Independent Component Analysis for Maternal-Fetal Electrocardiography

by

Kathryn L. Marcynuk

A Thesis submitted to the Faculty of Graduate Studies of

The University of Manitoba

in partial fulfilment of the requirements of the degree of

MASTER OF SCIENCE

Department of Electrical and Computer Engineering

University of Manitoba

Winnipeg, Canada

To my family and friends

Abstract

Separating unknown signal mixtures into their constituent parts is a difficult problem in signal processing called blind source separation. The ability to separate signal mixtures is relevant to a wide variety of areas including telecommunications, image processing, and medicine. One of the benchmark problems of blind source separation is the extraction of the fetal heartbeat from an electrocardiography (ECG) recording in which it is overshadowed by a strong maternal heartbeat. Successful extraction of the fetal heartbeat is more than just an academic problem. Fetal heart monitoring has become common practice in much of North America and is a deciding factor in whether or not to intervene during labour, so accurate fetal heart readings have social and economic impacts on the health care system.

This thesis presents a study of blind source signal separation technique called independent component analysis (ICA), in order to assess its suitability for the maternal-fetal ECG separation problem. This includes an analysis of ICA on both (i) deterministic and stochastic signals, and (ii) simulated and recorded ECG signals. The effectiveness of ICA is investigated by identifying measures to evaluate and quantify its success at extracting source signals.

The experiments presented in this thesis demonstrate that ICA is effective on linear mixtures produced from known source signals of either simulated or recorded ECGs. The success of the ICA extraction was assessed using energy, information theory, and fractal-based measures, as well as visual comparison and heart rate calculation. In all cases there were minimal changes between the source and extracted signals, and the heart rate was recovered accurately. ICA extraction of clinically recorded maternal-fetal ECGs mixtures, in which the source signals were unknown, were less conclusive and appeared to only recover the fetal heart rate.

Visual Abstract

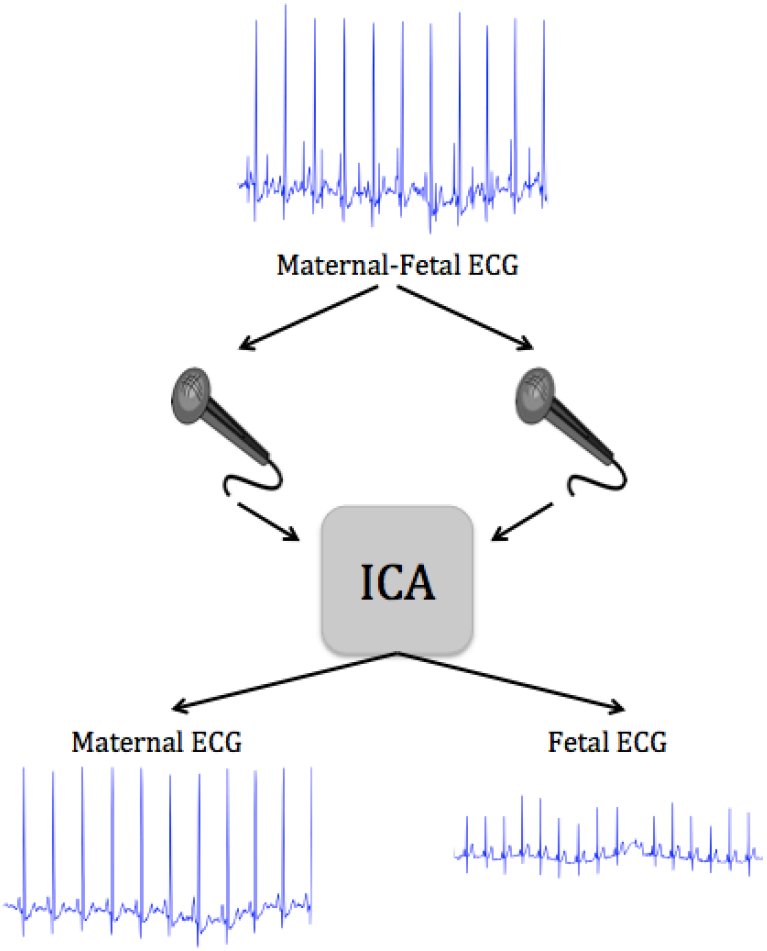

This diagram summarizes the independent component analysis (ICA) process that takes a mixture of signals, obtained by multiple-sensor recording, and separates them into independent sources.

Acknowledgements

This work would not have been possible without the help and support of an entire network of people. Thank you to Dr. Kinsner for welcoming me into the Delta Research group and for trusting me with this important project. I have learned so much from you, and I am excited to delve into new topics and continue this research. I would also like to thank my examining committee, Dr. Dean McNeill and Dr. Michael Domaratzki, for taking the time to truly read this work and provide insightful questions and areas to explore.

Thank you as well to the incredibly supportive Technical Communications group: Dr. Anne Parker, Norma Godavari, Babak Salimifard, and Aidan Topping. You have helped me grow and hone my communication skills, and I could not ask for a better group of people with whom to work.

I would also like to extend my gratitude to the members of the Delta Research group, past and present. Our current group (David, Greg, Hieu, Maryam, Rafi, Sina, and Soroush) explores so many diverse and engaging topics, and I continue to be motivated by each of you. Thank you as well to group alumnus Dario Schor, for your words of wisdom, encouragement, and constructive criticism - and knowing when I needed to hear some of each.

To my friends: thank you for allowing me to talk to you about this project well past the point that you were tired of it, yet never failing to ask how things were going. In particular, I would like to thank Katie Grabau, Kate Langrell, Talia Pankewycz, and Erika Mann who bore the brunt of these discussions.

I owe a special thank you to Matthew Kulasza for being my sounding board, shoulder to cry on, and one-man cheering squad. You continuously motivate me to do my best, while keeping me grounded and reminding me to take a break and have fun. I am so grateful that we were able to share in this experience together.

