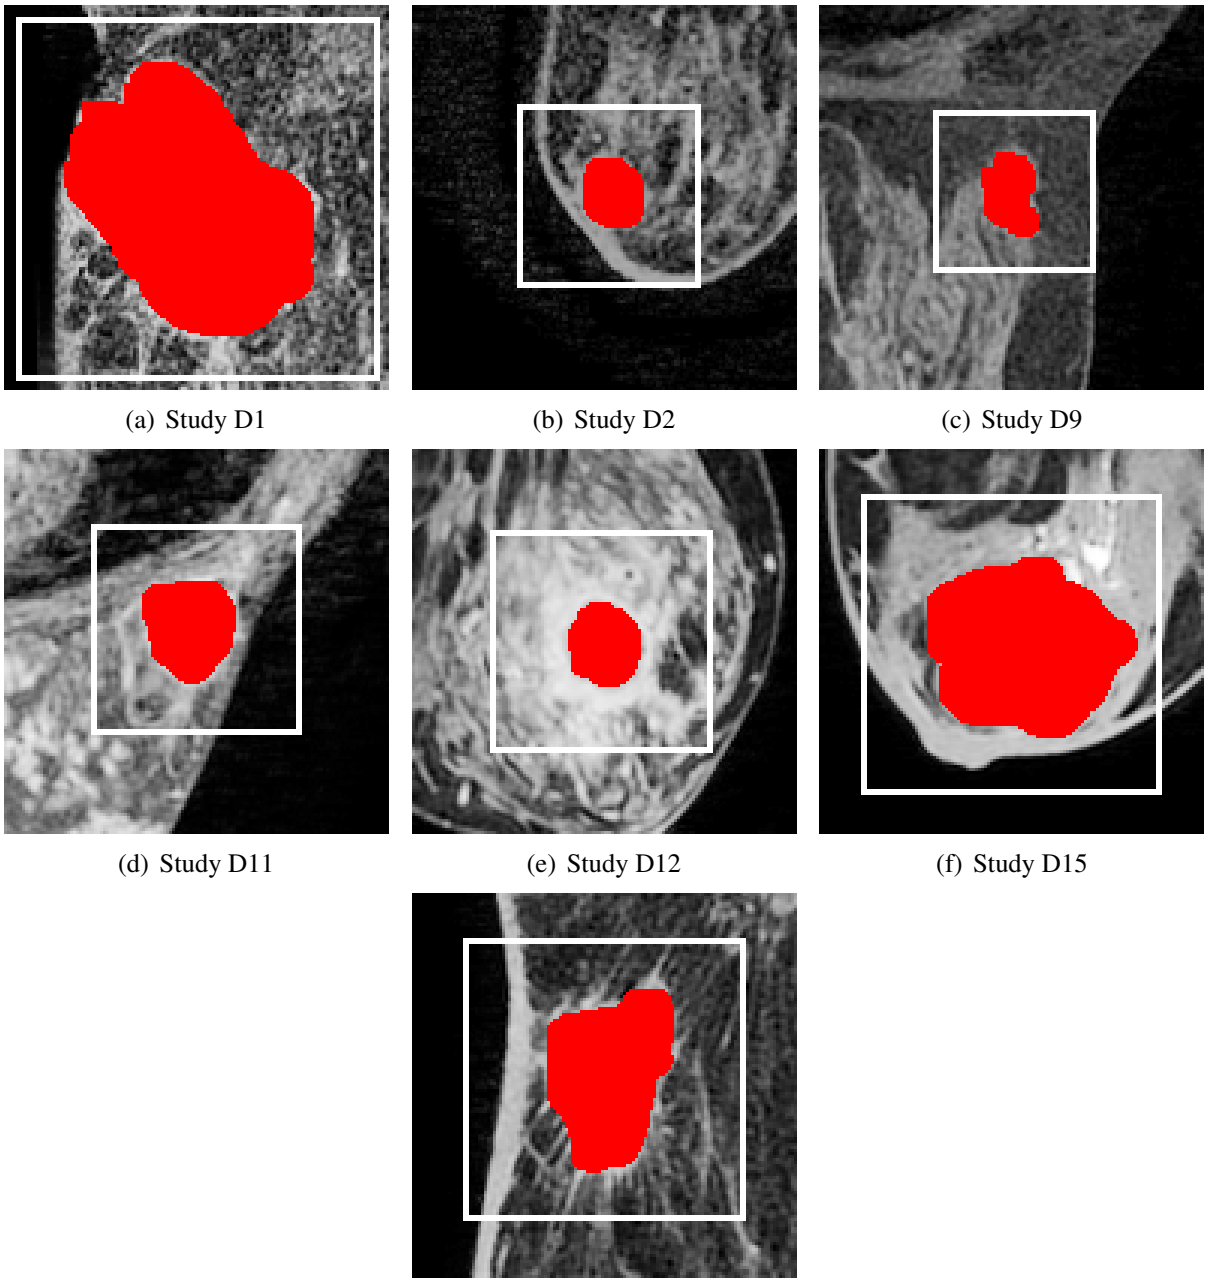

(a) Study D1 (b) Study D2 (c) Study D9

(d) Study D11 (e) Study D12 (f) Study D15

(g) Study D16

Figure 4.4: Manually-drawn tumor masks are in red

Each pixel time-series $C(t)$ was constructed as a 5-element feature vector. The pixels from each study were separately clustered with a K-means clustering algorithm. Through sequential search, $K = 7$ was found to generate enough distinct feature classes. Each feature class was sorted by maximum element value and assigned a color in a green-to-yellow-to-red color palette to create the images presented on Figure 4.5.

Note that this was not a classification step, but a segmentation aid that was exploited in the classifier training later in the dissertation. The pixels in the region did not always possess the maximum $C(t)$ value as shown in study D16. The clustering results were used to select regions that corresponded to the tumor area in the disease-positive studies. They were used as an alternate tumor mask for classifier training. Figure 4.6 shows these refined tumor masks. Additionally, a breast tissue mask was also created manually to eliminate the chest and the background pixels.

The results for each subsequent phase of the experiment compare the effect of using these two methods of ground truth selection.

4.4 3-Time Points Results

The 3-Time Points (3TP) method generates a three category classification used to indicate different tiers of microvascular permeability. Of the three curve categories, Type II and Type III enhancement curves indicate a "cancer-positive" condition, per pixel. Type I curves and non-enhancing curves indicate a "cancer-free" condition, per pixel.

The canonical implementation of the 3TP method generates images with Type I, II, and III curves as color-coded pixels. Any pixels that do not match one of the enhancement categories are left uncolored (black). This instance of the 3TP method uses the first, second, and fifth time points in each enhancement curve. Figure 4.7 shows these results.

The initial 3TP results, without post-processing, contained noise and artifacts throughout the image. Because the intensity of green, or blue pixels is computed from the percent change between the first and second time points, many background pixels show significant enhancement.