Finally, but certainly not least, I would like to thank my family for their unwavering support. Thank you to my parents for giving me the gifts of a love of knowledge and the education to pursue this passion. What five-year-old dreams of going to graduate school? Yet you continue to believe in me every step of the way. Lastly, completing this project would not have been nearly as enjoyable without my furry study buddies and "supervisors", Oden and Noelle, sitting by my side.

This work was funded by NSERC, the University of Manitoba Faculty of Graduate Studies, and the University of Manitoba Faculty of Engineering.

Contents

List of Figures

List of Tables

List of Acronyms

bpm	beats per minute	63, 70, 80, 81, 85, 87, 88, 93, 96, 101, 105, 110, 113, 122, 126
BSS	blind source separation	2, 4, 6–10, 13, 15–17, 25, 26, 84, 85, 87, 88, 90, 118
cdf	cumulative distribution function	19
ECG	electrocardiogram	2–7, 12, 17, 25, 27, 28, 37, 49, 62, 64–93, 97, 98, 101, 102, 106, 110, 111, 113–129, A1, B1
EEG	electroencephalography	25
fMRI	functional magnetic resonance imaging	25
HOS	higher order statistics	18

HRV	heart rate variability	65–70, 72, 74, A1, B1
ICA	independent component analysis	2, 4–22, 24–30, 33–35, 37–39, 41, 43, 45, 46, 48, 49, 51, 54, 55, 57, 58, 60, 61, 68, 75, 81–84, 88, 90, 92–94, 98, 101, 103, 105, 108, 110, 113–116, 118–129
ISP	intelligent signal processing	10, 46
MEG	magnetoencephalography	25
MLE	maximum likelihood estimator	21, 23
nifecgdb	Non-Invasive Fetal ECG Database	90, 122, B1
nsrdb	Normal Sinus Rhythm Database	88, 97, 106, B1
PCA	principal component analysis	12, 18, 26, 84
pdf	probably density function	17–19, 22
RMSE	root mean square error	27, 28, 31, 35, 39, 43, 45, 48, 51, 54, 57, 60, 96, 100, 105, 108, 113, 115–118, 121, 122, 126, 127, B1
RMSSD	root mean square of the successive differences	67
SNR	signal to noise ratio	9, 10, 46, 82

List of Symbols

Notation: Scalars are denoted by plain text, italics. Vectors are denoted by bold text.

$\mathbf{A}$	Mixing matrix in ICA	14, 16, 17, 22, 30, 39, 41, 46, 51, 52, 57, 58, 94, 110
D_β	Spectral fractal dimension	29, 69, 119
$\mathrm{g}(\cdot)$	Invertible function	19
$H(\mathbf{x})$	Shannon entropy of $\mathbf{x}$	21
$L(\mathbf{W})$	Likelihood function of $\mathbf{W}$	23
M	Number of mixture signals	14, 21, 23
N	Number of samples in a signal	23
n	Total number of RR-intervals in an ECG recording	67
$\mathbf{P}$	A permutation matrix	17
p_s	Joint pdf of the source signal vector$\mathbf{s}$	22
p_x	Joint pdf of the mixture signal vector$\mathbf{x}$	22
RR_i	The i^{th} RR-interval	67

$\mathbf{s}$	Source signal vector	13, 14, 16, 17, 22, 33, 94, 110
s_1	The first source signal	13, 31, 33, 39, 43, 46, 99, 100, 103
s_2	The second source signal	13, 31, 33, 39, 43, 46, 99, 103
$\mathbf{W}$	Unmixing matrix in ICA	14, 15, 17, 20–23
$\mathbf{W}^*$	Optimal unmixing matrix	22, 23
$\mathbf{x}$	Mixture signal vector	13–18, 22, 31, 33, 94, 110
x_1	The first mixture signal	13, 30, 33, 34, 39, 46, 51, 99, 103
x_2	The second mixture signal	13, 30, 33, 34, 39, 46, 51, 99, 103
$\mathbf{y}$	Extracted signal vector	14, 15, 18, 19, 33
y_1	The first extracted signal	31, 33, 39, 43, 51, 99, 100, 103
y_2	The second extracted signal	31, 33, 39, 51, 99, 103

Chapter 1

Introduction

"Distinguishing the signal from the noise requires both scientific knowledge and self-knowledge: the serenity to accept the things we cannot predict, the courage to predict the things we can, and the wisdom to know the difference."

– Nate Silver [Silv12]

In a crowded room with many people talking it is usually still possible to carry on a conversation without yelling. The human brain is remarkably adept at filtering out unwanted sounds while focusing on what is important. Ask a machine to do so and the results are typically disappointing, as anyone who has used automatic voice recognition programs in a noisy room knows. This is called the *cocktail party problem*, and it applies to more than just social gatherings. The ability to reliably separate signals is important to many fields including telecommunications, image processing, and medicine.

1.1 Problem Statement

It is rare in signal processing to work with isolated, purely individual signals. In most cases, physical signals are observed as part of a mixture of source signals and this mixture is typically further obscured by noise. This section introduces the specific *blind source separation* (BSS) problem of maternal-fetal heartbeat separation and why it is an important problem in signal processing. The signal separation technique of *independent component analysis* (ICA) is also introduced.

1.1.1 Motivation

In medicine the ability to measure the electrical activity of the body can provide useful diagnostic information. When electrical signals from the heart are measured it is called an *electrocardiogram* (ECG). Since ECGs can be non-invasive, they can be used to monitor a fetus before or during birth. However, the fetal heartbeat is obscured by the stronger maternal heartbeat, as well as by other electrical noise within the mother's body. The interfering signals make it difficult to reliably record a fetal ECG, and can falsely indicate problems leading to unnecessary medical intervention.

The ability to separate maternal and fetal ECGs accurately has social and economic impacts on the health care system. Maternal-fetal ECG recordings have been used by physicians since the middle of the 20th century and non-invasive, continuous electronic fetal heart monitoring is now standard practice during labour in North America [Buxt63][Wool90] [Davi93][Meni01]. The ECG recordings and heart rate data are used to assess fetal distress during birth, and are influential factors in the decision to perform an emergency Caesarean section [Lark58][Decl06][Clar07]. Of concern, the output from non-invasive fetal heart monitors can be difficult to interpret due to the relatively strong maternal ECG and noise interference from other biological electrical signals [Albe93][Rode08]. In practice it is not

always possible to reliably detect the features of the heart wave, including the main peaks that are used to calculate the fetal heart rate [Donk93]. Chapter 4 provides more information on the social and medical motivation behind this signal separation problem.

1.1.2 Problem Definition

Blind source separation addresses the problem of disentangling a signal of interest from a linear mixture when little to no information is known about the sources, although they are assumed to be independent. Often the signal mixing process is also unknown [Hyva01] [Ston04]. This is analogous to talking with someone in a noisy room, which is why it is often called the cocktail party problem. Even though there may be other conversations well within earshot, it is still possible for the people to focus their attention on one conversation.

While separating speech signals comes naturally to the human brain, it is still a challenging area of interest in other domains. Separating the fetal heartbeat from the other signals in a non-invasive ECG is a physical application of the cocktail party problem, and has become a benchmark problem in non-deterministic linear signal mixing [Hare60][Widr75][Vand87] [Schr96][Anan03]. Successes in separating maternal-fetal ECGs can also be applied to other blind source separation problems including audio, radio signal interference, and biological electrical signals [Robe01]. Blind source separation is not limited to one-dimensional signals, and can also be applied to signals such as images [Ston04].

1.1.3 Proposed Solution

Classic signal processing techniques, which rely on signal power, are not effective when the signal-to-noise ratio (SNR) is low. In these cases, other techniques such as adaptive or intelligent signal processing (ISP) must be used [Hayk01].

Independent component analysis (ICA) is a relatively new ISP technique to separate linearly mixed signals as recorded by more than one sensor simultaneously. ICA requires *higher order statistics* (HOS) in order to separate the signals, and the component signals must be statistically independent and non-Gaussian but need not be deterministic or probabilistic. These features make ICA particularly suited for blind separation of physical signals, which are typically non-deterministic and embedded in noise [DeLa00][Jung00][Pott01][Gadh06]. ICA exploits these assumptions to separate a signal mixture into statistically independent components of minimal mutual information or maximum non-Gaussianity [Robe01]. There is a wide range of applications for ICA, including telecommunications, stock market prediction, and biological electrical signal monitoring [Horn01]. The theory and conditions of ICA are described more fully in Chapter 2.

1.2 Thesis Formulation

This thesis is comprised of three parts: (i) an introduction to BSS problems, including the theory behind ICA and examples of its use; (ii) a study of ECG signal characteristics, as well as how maternal and fetal ECGs are combined; and (iii) an investigation into the maternal-fetal ECG problem as an application of ICA.

1.2.1 Thesis Statement

The goal of this thesis is to assess the suitability of ICA for ECG signals in the maternal-fetal ECG separation problem, and to investigate the effectiveness of ICA on this problem by identifying the successes and limitations of the approach.

1.2.2 Thesis Objectives

There are three main objectives in this thesis:

1. Study the effectiveness of ICA for blind source separation by:
   (a) implementing an ICA algorithm;
   (b) identifying metrics to evaluate and quantify the success of ICA at extracting source signals; and
   (c) experimenting with constructed deterministic and stochastic signals to test how source signal characteristics affect the outcome of ICA.
2. Study the characteristics of ECG signals and maternal-fetal ECG mixtures by:
   (a) implementing a model to simulate ECG signals;
   (b) obtaining recorded adult ECG signals and maternal-fetal ECG mixtures; and
   (c) simulating maternal-fetal ECG mixture signals.
3. Assess the effectiveness of ICA with the maternal-fetal ECG separation problem as a case study with both simulated and recorded data.

1.2.3 Research Questions

Blind source signal separation, and specifically maternal-fetal ECG separation, is a challenging problem in signal processing. The following is a list of the topics and research questions addressed in full or part in this thesis.

1. What are the limitations on the source signals that can be used with ICA?
   (a) What is the effect of source signal amplitude on ICA?

   (b) What is the effect of source signal frequency and phase differences on ICA?

2. How can the performance of ICA be justified and quantified?

   (a) How well do energy-based measures quantify the success of ICA?

   (b) What other measures can be used?

3. What are the characteristics of a healthy adult ECG?

   (a) How do the characteristics of adult and fetal ECGs differ?

   (b) How well do heart models accurately reflect ECG recordings?

4. What are the characteristics of maternal-fetal ECG mixtures?

   (a) How are the maternal and fetal ECGs combined?

   (b) How large is the maternal ECG relative to the fetal ECG?

5. What are the current methods of maternal-fetal ECG separation?

   (a) What features of the fetal ECG are important to capture accurately?

   (b) What are the specifications of current fetal ECG recording devices?

   (c) Can we reduce the number of "false positives" (i.e., the times when a pregnancy is considered unhealthy due to the fetal ECG, but it is actually healthy) by increasing the accuracy of fetal ECG separation?

6. How does the BSS problem apply to maternal-fetal ECG separation?

   (a) Is ICA applicable to the problem of maternal-fetal ECG separation?

   (b) To what other problems do BSS and ICA apply?

   (c) How well can ICA separate the maternal-fetal ECGs?

   (d) Can the fetal heart rate be accurately recovered in the signals separated by ICA?

1.3 Thesis Organization

This thesis addresses the problem of blind source separation, specifically in terms of maternal-fetal ECG separation, with ICA. An introduction to BSS problems is presented in Chapter 2, along with the fundamental concepts of ICA theory. This chapter also contains experiments to test and illustrate how well ICA works under a variety of conditions. Chapter 3 provides an overview of heart waveform characteristics and how ECG recordings are obtained. This chapter also serves as an introduction to current models that can be used to simulate heart signals. The discussion of ECGs is expanded into Chapter 4 with a review of past and current practices in maternal-fetal ECG separation. Experiments were designed to test ICA on simulated and recorded ECG data, and Chapter 5 presents the results. Five effectiveness measures are used in Chapter 6 to discuss the results of these experiments and to compare the trials with simulated ECG waveforms to those with recorded data. Chapter 7 provides a summary of the results, suggestions for future work, and concluding remarks.