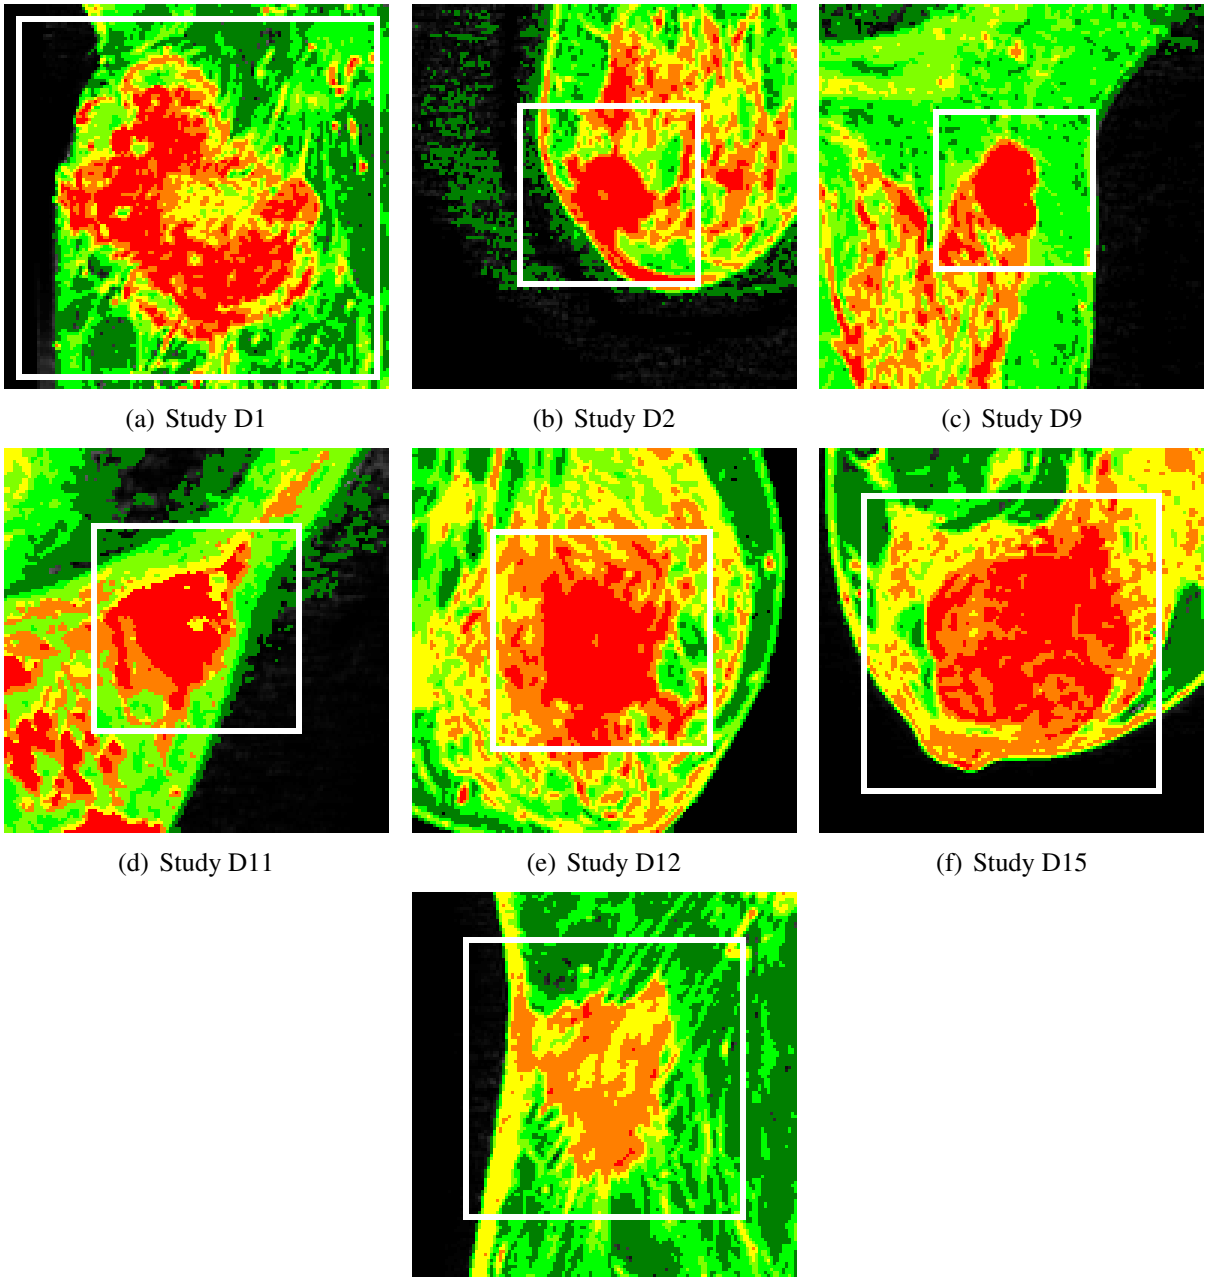

(a) Study D1 (b) Study D2 (c) Study D9

(d) Study D11 (e) Study D12 (f) Study D15

(g) Study D16

Figure 4.5: K-means cluster result, with k = 7

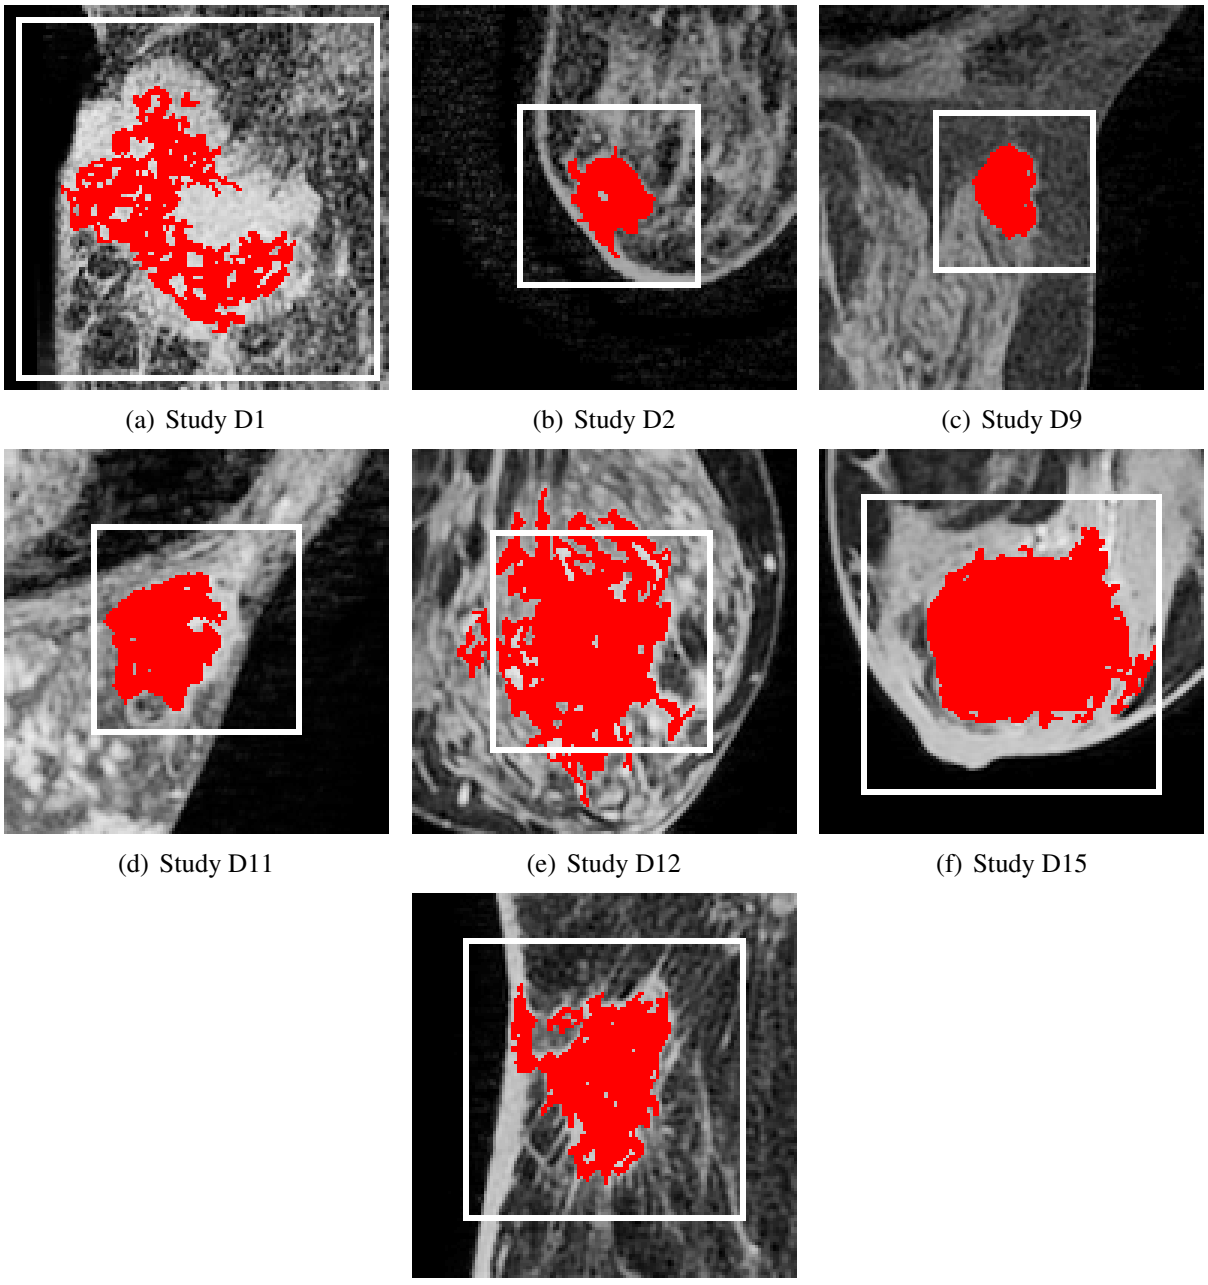

(a) Study D1 (b) Study D2 (c) Study D9

(d) Study D11 (e) Study D12 (f) Study D15

(g) Study D16

Figure 4.6: Tumor masks generated by clustering assisted method

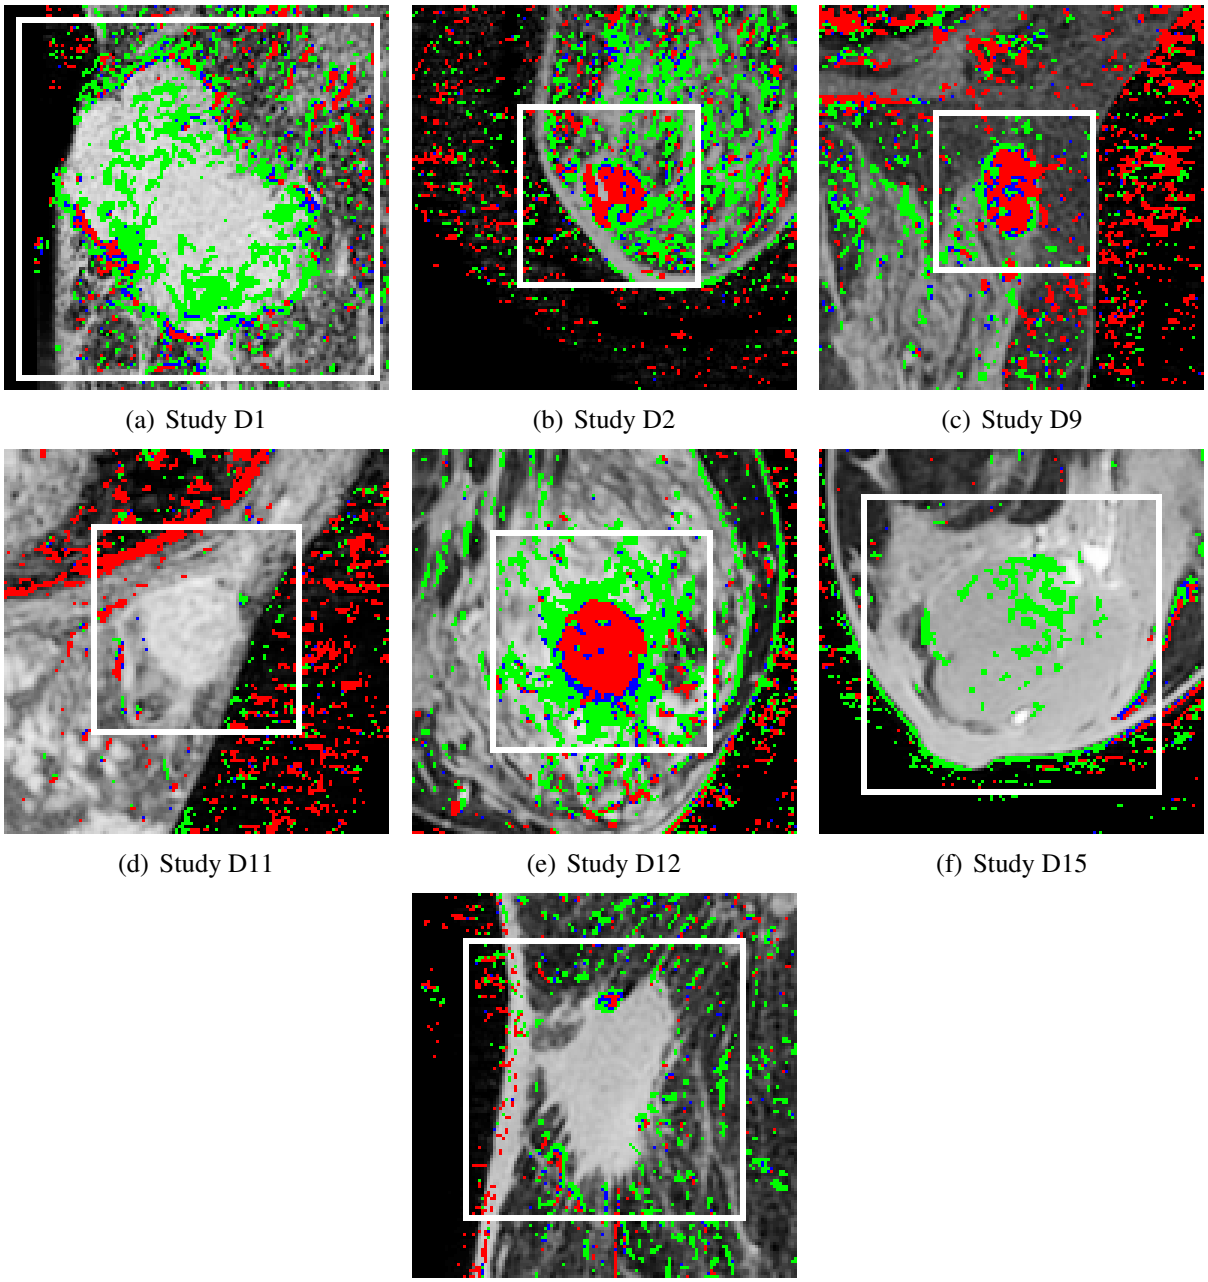

(a) Study D1 (b) Study D2 (c) Study D9

(d) Study D11 (e) Study D12 (f) Study D15

(g) Study D16

Figure 4.7: Initial 3TP Results

In Figure 4.7, this effect appears as halo of pixels that is either part of the background or within the empty chest cavity. The lesions are indicated by a white area-of-interest box. Study D1 shows a ring of Type I enhancement around a central area of non-enhancement, as does study D15. Studies D2, D9, and D12 show central areas of Type II enhancement, surrounded by Type I enhancement. Studies D11 and D16 show no enhancement in the lesion areas.

To approximate the results of the CADStream 3TP results, several operations were applied to these images. First, a threshold for minimum enhancement was applied, where any pixels with a 3TP value of less than 40% were set to zero. This threshold affected mostly background (non-tissue) pixels. After applying a threshold, each study still contained a large proportion of pixels that did not lie within the area-of-interest (AOI) but were identified as a Type I, II, or III curve. Some of these pixels belonged to blood vessels in the breast tissue. These blood vessel pixels could be distinguished by examining adjacent axial slices. The vessels show a connected structure in the adjacent images and "travel" along the vessel.

The second operation applied to the image to remove these "vessel-pixels" was a median filter. A 3×3 median filter was used to eliminate these single-pixel enhancements leaving larger, continuous enhancement regions. A larger 5×5 median filter was also attempted, but may have decimated too many pixels for the smaller lesions to be seen. For example, Study D16, only presented a small spot of type I, II, and III on Figure 4.7. In Figure 4.8(g), this area has been eliminated, leaving no significant group of pixels in the AOI. The second image operation applied was a 5×5 median filter that eliminated any enhancing area less than approximately 4-pixels in area. Finally, the 3TP images were clipped to the coronal plane, isolating only the breast tissue. Figure 4.8 presents the post-processed 3TP results.