Chapter 2

Independent Component Analysis (ICA)

It is rare in signal processing to work with clear, individual signals unless they are simulated. Most signals in the physical world are observed as part of a mixture that is combined with other signals and noise. In many cases it is costly or impossible to directly observe the signal of interest in isolation.

In this chapter, BSS problems and the most common techniques used to address them are introduced. One such technique, *independent component analysis* is described, beginning with a brief history and overview of current applications. The limitations of ICA are then discussed, including how they can be overcome for most BSS problems. This is followed by a high-level description of how ICA works and a brief introduction to the mathematics behind it.

The FastICA implementation of ICA was chosen for all of the experiments in this work. It is one of the more popular ICA algorithms due to its speed and reliability [Hyva00]. The

algorithm is introduced along with other ICA implementations and alternatives.

Finally, this chapter contains examples that show how ICA is used. These examples include experiments that highlight the effectiveness of ICA, as implemented in the FastICA algorithm, under a variety of conditions. ICA is explored using both deterministic and stochastic signals with varying attributes including their frequencies, amplitudes, and phases.

2.1 Blind Source Signal Separation Problems

BSS addresses the problem of disentangling independent signals from a mixture when little to no information is known about the sources. The mixing process is often unknown as well, increasing the difficulty of the problem.

These types of problems are also known as *cocktail party problems*, because they are analogous to talking with someone in a noisy room. Even though there may be other sounds such as talking, music, or movement, it is still possible for the human brain to sift through noise to carry on a conversation. While separating signals comes naturally to the human brain, it is still a difficult feat for machines. In every day life, this difficulty is shown in audio recognition programs such as Shazam or Siri that need clear, unobstructed input in order to work. Signal separation is a challenging area of research that has many applications in addition to speech recognition, including radio signal interference, biological electrical signal processing, and feature extraction.

Traditional signal processing relies on the relative power difference between the signal of interest and the surrounding noise. This ratio is called the *signal to noise ratio* (SNR) and is defined in Equation 2.1.

$$SNR = \frac{P_{signal}}{P_{noise}} \tag{2.1}$$

Signal processing that relies on the SNR uses the assumption that the signal of interest has more power than any of the other components. That is, anything other than the signal of interest is considered to be noise. However, this approach only works if there is a strong signal of interest and weak noise, therefore leading to a high SNR. What if the SNR is low, or if the signal is even weaker than the noise? In these cases, traditional signal processing techniques to filter out the noise are not feasible. Too much, or all, of the signal of interest would be filtered out in the process. Instead, new techniques that do not rely on the SNR must be used.

The term *intelligent signal processing* (ISP) describes a relatively new class of techniques that have been developed to address issues that can not be solved with classical signal processing techniques alone [Hayk01].

2.2 Overview of ICA

While ICA is still relatively new and not yet widely used, it is an active area of research due to the many BSS problems that exist in the world. In particular, ICA is of interest because, unlike many other BSS techniques, it does not assume that the source signals have Gaussian distributions. This is important since physical signals are typically not normally-distributed.

ICA can be applied to diverse fields: psychology, stock market prediction, radio interference, image processing, medicine and bioelectric signals, among others.

2.2.1 Historical Background of ICA

The concepts of ICA started to appear in the literature in the late 1980's. Christian Jutten and Jeanny Hérault are generally credited with first working on the technique, although under the less-specific label: "blind source separation" [Jutt91]. A number of other researchers at the time were also working on similar techniques, each from a slightly different angle [Como94].

In 1984, the term "independent component analysis" was first introduced by Pierre Comon in an article titled "Independent component analysis, A new concept?" published in *Signal Processing* [Como94]. This paper was also one of the first to present a thorough explanation and mathematical definition of ICA, including an algorithm to implement ICA in polynomial time.

2.2.2 The Effects of Signal Mixing

When non-Gaussian signals are linearly combined to produce new signals, any given mixture differs from the source signals in three significant ways [Ston04]:

1. **Independence:** while the source signals are independent, the mixtures are all dependent because they were formed from the same underlying signals.
2. **Gaussianity:** the sum of any non-Gaussian signals will produce a signal that has a more Gaussian histogram. Conversely, the individual original signals will be less Gaussian than any possible linear mixture of them.
3. **Complexity:** the sum of any two signals will be more complex than any of its constituent parts. The complexity can be measured as entropy.

ICA and related techniques exploit these effects to separate the mixtures into signal components. In particular, ICA focuses on the effects related to independence and Gaussianity by separating signals into maximally independent components or minimally Gaussian components.

2.2.3 Relationship to Principal Component Analysis (PCA)

ICA was built upon the earlier signal separation technique of *principal component analysis* (PCA), which is a less constrained signal separation technique. PCA relies on the property that the original source signals are not correlated, so it extracts uncorrelated signals from the mixtures. These extracted signals are called *principal components*.

In contrast, ICA focuses on separating out signals based on the property of independence. This is a more rigid requirement, as independence implies that the signals are not correlated. However, this assertion does not hold in the reverse: uncorrelation does not imply independence. In ICA the source signals are assumed to have no dependence, so independent components are extracted from the mixtures. This is a valid assumption in most physical world separation problems, including the maternal-fetal ECG separation problem [Pott08].