The 3TP method was also applied to the disease-negative images. The same threshold, median filter, and clipping operations were applied. The disease-negative cases showed some enhancement regions, but no structured enhancement as expected.

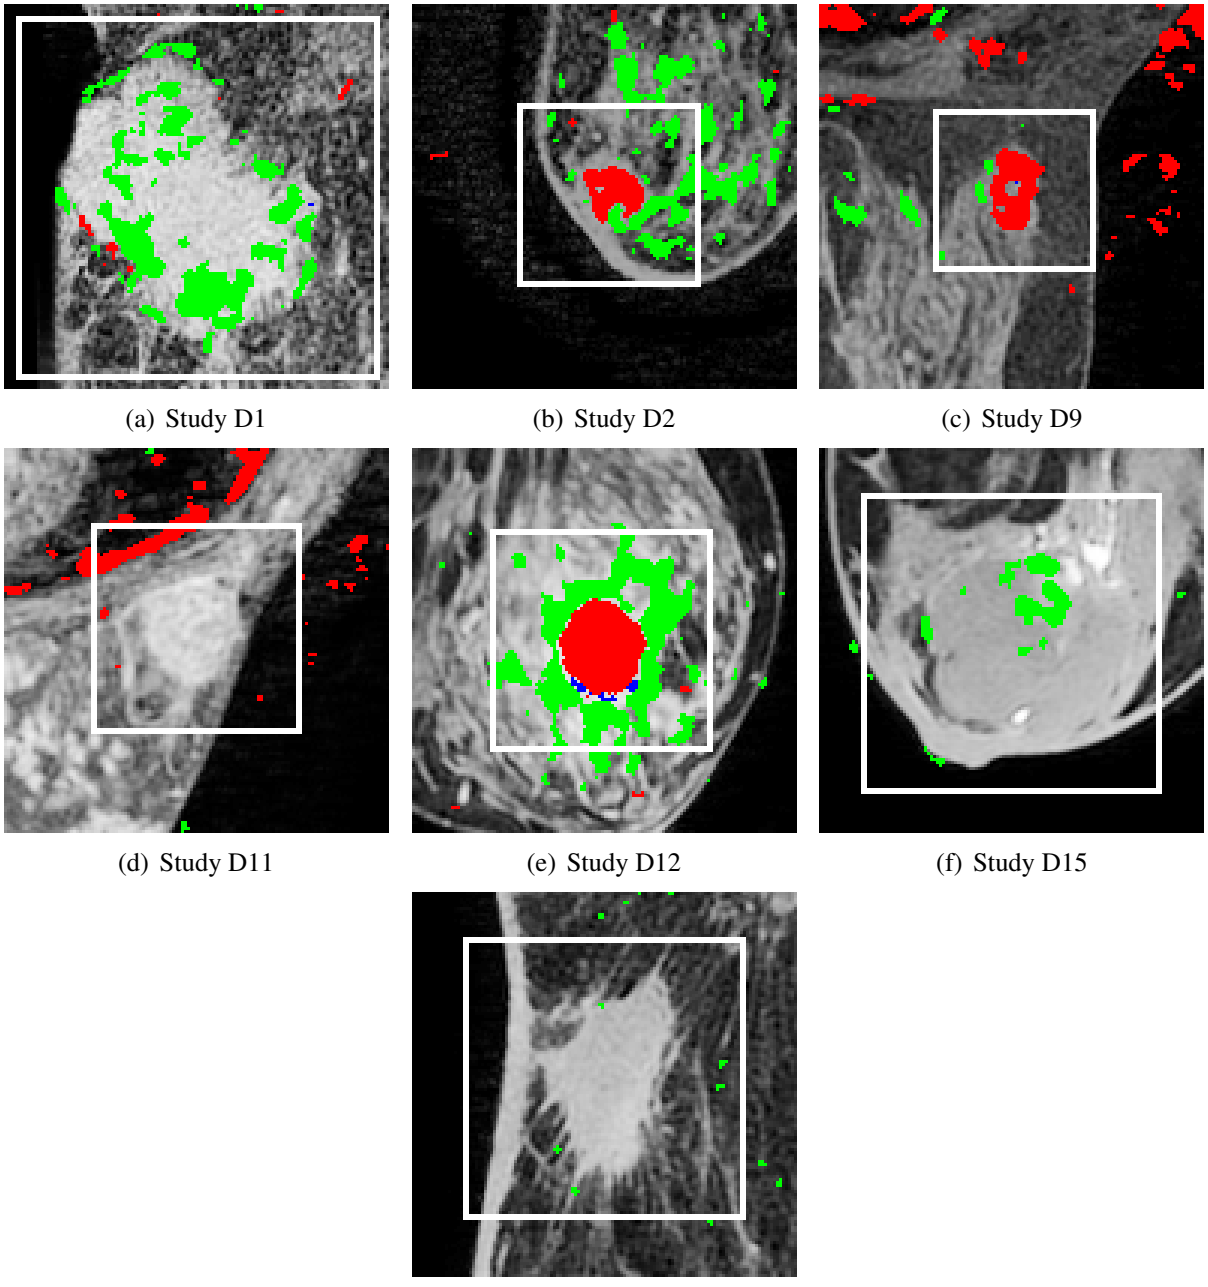

(a) Study D1 (b) Study D2 (c) Study D9

(d) Study D11 (e) Study D12 (f) Study D15

(g) Study D16

Figure 4.8: Disease-positive 3TP results thresholded, filtered with a 5×5 median operation

Table 4.1 summarizes the kinds of enhancement observed in each study. These results were most similar to the CADStream result, but illustrates the necessary post-processing steps to eliminate side effects such as noise and blood vessel enhancement.

Table 4.1: 3TP summary results

Study	Type I	Type II	Type III	Notes
D1		Yes		Type II periphery, center does not enhance
D2		Yes	Yes	Type III center
D9			Yes	Type III center
D11				No enhancement
D12		Yes	Yes	Type III center, Type II boundary
D15		Yes		Type II center
D16				Minimal enhancement

The performance of the 3TP method as a classifier was computed by counting all Type II and Type III curves as disease-positive classifications. Type I curves and non-enhancing pixels were counted as disease-negative classifications. The predicted class was compared to the tumor mask, and the performance statistics were computed. Balanced success was used to measure the performance of the 3TP method due to the classes sizes being very unbalanced. In the training data, there were 11,746 disease-positive examples (7.8390%) and 138,094 disease-negative examples (92.161%). The balanced success rate was 52.617%, specificity 95.664%, and sensitivity 9.5692%.

4.5 Perfusion Parameter Recovery

Kernel recovery was performed on each of the stacked volumes with the following three approaches:

1. pseudoinverse with a sampled arterial input function,
2. pseudoinverse with a gamma variate curve, and
3. iterative blind deconvolution with a non-negative kernel constraint.

For each disease positive study, a number of randomly chosen samples for three areas were sampled: within the lesion boundary, outside of the lesion boundary, and inside the heart. Ten samples for each area were gathered. Figure 4.9 shows the signal intensity plotted versus time in seconds.

Studies D1, D2, D9, D12, D15, and D16 show that tumor pixels reach a higher level of enhancement than non-tumor pixels. Study D11 presents itself as a smaller signal strength.

The first deconvolution method used the sampled heart signal as the arterial input function. The pseudoinverse operation was used to solve for the kernel. Each tissue kernel in these examples showed an initial decrease in the value of the signal strength to a minimum point, followed by an increase in value. According to the model of perfusion previously presented, the upslope of the kernel function is not explained by the initial enhancement value, but is likely due to the extended period of time that the signal lasts.

The second deconvolution used a fixed input function based on a gamma variate curve translated and scaled to fit the time scale and magnitude of the signal. Pseudoinversion was also used for the kernel recovery. These results are shown on Figure 4.11.

The third deconvolution method employed blind deconvolution. The averaged cardiac signal was used to initialize the blind convolution. Two constraints were applied on the kernel at each iteration. First, the kernel was enforced to be non-negative (i.e., negative values were set to zero). Second, the kernel series at each iteration was adjusted to non-increasing (i.e., any two adjacent points on the kernel that exhibited an increase were replaced by the average of the two values). The kernel result converged after six iterations. Figure 4.12 shows the results of the blind deconvolution.

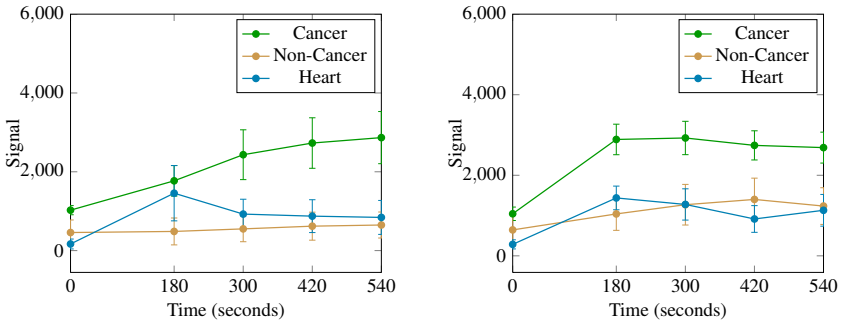

(a) D1

(b) D2

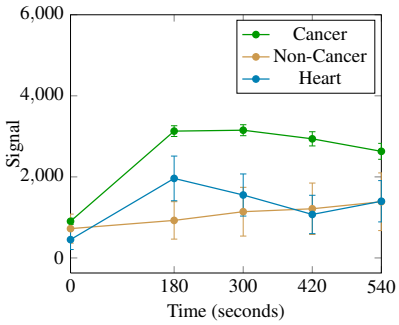

(c) D9

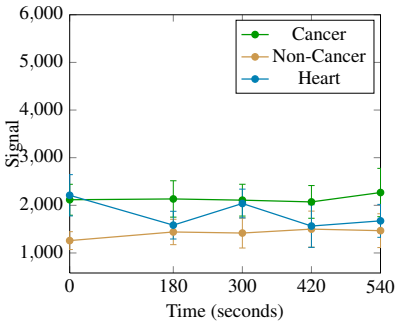

(d) D11

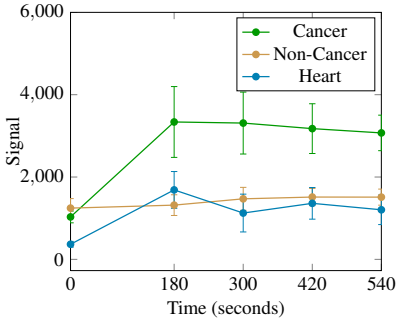

(e) D12

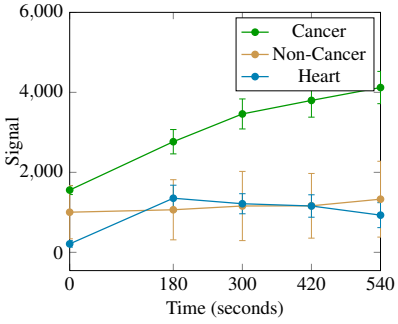

(f) D15

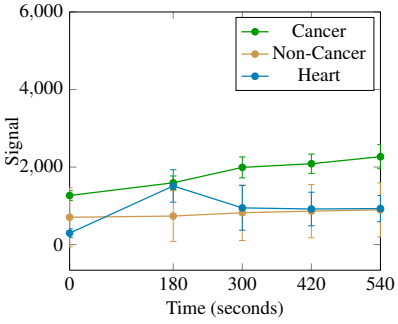

(g) D16

Figure 4.9: Averaged signal samples from disease-positive studies

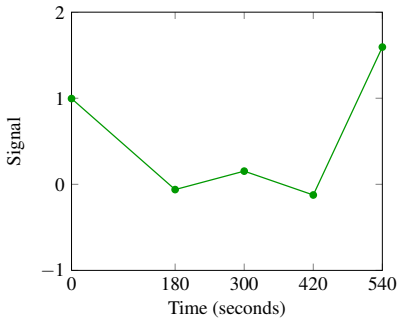

(a) D1

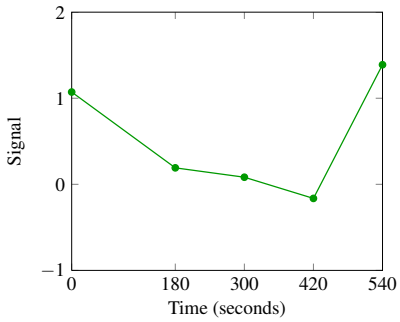

(b) D2

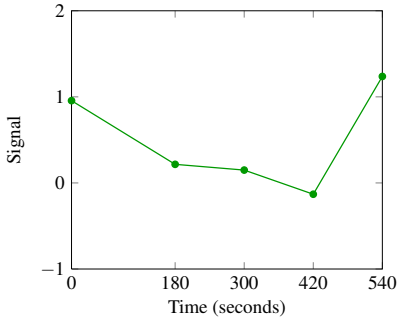

(c) D9

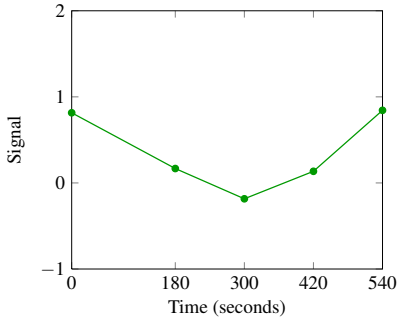

(d) D11

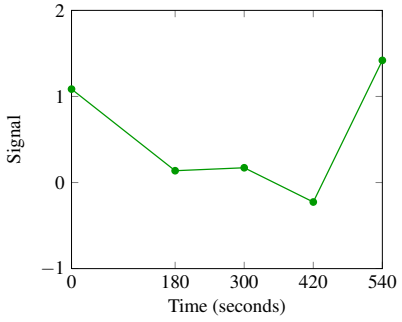

(e) D12

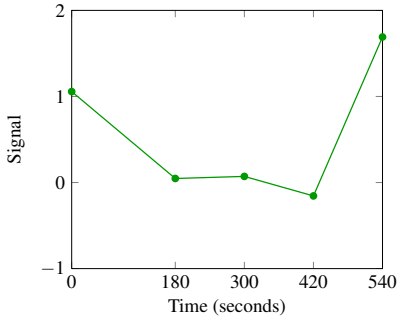

(f) D15

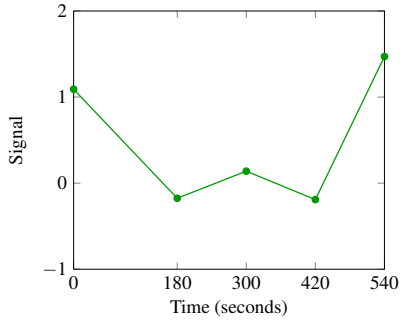

(g) D16

Figure 4.10: Kernel recovered with artery

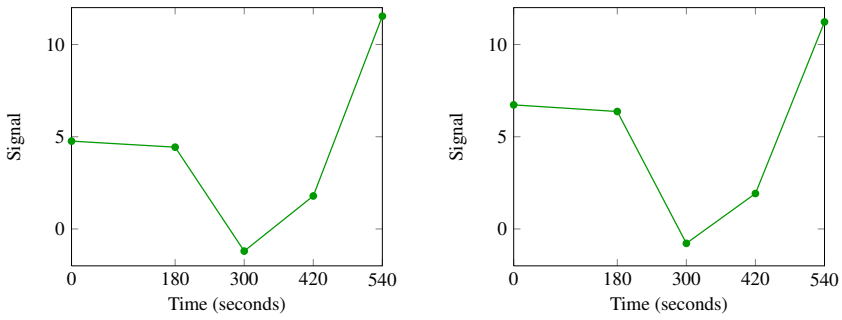

(a) D1

(b) D2

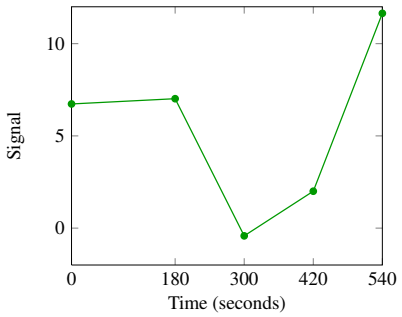

(c) D9

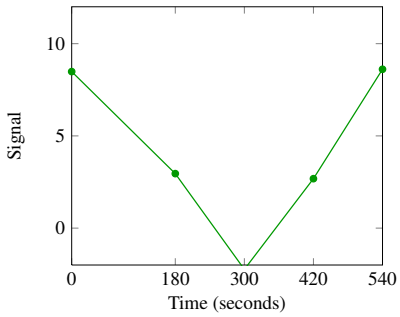

(d) D11

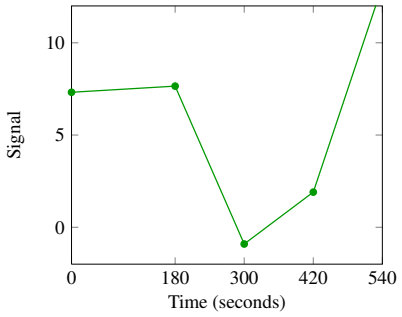

(e) D12

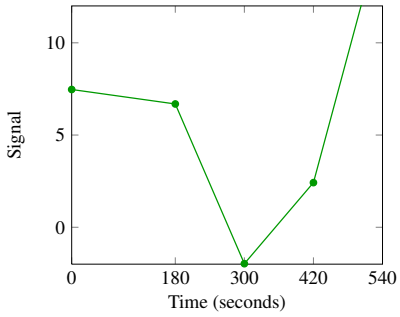

(f) D15

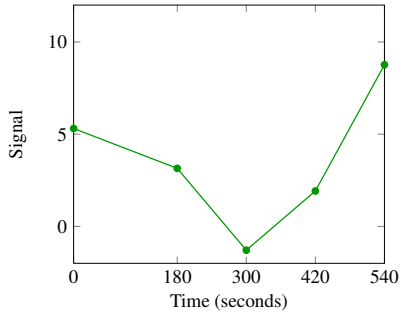

(g) D16

Figure 4.11: Kernel recovered with gamma variate

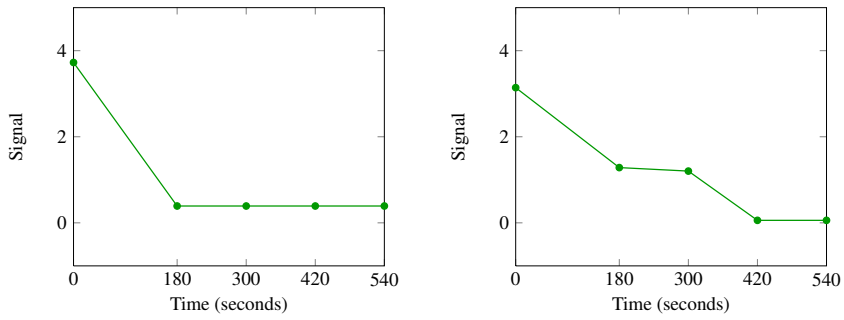

(a) D1

(b) D2

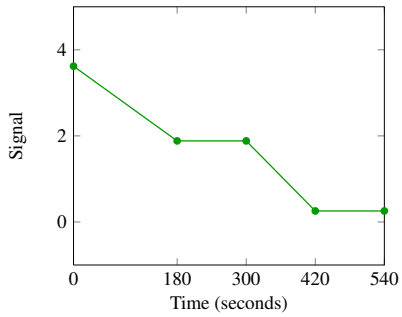

(c) D9

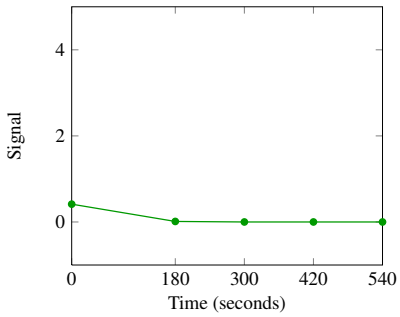

(d) D11

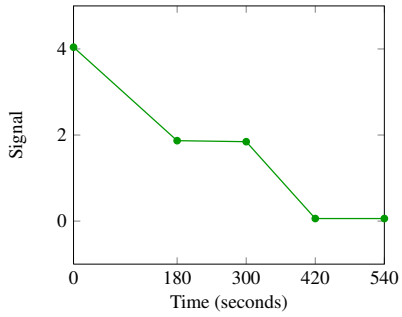

(e) D12

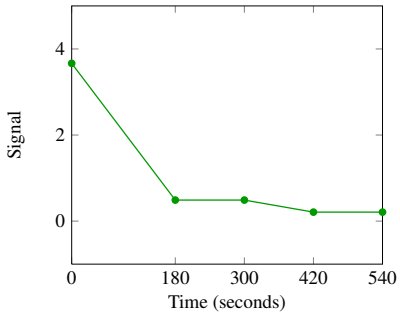

(f) D15

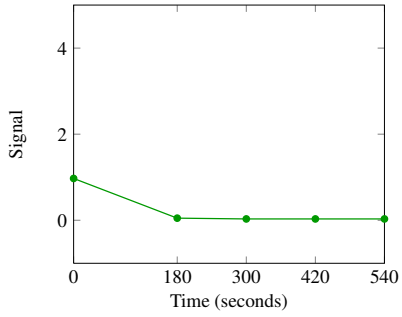

(g) D16

Figure 4.12: Kernel recovered with gamma variate plus blind deconvolution

4.6 Support Vector Machine Classification

With two different methods of tumor identification and three different approaches to deconvolution, a tissue kernel function was computed from each pixel's time series using the training data, and marked with a "Cancer"/"Non-Cancer" training value. For each classification trial, the recovered kernel, $K(t)$ and the original signal data, $C(t)$ were composed into feature vectors in the following patterns.

1. $\langle C(t) \rangle$
2. $\langle C(t), K(t) \rangle$
3. $\langle K'(t), \int K(t), max(K(t)) \rangle$

The first feature vector composition included only the enhancement signal. The second feature vector composition appended the recovered kernel to the feature vector. The third feature vector composition included the first derivative of the kernel vector and derived perfusion parameters of flow and volume.