2.3 Description of ICA

ICA works by separating signal mixtures into two or more independent components. Typically the source signals are time-varying signals that are sampled over some finite period of time. If the mixture is composed of two source signals with N samples each, the source

signals can be defined as s_1 (Equation 2.2a) and s_2 (Equation 2.2b).

$$s_1 = (s_1^1, s_1^2, s_1^3 \dots s_1^N) \tag{2.2a}$$

$$s_2 = (s_2^1, s_2^2, s_2^3 \dots s_2^N) \tag{2.2b}$$

For brevity, any number of these source signals can be combined into a *source vector*. For the example with two source signals, the source vector $\mathbf{s}$ is defined in Equation 2.3.

$$\mathbf{s} = \begin{bmatrix} s_1 \\ \\ s_2 \end{bmatrix} \tag{2.3}$$

In general, for any number of source signals the vectors and matrices in the ICA mixing model will increase in size. The signals in $\mathbf{s}$ are linearly combined with no additive constant to form mixture signals, x_1 and x_2. The assumption that the source signals have been linearly combined is a crucial idea in ICA. This means that x_1 and x_2 are weighted sums of the source signals, as shown in Equations 2.4a and 2.4b. The scalar multipliers a, b, c and d may be positive or negative, and any or none of them may be equal.

$$x_1 = as_1 + bs_2 \tag{2.4a}$$

$$x_2 = cs_1 + ds_2 \tag{2.4b}$$

The mixture signals form a new vector $\mathbf{x}$ (Equation 2.5), which will be the input to the ICA algorithm. In BSS these are the only signals that we are able to observe, as the original signals are inaccessible. Therefore $\mathbf{x}$ may also be called the recorded or observed signals.

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} as_1 + bs_2 \\ cs_1 + ds_2 \end{bmatrix} \tag{2.5}$$

Typically this mixing process is described in terms of matrix multiplication, rather than calculating each element of $\mathbf{x}$ separately. To do this, an $M \times M$ *mixing matrix* $\mathbf{A}$ is defined that contains all of the scalar multipliers used to form the mixture signals. The 2×2 matrix $\mathbf{A}$ is defined in Equation 2.6 for the current example.

$$\mathbf{A} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \tag{2.6}$$

All of this leads to the main equation that characterizes the signal mixing model used in ICA, given in Equation 2.7.

$$\mathbf{x} = \mathbf{A}\mathbf{s} \tag{2.7}$$

Alternatively, the signal mixing model can be expressed with time varying signals explicitly, as in Equation 2.8 [Gian99]. In this equation the $\mathbf{a_i}$ components are the columns of mixing matrix $\mathbf{A}$.

$$\mathbf{x}(t) = \mathbf{A}\mathbf{s}(t) = \sum_{i=1}^{M} \mathbf{a}_i s_i(t) \tag{2.8}$$

ICA attempts to recover the mixing matrix $\mathbf{A}$ and the original signals $\mathbf{s}$ when only the mixed signals $\mathbf{x}$ are known. To do this, it attempts to find an *unmixing matrix* $\mathbf{W}$ such that:

$$\mathbf{y} = \mathbf{W}\mathbf{x} \tag{2.9}$$

where $\mathbf{y}$ is the vector of extracted signals that are expected to be equal to $\mathbf{s}$.

As previously stated, ICA separates $\mathbf{x}$ into component signals of:

1. minimum Gaussianity, or
2. maximum statistical independence.

Finding these signals can be thought of as searching the set of all possible component signals to find the subset that match this criteria. This search can be done in a brute force way, but one of the active areas of research within ICA is how to improve the efficiency of the search and make it more intelligent.

The particular search algorithm used to find **y** and **W** impacts the effectiveness of ICA in finding a solution. For example, a gradient-based search can become "stuck" at local minima or maxima, rather than converging to the best solution.

2.3.1 Considerations of ICA

ICA is designed to work on signal mixtures when there is no a priori information about the source signals. However, to do this there are a few assumptions on the source signals and mixing process that influence which problems ICA can be applied to properly. As well, ICA comes with some limitations that are inherent to the technique and must be taken into account when selecting a BSS method [Hyva00][Ston04][Gadh04].

2.3.1.1 Constraints and Assumptions

ICA makes four assumptions about the source signals and mixing process that constrain how it can be used:

1. **Independence:** All source signals are assumed to be statistically independent.
2. **Gaussianity:** ICA only works on non-Gaussian source signals, although it has been shown to work with a maximum of one Gaussian signal [Hyva00]. Many physical world signals of interest are non-Gaussian, so this assumption does not greatly limit ICA.
3. **Number of recorded signals:** The number of recorded signals in **x** must be greater

than or equal to the number of source signals that ICA is attempting to extract.

4. **Linear mixing:** The ICA technique assumes that the source signals have been linearly mixed to produce the recorded signals. This means that each of the source signals can only be multiplied by a scalar and added together to produce the recorded signals. The recorded signals can not be the product of source signals under this assumption.

These four constraints are critical to the proper use and success of ICA. The assumptions are valid for many practical BSS problems, making ICA an important signal separation technique. Applications meeting these criteria are discussed further in Section 2.3.4.

2.3.1.2 Limitations of ICA

Even when the constraints outlined in Section 2.3.1.1 are met, there are still some limitations on the signals that ICA can extract from mixtures. In particular, ICA is unable to determine the original energies or order of the source signals [Hyva00], and does not take noise into account [Ston04].

Using the ICA model, shown in Equation 2.7, it is not possible to recover the **energies** or variances of the original signals. Furthermore, the original signals may be extracted upside-down. This is because the recorded signals vector, $\mathbf{x}$, is the product of two unknown terms: the source signals vector, $\mathbf{s}$, and the mixing matrix, $\mathbf{A}$. Any scalar multiple of one of the sources could be cancelled out by the mixing matrix without otherwise affecting the signal. So most ICA methods assume that the source signals have a variance of one. The extracted signals may also be upside-down relative to the source signals. This is because both the signal and a column of the mixing matrix could be multiplied by a factor of -1 without affecting the model [Hyva00].

ICA is also unable to recover the **original ordering** of the source signals. As with the energies, this is because both the source vector, $\mathbf{s}$ and the mixing matrix, $\mathbf{A}$, are unknown. So, the order of the signals in $\mathbf{s}$ could be changed and still produce the same vector $\mathbf{x}$ with a corresponding change in the mixing matrix. For example, if a permutation matrix is defined as $\mathbf{P}$, then $\mathbf{x} = (\mathbf{A}\mathbf{P}^{-1})(\mathbf{P}\mathbf{s})$ is equivalent to Equation 2.7 where $\mathbf{A}\mathbf{P}^{-1}$ is the new mixing matrix and $\mathbf{P}\mathbf{s}$ is the new vector of rearranged source signals [Hyva00].