Three different SVM classification kernels were selected to classify each feature vector formulation: a linear kernel, a polynomial kernel of order 2, and a Gaussian kernel. Ten-fold stratified cross validation was used to measure the predictive performance of the SVM-based classifier on each feature vector formulation. In each iteration of cross validation, the collection feature vectors were randomly partitioned into a training and testing sets. Because it is *stratified* cross validation, the number of examples for each class in each partition are kept proportional to the size of each class in the total set. This is a preferred method of cross-validation, because the number of "Cancer" pixels are small in comparison to the number of "Non-Cancer" pixels. Each feature vector was normalized to unit length prior to classification.

The first classification trial used feature vectors consisting only of the original signal data. One copy of the training/testing set was marked with manually drawn tumor masks. The second set was marked with clustering-assisted tumor masks. The SVM was trained using a linear kernel and

the results are shown in Table 4.2. After 10-fold cross validation, the data set marked using the assisted tumor masks showed an increase in balanced success, sensitivity, and specificity.

Table 4.2: Drawn tumor masks vs. assisted, feature vector 1, linear kernel

Statistic	Drawn	Assisted	$\Delta\%$
Balanced Accuracy	69.33%	72.95%	+3.62%
Sensitivity	75.82%	76.78%	+0.96%
Specificity	62.84%	69.12%	+6.28%

The second trial included only the training/testing set marked with cluster-assisted tumor masks. Kernel recovery was performed using pseudoinverse deconvolution with a sampled cardiac arterial input function. The original signal was combined with the recovered kernel as described in feature vector pattern 2. The 10-fold cross validation was performed three times, varying the classification kernel each time. The results of this trial are show in Table 4.3.

Table 4.3: Linear, quadratic, and Gaussian kernels, using assisted mask, feature vector 2

Statistic	Linear	Poly = 2	Gaussian
Accuracy	72.93%	74.61%	75.75%
Sensitivity	76.74%	76.46%	77.69%
Specificity	69.13%	72.76%	73.81%

Appending the kernel function to the feature vector with a linear classifier does not show an improvement beyond one decimal place. However, between classification kernels, balanced success, sensitivity, and specificity improve when using a quadratic classification kernel, and improve even more using a Gaussian kernel.

In the third trial, the third formulation of the feature vector was trained and tested. Each feature vector was marked using an assisted tumor mask, and a Gaussian kernel was employed. The results are shown in Table 4.4. The results with this feature vector formulation shows an

increase in balanced success and sensitivity, but a decrease in specificity when compared to the Gaussian from Table 4.3.

Table 4.4: Gaussian kernel, assisted mask, feature vector 3

Statistic	Gaussian
Accuracy	76.79%
Sensitivity	82.52%
Specificity	71.07%

4.7 Application

To evaluate and visualize the performance of the SVM classifier using the derived tissue kernel feature vector, several different schemes for training and classification were applied to the data. Leave-one-out cross validation (LOOCV) was too computationally expensive to be practiced. Instead, the set of examples was divided by patient. Of the seven patients with tumors present, a subset of 1 to 6 were chosen to train the SVM. This result was then used to classify the remaining patients. Figure 4.13, shows the classification result from one example run in which feature vectors generated from study D12 were used as the training set. The remaining studies were classified with the support vector machine that resulted from this. The color coding used in Figure 4.13 to represent true-positive red, true-negative green, false-positive blue, and false-negative yellow.