The final major limitation of ICA is that it does not take additive **noise** into account [Pott02]. The ICA model modified to explicitly include additive noise could be $\mathbf{x} = \mathbf{A}\mathbf{s} + \mathbf{n}$, where $\mathbf{n}$ is the noise vector. However, this new model is not compatible with the unmixing process in ICA, so additive noise is typically ignored for simplicity [Como94].

Fortunately, for many BSS problems the ambiguity in the signal energy and orientation is inconsequential, or can be resolved in post-processing. For example, in the maternal-fetal ECG separation problem the extracted signal can simply be flipped if it is extracted upside-down because the general ECG waveform is already known. Furthermore, the assumption that the additive noise is zero, or minimal, is also acceptable in many BSS problems. ICA is not a good signal separation candidate for problems in which these limitations can not be ignored.

2.3.2 Principles of ICA

The fundamental concept behind ICA is that it takes mixtures of independent signals and separates them into mutually independent components by finding an unmixing matrix, $\mathbf{W}$. Intuitively, if two signals are independent that means that they do not influence each other. That is, any knowledge about one of the signals provides no information about the other. This means that the joint *probably density function* (pdf) of independent signals is

equal to the product of their individual pdfs, as shown in Equation 2.10.

$$p_{X,Y}(x,y) = p_X(x)p_Y(y) \tag{2.10}$$

Signal independence is a more stringent than constraining the signals to be uncorrelated, as in PCA. Two signals are uncorrelated if their *covariance* is equal to zero, as defined in Equation 2.11.

$$E[X,Y] - E[X]E[Y] = 0 \tag{2.11}$$

Signals that are independent are also uncorrelated, but the reverse does not hold true. Uncorrelated signals are not guaranteed to be independent because correlation does not take *higher order statistics* (HOS) into account.

2.3.2.1 Infomax and ICA

Independence can not be measured in and of itself, so ICA uses other properties related to statistical independence. The *infomax* principle was first described by Ralph Linsker [Lins89] and later applied to ICA by Anthony Bell [Bell95]. Infomax shows that signal independence can be maximized in ICA by finding an unmixing matrix that:

1. maximizes the entropy of the $\mathbf{y}$ components, or
2. maximizes the mutual information between the extracted signals, $\mathbf{y}$, and the signal mixtures, $\mathbf{x}$.

In information theory *entropy* describes the "randomness" of a signal, or how much a priori uncertainty there is about the signal. A signal with a uniform pdf has maximum entropy, because every value of the signal is equally likely. Further, a set of signals with a

uniform joint pdf will have maximum joint entropy [Ston04]. Entropy is most commonly defined as the Shannon entropy, which is defined in Equation 2.12 for signals of known pdf [Cove06].

$$H(X) = -\sum_{t=1}^{N} p(x_t) \log p(x_t) \tag{2.12}$$

The entropy can also be calculated from the data points of an observed signal, as in Equation 2.13 [Ston04].

$$H(X) = -\frac{1}{N} \sum_{t=1}^{N} \ln p_X(X^t) \tag{2.13}$$

Mutual information is a measure of how similar two variables or signals are, and is related to entropy as shown in Equation 2.14. The term $H(Y|X)$ is the entropy of Y when X is fully known.

$$I(Y;X) = H(Y) - H(Y|X) \tag{2.14}$$

Due to this relationship, maximizing the mutual information $I(Y;X)$ is the same as just maximizing the entropy of the observed signals $H(Y)$ [Bell95]. So maximizing the entropy of the ICA output is sufficient to extract mutually independent signals without needing to know $H(Y|X)$.

As stated previously, a signal with a uniform pdf has maximum entropy. So to maximize the entropy of the ICA output, $\mathbf{y}$, a new variable $\mathbf{Y} = \mathrm{g}(\mathbf{y})$ can be defined. The function $\mathrm{g}(\cdot)$ is the *cumulative distribution function* (cdf) of $\mathbf{y}$, making the distribution of $\mathbf{Y}$ uniform according to the *probability integral transform* [ORei73]. An invertible function, such as $\mathrm{g}(\cdot)$, applied to signals with maximum entropy will produce signals that are also mutually

independent [Ston04].

Using the definition of entropy of an observed signal from Equation 2.13, the entropy of $\mathbf{Y} = \mathrm{g}(\mathbf{y})$ is defined in Equation 2.15.

$$H(\mathbf{Y}) = -\frac{1}{N}\sum_{t=1}^{N} \ln p_Y(\mathbf{Y}^t) \tag{2.15}$$

To ensure independence, $H(\mathbf{Y})$ needs to be maximized. It can be shown that the probabilities of the $\mathbf{Y}$ output are related to the signals extracted by ICA and the source signals by Equation 2.16 [Ston04].

$$p_Y(\mathbf{Y}^t) = \frac{p_y(\mathbf{y}^t)}{p_s(\mathbf{y}^t)} \tag{2.16}$$

Recall that the signals extracted by ICA are related to the source signals through unmixing matrix $\mathbf{W}$ with the relation $\mathbf{y} = \mathbf{W}\mathbf{x}$. Therefore, the probabilities of the signals extracted by ICA are also directly related to the source signals. This relationship is shown in Equation 2.17.

$$p_y(\mathbf{y}^t) = \frac{p_x(\mathbf{x}^t)}{|\mathbf{W}|} \tag{2.17}$$

Substituting the results of Equations 2.16 and 2.17 back into Equation 2.15 produces Equation 2.18a. Rearranging and separating the right-hand side of Equation 2.18a, shows

that one of the terms is simply the entropy of the source signals.

$$H(\mathbf{Y}) = -\frac{1}{N}\sum_{t=1}^{N} \ln \frac{p_x(\mathbf{x}^t)}{|\mathbf{W}| p_s(\mathbf{y}^t)} \tag{2.18a}$$

$$= \frac{1}{N}\sum_{t=1}^{N} \ln p_s(\mathbf{y}^t) + \ln|\mathbf{W}| - \frac{1}{N}\sum_{t=1}^{N} \ln p_x(\mathbf{x}^t) \tag{2.18b}$$

$$= \frac{1}{N}\sum_{t=1}^{N} \ln p_s(\mathbf{y}^t) + \ln|\mathbf{W}| - H(\mathbf{x}) \tag{2.18c}$$