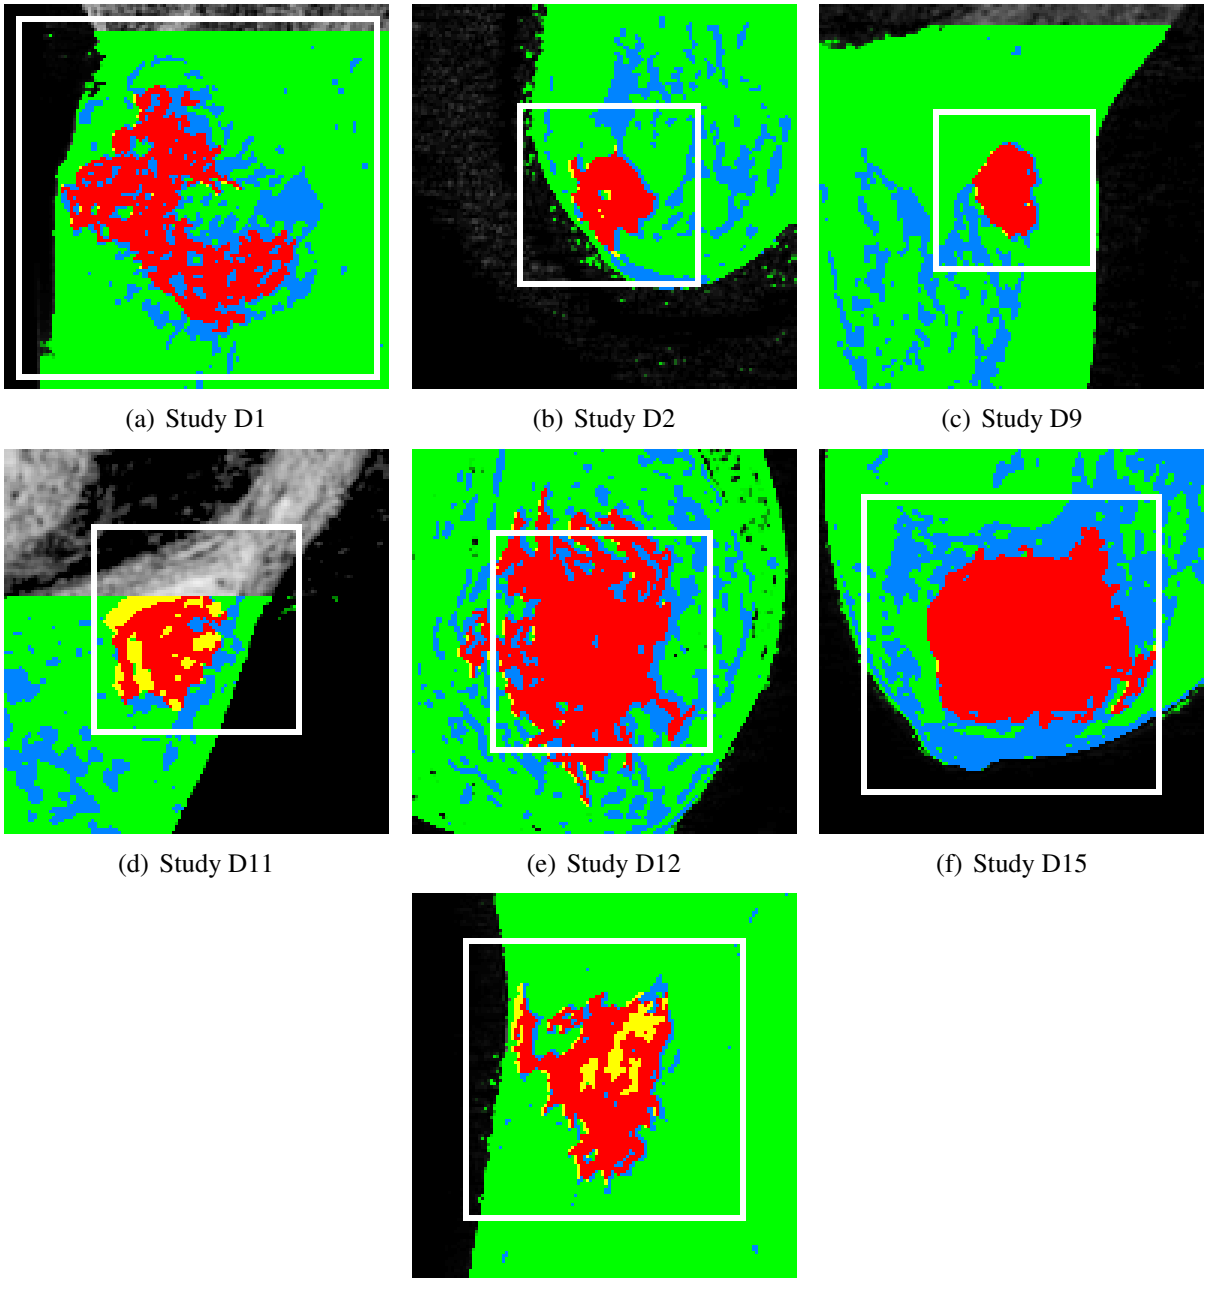

(a) Study D1 (b) Study D2 (c) Study D9

(d) Study D11 (e) Study D12 (f) Study D15

(g) Study D16

Figure 4.13: SVM classification

Chapter 5

Discussion

5.1 3TP Performance

The 3TP method produces majority classifications where most of the pixels are identified with either no-enhancement or Type I curves. This is a common occurrence in the clinical application of the 3TP method. Tumors which are identified by pathology do not exhibit Type II or Type III enhancement curves. In light of this lack of detection sensitivity demonstrated by the first phase of this research, there is a reliance on alternate modalities to improve a diagnosis. In fact, ultrasound and conventional X-ray mammography are examined in greater detail where morphology, not enhancement behavior, is the primary diagnostic factor. With the standard tumor morphology lexicon, the radiologist looks for common patterns such as spiculations.

In Figure 4.7, Studies D1, D11, D15, and D16 do not contain any Type II or Type III enhancement curves. It should be noted that the cardiac space has Type III enhancement curves. This is expected as the contrast agent circulates through the heart. Datasets D2, D9, and D12 do exhibit Type III enhancement curves. The balanced performance of the 3TP method, even when the washout tolerance was adjusted higher, has a very low detection sensitivity. These results were visually compared to the Dynamics View in CADStream. Indeed, the same curve classifications were made.

The Dynamics View in CADStream applied some additional filtering mimicked by the 5×5 median filtering applied to the 3TP results. As in CADStream, this reduced many of the artifacts attributable to signal noise. This is most apparent in study D9 comparing Figure 4.7(c) and Figure 4.8(c). There are many single-pixel "speckles" that are not part of a tumor or vasculature where enhancement would be expected.

The balanced success rate of 3TP, 52.617%, reinforces that this method alone does not yield much value in automated diagnosis. The low success rate could not be attributed to the limited temporal resolution of the data, because the 3TP method only uses three time points.

5.2 Cluster-Assisted Tumor Masks

The specification of ground truth using the software developed for the experiment showed a difference between the manual method and the cluster-assisted method. The manual method involved the radiologist "painting" a mask over the approximate boundary of the tumor. The radiologist could control the resolution of the marking tool and increase its zoom level to approximate the boundary of the tumor. Other methods, such as 3-D mammography, could provide higher spatial resolution to draw a more detailed boundary; however, the scope of this method is limited to the single modality of contrast-enhanced MRI.

The K-means clustering algorithm differentiated groups of pixels without a model of behavior. In visualizing the clustering result, however, it was useful to apply the central volume theorem. Generally, areas with higher perfusion should yield a higher maximum value for $K(t)$. A color palette of green to yellow to red was used to indicate cluster classes that have a potentially high level of blood flow. This result alone is not that useful for tumor classification. In Figure 4.5, the tumors in D2, D9, and D11 have foci that need a manual discrimination from other pixels belonging to the same cluster. Additional context was required to eliminate other pixels from the tumors mask. From manual inspection of adjacent axial layers, some of these pixels were identified with vasculature. In addition, other pixels were eliminated based on information from in the radiologist's report. For example, some lesions were identified in reports as non-cancerous cysts or fibroadenoma.

The tumor in Figure 4.5, D16, did not show the typical enhancement pattern, as can be seen in the 3TP result on Figure 4.8. This tumor also did not exhibit the highest degree of enhancement as the other studies; however, the clustering algorithm still grouped these pixels. With morphological evidence from the radiologist's report, this was identified as a tumor.

The cluster-assisted method provided much more detail in the tumor mask. Selecting contiguous areas of the same cluster group improved the specification of an actual tumor over manually drawn masks in the following two ways:

1. First, the tumor mask can be reproducibly created and is more consistent.

2. The model-free nature of the method allows some flexibility in determining which pixels are part of the tumor.

Comparing the results of 10-fold cross validation between the manually drawn tumor masks and the cluster-assisted method showed improvements in success rate and also in specificity and sensitivity. Incorporating pathology, context clues, and morphology is the preferred method of specifying the actual tumor than by simply drawing the tumor masks.

5.3 Methods of Perfusion Recovery

The accurate recovery of perfusion is inherently difficult because the arterial input function, $C_{AIF}(t)$, cannot be measured at each pixel of computation. Therefore, an arterial input function must be obtained in a way other than direct measurement. Breast tissue is also difficult to recover perfusion in that few large arteries/veins are adjacent to where ductal carcinoma typical occurs. This is as opposed to cerebral tissue where such physiology is present. Therefore, three different approaches of arterial input selection were attempted.

1. Deconvolution with a fixed model curve as the arterial input function is often used in simulated models of perfusion. The fixed set of parameters were chosen for the gamma variate curve. The subsequent deconvolution result, placed in a feature vector, provided results that out-performed classification with the control case.

2. The second approach to perfusion recovery used a sampled arterial input function from an averaged sample of pixels from the heart. When incorporated into a feature vector in the same fashion, it demonstrated higher performance versus the model curve.

3. Blind deconvolution was expected to have higher performance. Beginning with the sampled input function as a seed, each iteration generated a new pair of kernel and input function estimates. The rules used to refine the results at each iteration intended to isolate the initial enhancement, where the tissue kernel is monotonically decreasing. When added into

a feature vector, this had the least performance of the two other methods of tissue kernel recovery.

5.4 Performance of Adding Perfusion Data in Classification

The 3TP method is a standard technique to map the enhancement properties of tissue and provide scale of microvascular permeability. The index of microvascular permeability is categorized into one of three tissue classes, and applied as a diagnostic visualization in medical software such as CADStream. In practice, the 3TP image product is not used as a diagnostic option in isolation. Clinical radiologists improve their diagnostic accuracy by considering other forms of evidence simultaneously. While this does improve diagnostic performance, the analysis time required per study increases.

Classification algorithms can improve diagnostic performance by including the entire signal curve, rather than only three time points. Perfusion recovery is used to quantify these physiological parameters.

Relative perfusion parameters of flow, mean transit time, and volume are derived from the kernel function of a unit of tissue by applying the convection-diffusion model of indicator particle exchange. However, there are model-free properties of the kernel function that can be computed and added to a feature vector for classification.

5.5 Conclusions

This research explored all the steps of breast cancer computer-aided diagnosis: image analysis, tissue characterization, and cancer classification. Instead of relying on simulated data, this experiment used real MRI studies. Many challenges were exposed and addressed with this "real-world" context. We identify three main contributions from these results:

1. The addition of relative perfusion parameters via deconvolution improves classification performance versus using only the enhancement information.

2. This approach to perfusion-based tissue classification out-performs the 3-Time Points method employed in commercial-grade software tools.

3. Using an enhancement-based clustering method for segmentation to specify ground truth improves classifier performance by reducing false training examples.

5.6 Future Work

The experimental method developed in this research exploited the data from the only diagnostic protocol available, VIBRANT. Even with this grade of data, the results demonstrate some of the improvements that are possible in computer aided breast cancer diagnosis. There are several steps that warrant more exploration.

The effects of the low time resolution and extended period of measurement had an impact on the perfusion recovery results. The method described here could have taken advantage of many more time points as part of the image acquisition process. Most simulated processes use a time step as small as one second. The experimental data used in this experiment had entire volumes gathered at each time step. If spatial resolution could be sacrificed for temporal resolution, then the perfusion results may have better performance.

Conversely, better spatial information could have been exploited by incorporating a way to look at morphology. For example, after characterizing tumor pixels by perfusion, then a second stage could be used to cluster pixels and generate a very accurate boundary for each cluster. A way to automatically characterize the morphology of the cluster boundary could then be used to provide a higher confidence prediction.

Texture has been incorporated into classifiers to detect patterns of malignant tumors. The typical implementation is to transform a moving window of pixels into a feature vector, and then perform training and classification on that population. Taking the results of this experiment, the relative perfusion of several adjacent pixels within a window could be combined simultaneously in a feature vector. Furthermore, better information, such as the gradient vector, could also be

considered. This could potentially be used to identify the boundary of a tumor, and not rely on clustering pixels that are independently computed.

One next step could be to implement elements of this research into a software-based tool in a clinical setting. One challenge of developing such a tool is to communicate what the classifier output means. One way of building trust in the diagnostic metric is to incorporate a measure of uncertainty. By attributing a confidence level for each prediction, one can better interpret the results as a diagnostic measure.

Bibliography

[AD88] GR Ayers and J.C. Dainty. Interative blind deconvolution method and its applications. *Optics letters*, 13(7):547, 1988.

[ARS98] K. Alsabti, S. Ranka, and V. Singh. An efficient k-means clustering algorithm. In *First Workshop on High-Performance Data Mining*, volume 10. Citeseer, 1998.

[AWY+99] C.C. Aggarwal, J.L. Wolf, P.S. Yu, C. Procopiuc, and J.S. Park. Fast algorithms for projected clustering. *ACM SIGMOD Record*, 28(2):61–72, 1999.

[BBH+99] G. Brix, M.L. Bahner, U. Hoffmann, A. Horvath, and W. Schreiber. Regional blood flow, capillary permeability, and compartmental volumes: Measurement with dynamic ct–initial experience. *Radiology*, 210:269–276, 1999.

[BF98] P.S. Bradley and U.M. Fayyad. Refining initial points for k-means clustering. In *Proc. 15th International Conf. on Machine Learning*, volume 727. Morgan Kaufmann, San Francisco, CA, 1998.

[BH07] A. Ben-Hur. PyML: Release 0.6. 11.(Colorado State University. Fort Collins, CO), 2007.

[BKF96] J. A. Baker, P. J. Kornguth, and C. E. Floyd. Breast imaging reporting and data system standardized mammography lexicon: Observer variability in lesion description. *American Journal of Roentgenology*, 166:773–778, September 1996.

[Bur98] C.J.C. Burges. A tutorial on support vector machines for pattern recognition. *Data mining and knowledge discovery*, 2(2):121–167, 1998.

[CGC02] F. Clamante, D.G. Gadian, and A. Connelly. Quantificaiton of perfusion using bolus tracing magnetic resonance imaging in stroke. *Stroke*, 33:1146–1151, January 2002.

[CL01] C.C. Chang and C.J. Lin. LIBSVM: a library for support vector machines, 2001.

[CP04] D. J. Collins and A. R. Padhani. Dynamic magnetic resonance imaging of tumor perfusion. *IEEE Engineering in Medicine and Biology*, October 2004.

[CWM+03] R.F. Chang, W.J. Wu, W.K. Moon, Y.H. Chou, and D.R. Chen. Support vector machines for diagnosis of breast tumors on US images. *Academic radiology*, 10(2):189–197, 2003.

[CWMC03] R.F. Chang, W.J. Wu, W.K. Moon, and D.R. Chen. Improvement in breast tumor discrimination by support vector machines and speckle-emphasis texture analysis. *Ultrasound in Medicine and Biology*, 29(5):679–686, 2003.

[Dal98] H Daldrup. Correlation of dynamic contrast-enhanced MR imaging with histologic tumor grade: Comparison of macromolecular and small-molecular contrast media. *American Journal of Roentgenology*, 171:941–949, 1998.

[DLRH+03] H.E. Daldrup-Link, J. Rydland, T.H. Helbich, A. Bjønerud, K. Turetschek, J.A. Kvistad, E. Kaindl, T.M. Link, K. Staudacher, D. Shames, R.C. Brasch, O. Haraldseth, and E.J. Rummeny. Quantification of breast tumor microvascular permeability with feruglose-enhanced mr imaging: Inital phase II multicenter trial. *Radiology*, 229:885–892, 2003.

[ELPC+03] D.B. Ellegala, H. Leong-Poi, J.E. Carpenter, A.L. Klibanov, S. Kaul, M.E. Shaffrey, J. Sklenar, and J.R. Lindner. Imaging tumor angiogenesis with contrast ultrasound and microbubbles targeted to $\alpha_v\beta_3$. *Circulation*, 108:336–341, 2003.

[ENYW+02] I. El-Naqa, Y. Yang, MN Wernick, NP Galatsanos, and RM Nishikawa. A support vector machine approach for detection of microcalcifications. *IEEE Transactions on Medical Imaging*, 21(12):1552–1563, 2002.

[EWL+94] J.G. Elmore, C.K. Wells, C.H. Lee, D.H. Howard, and A.R. Feinstein. Variability in radiologists' interpretations of mammograms, 1994.

[FHSK+05] E. Furman-Haran, E. Schechtman, F. Kelcz, K. Kirshenbaum, and H. Degani. Magnetic resonance imaging reveals functional diversity of the vasculature in benign and malignant breast lesions. *Cancer*, 104:708–718, 2005.

[FON08] Niall Fitzgerald, Finbarr O'Sullivan, and George Newman. Maximum likelihood deconvolution of dynamic contrast mri data. In *Nuclear Science Symposium Conference Record, 2008. NSS '08. IEEE*, pages 4482–4488, Oct. 2008.

[GPM+03] M.S. Gee, W.N. Procopio, S. Makonnen, M.D. Feldman, N.M. Yeilding, and W.M.F. Lee. Tumor vessel developement and maturation impose limits of the effectiveness of anti-vascular therapy. *American Journal of Pathology*, 162(1), January 2003.

[HFE+94] S.E. Harms, D.P. Flamig, W.P. Evans, S.A. Harries, and S. Brown. MR imaging of the breast: Current status and future potential. *American Journal of Roentgenology*, 163:1039–1047, 1994.

[Hor97] J.P. Hornak. The basics of NMR. *Online Textbook*, 1997.

[HWU+10] K. Hoyt, J.M. Warram, H. Umphrey, L. Belt, M.E. Lockhart, M.L. Robbin, and K.R. Zinn. Determination of Breast Cancer Response to Bevacizumab Therapy Using Contrast-Enhanced Ultrasound and Artificial Neural Networks. *Journal of Ultrasound in Medicine*, 29(4):577, 2010.

[JHHMS04] M. Jerosch-Herold, X. Hu, N.S. Murthy, and R.T. Seethamraju. Time delay for arrival of MR contrast agent in collateral-dependent myocardium. *IEEE Transactions on Medical Imaging*, 23(7), July 2004.

[JMW+05] A. Jemal, T. Murray, E. Ward, A. Samuels, R.C. Tiwari, A. Ghafoor, E.J. Feuer, and M.J. Thun. Cancer statistics, 2005. *CA: A Cancer Journal for Clinicians*, 55:10–30, 2005.

[KBKS04] S.L. Keeling, R. Bammer, T. Kogler, and R. Stolberger. On the convolution model of dynamic contrast enhanced magnetic resonance imaging and nonparametric deconvolution approaches. *Optimization and Control, University of Graz*, April 2004.

[KBSRS+02] Rachel Katz-Brull, Dalia Seger, Dalia Rivenson-Segal, Edna Rushkin, , and Hadasa Degani. Metabolic markers of breast cancer: Enhanced choline metabolism and reduced choline-ether-phospholipid synthesis. *Cancer Research*, 62:1966–1970, 2002.

[KCS03] T.S. Koh, L.H. Cheong, and Y.C. Soh. A physiologic model of capillary-tissue exchange for dynamic constrast-enhanced imaging of tumor microcirculation. *IEEE Trans. on Biomed. Eng.*, 50(2):159–167, February 2003.

[Kis01] V.G. Kiselev. On the theoretical basis of perfusion measurements by dynamic susceptibility contrast MRI. *Magnetic Resonance in Medicine*, 46:1113–1122, 2001.

[KKG+99] Matjac Kukar, Igor Kononenko, Ciril Groselj, Katarina Kralj, and Jure Fettich. Analysing and improving the diagnosis of ischaemic heart disease with machine learning. *Artificial Intelligence in Medicine*, 16(1):25–50, 1999.

[KP03] C. Kotropoulos and I. Pitas. Segmentation of ultrasonic images using support vector machines. *Pattern Recognition Letters*, 24(4-5):715–727, 2003.

[KRV+00] K.A. Kvistad, J. Rydland, J. Vainio, H.B. Smethurst, S. Lundgren, H.E. Fjøsne, and O. Haraldseth. Breast lesions: Evaluation with dynamic contrast-enhanced t1-weighted first-pass perfusion MR imaging. *Radiology*, 216:545–553, 2000.

[KT10] A. Kassner and RE Thornhill. Texture Analysis: A Review of Neurologic MR Imaging Applications. *American Journal of Neuroradiology*, 31(5):809, 2010.

[Kuh99] C. K. Kuhl. Dynamic breast MR imaging: Are signal intensity time course data useful for differential diagnosis of enhancing lesions? *Radiology*, 211:101–110, 1999.

[KWCL04] T.S. Koh, X.Y. Wu, L.H. Cheong, and C.C.T. Lim. Assessment of perfusion by dynamic contrast-enhanced imaging using a deconvolution approach based on regression and singular value decomposition. *IEEE Transactions on Medical Imaging*, 23(12), December 2004.

[LAC01] C. Levine, K. Armstrong, and S. Chopra. *Diagnosis and management of specific breast abnormalities*. MD: Agency for Healthcare Research and Quality, September 2001.

[Lea04] Jeffery J. Leader. *Numerical Analysis and Scientific Computation*. Pearson Addison-Wesley, 2004.

[LKP+08a] Sang Ho Lee, Jong Hyo Kim, Jeong Seon Park, Jung Min Chang, Sang Joon Park, Yun Sub Jung, and Woo Kyung Moon. Computerized segmentation and classification of breast lesions using perfusion volume fractions in dynamic contrast-enhanced mri. In *BioMedical Engineering and Informatics, 2008. BMEI 2008. International Conference on*, volume 2, pages 58–62, May 2008.

[LKP+08b] Sang Ho Lee, Jong Hyo Kim, Jeong Seon Park, Jung Min Chang, Sang Joon Park, Yun Sub Jung, Sungho Tak, and Woo Kyung Moon. Texture analysis of lesion perfusion volumes in dynamic contrast-enhanced breast mri. In *Biomedical Imaging: From Nano to Macro, 2008. ISBI 2008. 5th IEEE International Symposium on*, pages 1545–1548, May 2008.