Since the mixing matrix $\mathbf{W}$ that maximizes the entropy of the extracted signals is what is being sought, and the mixing matrix can not affect the source signals, $H(\mathbf{x})$ can be ignored. The entropy of the extracted signals can now be defined more simply as in Equation 2.19.

$$h(\mathbf{Y}) = \frac{1}{N}\sum_{t=1}^{N} \ln p_s(\mathbf{y}^t) + \ln|\mathbf{W}| \tag{2.19}$$

If the M source signals are assumed to be independent, then the function can be further simplified as in Equation 2.20.

$$h(\mathbf{Y}) = \frac{1}{N}\sum_{i=1}^{M}\sum_{t=1}^{N} \ln p_s(y_i^t) + \ln|\mathbf{W}| \tag{2.20}$$

The optimal unmixing matrix will maximize the entropy of $\mathbf{Y} = \mathrm{g}(\mathbf{y})$, therefore ensuring that the signals in $\mathbf{Y}$ are independent. This means that the extracted signals $\mathbf{y} = \mathrm{g}^{-1}(\mathbf{Y})$ are also independent.

2.3.2.2 Maximum Likelihood Estimation of ICA

As shown by [Ston04] and others, the above method for maximizing the entropy of the ICA output is a *maximum likelihood estimator* (MLE) for the original signals. The observed

signals in ICA are considered to be the product of the source signal vector with a mixing matrix, or $\mathbf{x} = \mathbf{As}$ (Equation 2.7). Thus the pdfs of $\mathbf{s}$ and $\mathbf{x}$ are directly related. If the source signals have joint pdf p_s and the observed signals have pdf p_x, then they are related as shown in Equation 2.21.

$$p_x(\mathbf{x}) = p_s(\mathbf{s}) \left| \frac{\delta \mathbf{s}}{\delta \mathbf{x}} \right| \tag{2.21}$$

If the mixing matrix, $\mathbf{A}$, was known then the source signals could be obtained directly from the mixtures by using the inverse of $\mathbf{A}$ as an unmixing matrix, as described in Equation 2.22.

$$\mathbf{s} = \mathbf{A}^{-1}\mathbf{x} \tag{2.22}$$

However, in blind source separation problems the mixing matrix is, of course, not known. Instead, the unmixing matrix is defined as $\mathbf{W}$, and the unmixing matrix that is equal to the inverse of $\mathbf{A}$ is then called the *optimal unmixing matrix*, $\mathbf{W}^*$. The optimal unmixing matrix is directly related to the source and mixture signals as shown in Equation 2.23.

$$|\mathbf{W}^*| = \left| \frac{\delta \mathbf{s}}{\delta \mathbf{x}} \right| \tag{2.23}$$

This optimal unmixing matrix can then be substituted into Equation 2.21, to produce Equation 2.24.

$$p_x(\mathbf{x}) = p_s(\mathbf{s}) \, |\mathbf{W}^*| \tag{2.24}$$

In practice the optimal unmixing matrix is unknown and the estimated unmixing matrix $\mathbf{W}$ is used to obtain estimated signals $\mathbf{y} = \mathbf{Wx}$ (Equation 2.9). The estimated components

are then substituted into Equation 2.24 to obtain:

$$p_x(\mathbf{x}) = p_s(\mathbf{y})|\mathbf{W}| \tag{2.25a}$$

$$= p_s(\mathbf{W}\mathbf{x})|\mathbf{W}| \tag{2.25b}$$

$$= L(\mathbf{W}) \tag{2.25c}$$

$L(\mathbf{W})$ in Equation 2.25 is called the *likelihood function* of $\mathbf{W}$. The unmixing matrix that maximizes $L(\mathbf{W})$ is the *maximum likelihood estimator* of $\mathbf{W}^*$. The likelihood function can be simplified by assuming that the M source signals are independent (Equation 2.26b), and then further simplified by assuming that the N samples in each source signal are also independent (Equation 2.26c).

$$L(\mathbf{W}) = p_s(\mathbf{W}\mathbf{x})|\mathbf{W}| \tag{2.26a}$$

$$= \prod_{i=1}^{M} p_s(\mathbf{w}_i^T\mathbf{x})|\mathbf{W}| \tag{2.26b}$$

$$= \prod_{i=1}^{M}\prod_{t=1}^{N} p_s(\mathbf{w}_i^T\mathbf{x}^t)|\mathbf{W}| \tag{2.26c}$$

For convenience, the likelihood function can be converted into the *log likelihood function* by taking the logarithm of both sides (Equation 2.27a). The likelihood function can also divided by N without affecting the MLE (Equation 2.27b). Doing so produces Equation 2.27c, which is the same as the right hand side of the equation from the infomax method

(Equation 2.20).

$$\ln L(\mathbf{W}) = \sum_{i=1}^{M}\sum_{t=1}^{N} \ln p_s(\mathbf{w}_i^T \mathbf{x}^t) + N \ln |\mathbf{W}| \tag{2.27a}$$

$$\frac{1}{N} \ln L(\mathbf{W}) = \frac{1}{N} \sum_{i=1}^{M}\sum_{t=1}^{N} \ln p_s(\mathbf{w}_i^T \mathbf{x}^t) + \ln |\mathbf{W}| \tag{2.27b}$$

$$= \frac{1}{N} \sum_{i=1}^{M}\sum_{t=1}^{N} \ln p_s(y_i^t) + \ln |\mathbf{W}| \tag{2.27c}$$

Therefore, maximizing the entropy of the extracted signals in the infomax method produces a maximum likelihood estimator.