[LLBT+05] Q.G. Lussanet, S. Langereis, R.G.H. Beets-Tan, M.H.P. van Genderen, J.M.A. van Engelsoven, and W.H. Backes. Dynamic contrast-enhanced MR imaging kinetic parameters and molecular wieght of dendritic contrast agenst in tumor angiogenesis in mice. *Radiology*, 235:65–72, 2005.

[LMS+05] Z. Li, Q. Meng, C. Sun, D. Xu, M. Fan, X. Yang, and D. Chen. Tumor angiogenesis and dynamic CT in colorectal carcinoma: Radiologic-pathologic correlation. *World J Gastroenterol*, 11(9):1287–1291, 2005.

[MCE00] V.M. Mai, Q. Chen, and R.R. Edelman. Absolute quantification of pulmonary rates using flow-sensitive alternating inversion recovery with an extra radiofrequency pulse (FAIRER). *Proceedings of the International Society of Magnetic Resonance Medicine*, 8, 2000.

[MCT+02] W.L. Monsky, C.M. Carreira, Y. Tsuzuki, T. Gohongi, D. Fukumura, and R.K. Jain. Role of Host Microenvironment in Angiogenesis and Microvascular Functions in Human Breast Cancer Xenografts: Mammary Fat Pad versus Cranial Tumors 1, 2002.

[Mil03] K.A. Miles. Perfusion CT for the assessment of tumor vascularity: which protocol? *The British Journal of Radiology*, 76, 2003.

[MM08] O.L. Mangasarian and D.R. Musicant. Lagrangian support vector machine, July 1 2008. US Patent 7,395,253.

[MSSS+96] M Müller-Schimpfle, P. Stoll, W. Stern, S. Kurz, F. Dammann, and C.D. Claussen. Do mammography, sonography, and mr mamography have a diagnostic benefit compared with mamography, and sonography? *American Journal of Roentgenology*, 168, May 1996.

[MSTK99] Teemu Makitie, Paula Summanen, Ahti Tarkkanen, and Tero Kivela. Microvascular Density in Predicting Survival of Patients with Choroidal and Ciliary Body Melanoma. *Invest. Ophthalmol. Vis. Sci.*, 40(11):2471–2480, 1999.

[NDDT+03] G.C. Newman, E. Delucia-Deranja, A. Tudorica, F.E. Hospod, and C.S. Patlak. Cerebral blood volume measurements by T_2^*-weighted MRI and contrast infusion. *Magnetic Resonance in Medicine*, 50:844–855, 2003.

[OC98] L.J. Old and Y. Chen. New paths in human cancer serology. *Journal of Experimental Medicine*, 187(8):1163–1167, April 1998.

[Ore98] S G Orel. High-resolution MR imaging for the detection, diagnosis, and staging of breast cancer. *RadioGraphics*, 18:903–912, 1998.

[Pad99] A. R. Padhani. Dynamic contrast-enhanced MRI studies in human tumours. *The British Journal of Radiology*, 72(1999):427–431, 1999.

[Pad03] A.R. Padhani. MRI for assessing antivascular cancer treatments. *The British Journal of Radiology*, 76, 2003.

[PAK+02] J. Perkiö, H.J. Aronen, A. Kangasmäki, Y. Liu, J. Karonen, S. Savolainen, and L. Østergaard. Evaluation of four postprocessing methods for determnation fo cerebral blood volume and mean transit time by dynamic susceptibility contrast imaging. *Magnetic Resonance in Medicine*, 47:973–981, 2002.

[PFL05] A. Papadopoulos, DI Fotiadis, and A. Likas. Characterization of clustered microcalcifications in digitized mammograms using neural networks and support vector machines. *Artificial Intelligence in Medicine*, 34(2):141–150, 2005.

[POM+09] B. Preim, S. Oeltze, M. Mlejnek, E. Groeller, A. Hennemuth, and S. Behrens. Survey of the visual exploration and analysis of perfusion data. *Visualization and Computer Graphics, IEEE Transactions on*, 15(2):205–220, March-April 2009.

[PP99] D.L. Pham and J.L. Prince. An adaptive fuzzy c-means algorithm for image segmentation in the presence of intensity inhomogeneities. *Pattern Recognition Letters*, 20(1):57–68, 1999.

[PXP00] D.L. Pham, C. Xu, and J.L. Prince. A Survey of Current Methods in Medical Image Segmentation. *Annual Review of Biomedical Engineering*, 2(1):315–337, 2000.

[RP95] WG Roberts and GE Palade. Increased microvascular permeability and endothelial fenestration induced by vascular endothelial growth factor, 1995.

[RSR04] A. Rosset, L. Spadola, and O. Ratib. OsiriX: an open-source software for navigating in multidimensional DICOM images. *Journal of Digital Imaging*, 17(3):205–216, 2004.

[SAS00] M. Stuhrmann, R. Aronius, and M. Schietzel. Tumor vascularity of breast lesions: Potentials and limits of contrast-enhanced doppler sonography. *American Journal of Roentgenology*, 175:1585–1589, 2000.

[STCD04] R. Srikanchana, D. Thomasson, P. Choyke, and A. Dwyer. A comparison of pharmacokinetic models of dynamic contrast enhanced mri. *Computer-Based Medical Systems*, pages 361–366, 2004.

[SWDH02] V.J. Schmithorst, M. Wilke, B.J. Dardzinski, and S.K. Holland. Correlation of White Matter Diffusivity and Anisotropy with Age during Childhood and Adolescence: A Cross-sectional Diffusion-Tensor MR Imaging Study 1, 2002.

[TLN05] T. Twellmann, O. Lichte, and T.W. Nattkemper. An adaptive tissue characterization network for model-free visualization of dynamic contrast-enhanced magnetic resonance image data. *IEEE Transactions on Medical Imaging*, 24(10):1256–1266, October 2005.

[VM03] SVN Vishwanathan and M.N. Murty. Ssvm: A simple svm algorithm. In *Proceedings of the International Conference on Machine Learning*, 2003.

[WsB+03] O. Wu, L. Østergaard, T. Benner, B.R. Rosen, R.M. Weisskoff, and A.G. Sorensen. Tracer arrival timing-insensitive technique for estimating flow in MR perfusion-weighted imaging using singular value decomposition with a block-circulant deconvolution matrix. *Magnetic Resonance in Medicine*, 50:164–174, 2003.

[WSC+99] D. Weinstein, S. Strano, P. Cohen, S. Fields, J.M. Gomori, and H. Degani. Breast fibroadenoma: Mapping of pathophysiologic features with three-time-point, constrast-enhanced mr. imaging–pilot study. *Radiology*, 210:233–240, 1999.

[WsK+03] O. Wu, L. Østergaard, W.J. Koroshetz, L.H. Schwamm, J. O'Donnell, P.W. Schaefer, B.R. Rosen, R.M. Weisskoff, and A.G. Sorensen. Effects of tracer arrival time on flow estimates in MR perfusion-weighted imaging. *Magnetic Resonance in Medicine*, 50:856–864, 2003.

[WSM93] W.H. Wolberg, W.N. Street, and OL Mangasarian. Machine learning techniques to diagnose breast cancer from image-processed nuclear features of fine needle aspirates. *Cancer Letters*, 77(2):163–172, 1993.

[Yeu01] D. Yeung. Human breast lesions: Characterization with contrast-enhanced in vivo proton MR spectroscopy–initial results. *Radiology*, 220:40–46, 2001.

[YLL+05] T. E. Yankeelov, J. J. Luci, M. Lepage, R. Li, L. Debusk, P.C. Lin, R. R. Price, and J. C. Gore. Quantitative pharmacokinetic analysis of DCE MRI data without and arterial input function: a reference region model. *Magnetic Resonance Imaging*, 23:519–529, February 2005.

[ZACS01] R. Zoorob, M.P.H. R. Anderson, C Cefalu, and M. Sidani. Cancer screening guidelines. *American Family Physician*, 63(6), 2001.

[Zie62] K.L. Zierler. Theoretical basis of indicator-dilution method of measuring flow and volume. *Circulation Research*, 10:363–407, March 1962.

[Zie63] K.L. Zierler. Theory of use of indicators to measure blood flow and extracellular volume and calculation of transcapillary movement of tracers. *Circulation Research*, 12:464–471, May 1963.

Appendix: Datasets

Table A.1: Datasets for study, 15 ductal carcinoma, 6 negative, 9 unknown. Axis refers to ”clock-face” position relative to the areola center

Dataset #	Patient #	Left/Right	Axis	Pathology
D1	P21	Right	9	ductal carcinoma
D2	P5	Right	9	ductal carcinoma
D3	P6	N/A	N/A	negative
D4	P16	N/A	N/A	negative
D5	P2	Right	8	unknown
D6	P11	Left	12	unknown
D7	P17	Right	1	unknown
D8	P25	Left	3	ductal carcinoma
D9	P19	Left	3	ductal carcinoma
D10	P3	Right	12	ductal carcinoma
D11	P26	Left	Everywhere	ductal carcinoma
D12	P10	Left	Retroareolar	ductal carcinoma
D13	P4	Right	11	unknown
D14	P7	N/A	N/A	negative
D15	P12	Left	Retroareolar	ductal carcinoma
D16	P22	Right	9	ductal carcinoma
D17	P27	Left	Everywhere	ductal carcinoma
D18	P8	N/A	N/A	negative
D19	P13	Left	7	unknown
D20	P24	Right	10	ductal carcinoma
D21	P23	Left	Everywhere	ductal carcinoma
D22	P29	Right	Retroareolar	unknown
D23	P9	N/A	N/A	unknown
D24	P15	Right	12	ductal carcinoma
D25	P14	Left	7	ductal carcinoma
D26	P28	Right	Retroareolar	unknown
D27	P18	Left	6	negative
D28	P20	Right	12	unknown
D29	P1	Right	Retroareolar	ductal carcinoma

Vita

Bruce Yu-Sun Lin was born in Houston, Texas, in September, 1980. He received his Bachelor of Science in computer science at Louisiana State University in Baton Rouge, Louisiana, in 2002 as a University Medalist. In August 2003, he began graduate studies in computer science continuing at Louisiana State University as a Board of Regents Fellowship recipient. He will receive his doctorate in computer science in December 2010. He is currently employed as a research scientist by the United States Naval Research Laboratory at Stennis Space Center.