2.3.3 FastICA and Other ICA Algorithms

One of the most widely used implementations of ICA is the *FastICA* algorithm introduced by Aapo Hyvärinen in 1999 [Hyva99]. FastICA was originally introduced in the context of neural network research. However, the implementation has been readily adopted in general ICA research in large part due to its computational efficiency [Gian98]. This algorithm earns its name by having a quadratic convergence, which is much faster compared to the linear convergence of the previous gradient descent-based methods. In some cases the FastICA algorithm has even been shown to have a cubic convergence [Hyva00]. As well, the algorithm does not require a large amount of memory in order to run.

Another aspect that sets FastICA apart from other ICA algorithms is that it does not require any prior knowledge or an estimate of probability distributions of the source signals. As well, while the independent components are typically estimated all at once, FastICA provides an option to allow for the extraction of one independent component at a time.

FastICA is available as an open source algorithm with code in a variety of formats including Matlab, C++, and Python [Hyva13]. The current Matlab version was updated in 2013 and also includes an optional graphical user interface. FastICA is used in the experiments contained in this thesis and was chosen due to the advantages listed above. In this thesis the current Matlab version of FastICA is used without the graphical user interface.

Other ICA algorithms include: JADE [Card93], MRMI [Hild01], RADICAL [Lear03], and MILICA [Stog04], among others [Bell95][Miho02].

2.3.4 Applications of ICA

BSS problems that are suitable for ICA abound in a wide range of science and engineering fields. The following is a list of some of areas where ICA is of interest:

- **Speech processing:** This is the classic problem in ICA, and is also often called the "cocktail party problem" as introduced in Section 2.1 [Hayk05].
- **Medicine:** ICA has been applied to a variety of biomedical signal mixtures including electrical and medical scans. Some examples include *electroencephalography* (EEG) and *magnetoencephalography* (MEG) recordings of brain waves [Viga00][Vale14]; *functional magnetic resonance imaging* (fMRI) scans [Calh03]; and maternal-fetal ECGs.
- **Telecommunications:** Radio signal interference is a concern that has been addressed by ICA [Ston04][Besi14], including CDMA (Code Division Multiple Access) methods [Rist99] and array processing techniques such as beam forming or spatial filtering [Card93].
- **Financial data:** Recently ICA has been applied to financial data to look for underlying patterns, including stock market predictions [Oja00][Hyva00].

- **Image processing:** The principles of ICA are not limited to one-dimensional signals. They can also be applied to multidimensional signals such as images, often as a noise reduction technique [Karh97][Hyva00].

2.3.5 Alternatives to ICA

There are many other BSS techniques that have influenced ICA. Some may be used together with ICA as a pre-processing step, while others are currently competing areas of research [Hyva00][Ston04][Pott08]. The most common alternatives to ICA include:

- principal component analysis (PCA),
- factor analysis (FA),
- projection pursuit,
- wavelets,
- blind deconvolution, and
- artificial neural networks (ANNs).

2.3.6 Measures of the Effectiveness of ICA

There are multiple methods that can be used to evaluate how closely the signals extracted by ICA resemble the original source signals. These methods include measures based on energy, information theory, and fractals. Some measures are only applicable in simulated experiments, when the source signals are known and can be used for comparison. In true BSS problems when the source signals are not accessible, the extracted signals can not be compared directly to the sources.

In simulation experiments the limitations of ICA addressed in Section 2.3.1.2 can be accounted for because the source signals are known. A signal extracted upside-down can be inverted, and the amplitude can be adjusted after comparing the ranges of the extracted and source signals. This post-processing allows for a direct comparison of the independent components found by ICA to the source signals.

2.3.6.1 Visual Comparison

A common, intuitive approach to assess the success of any signal separation technique is to visually compare the source signals to the extracted signals. The source and extracted signals can be compared by looking at their amplitudes or normalized histograms. This method is only possible if the source signals, or good estimations of the source signals, are known. Thus, it is most useful with simulated data with known source signals to test a separation technique such as ICA, rather than to verify the results of true blind source separation problems.

However, in experimental blind source separation problems where the sources are unknown the visual comparison method can be used as a negative test. With a mixture of unknown ECG signals, for example, the extracted signals would be expected to exhibit the periodic waves that are characteristic of heart activity. If the extracted signals do not have these periodic waves then the separation method must not have been successful. The reverse does not hold true, however, as periodic waves in the extracted signals do not guarantee that the separation was successful.

2.3.6.2 Energy-based Measures

One of the most common tools to quantitatively compare expected and observed values is the *root mean square error* (RMSE). This measure is only applicable if the original data

is known or can be estimated, and is therefore applicable only to simulated experiments of maternal-fetal ECGs. The RMSE for the i^{th} signal in the ICA model is defined in Equation 2.28.

$$RMSE_i = \sqrt{\left(\frac{\sum_{t=1}^{N}(s_i^t - y_i^t)^2}{N}\right)} \tag{2.28}$$

The RMSE is both simple to calculate and intuitive, making it a useful measure. However, it can be misleading as it only takes the energy of the signals into account.

2.3.6.3 Information-based Measures

If the distributions of the original signals are known or can be estimated, then the *entropy* can be calculated and compared to that of the extracted signals. In information theory, entropy is a measure of the "degree of information that the observation of the variables gives" [Hyva00]. Signals that are more "random" have a higher entropy, as there is less a priori certainty about their values. There are other definitions of information that can be used for the same purpose of evaluating complexity, however entropy is one of the simplest and is therefore widely used [Kins10].

The most common form of entropy is the *Shannon entropy*, which is a specific form of the more generalized Rényi entropy [Cove06]. Shannon entropy is defined in Equation 2.29 for a random signal, X.

$$H(X) = -\sum_{i=1}^{N} p(x_i) \log p(x_i) \tag{2.29}$$

Entropy is based on the probability distribution of the signal, where $p(x_i)$ is the probability that $X = x_i$.

2.3.6.4 Fractal-based Measures

If the original signals are known or can be estimated, then the fractal dimension can be used to validate ICA. Fractal analysis belongs to a larger class of analytical tools that is termed polyscale analysis [Kins05]. For time series, such as biomedical signals, the *spectral fractal dimension* (D_β) can be used to reveal underlying information that is not available in single-scale analysis [Kins13]. The spectral fractal dimension for a stationary time series is defined in Equation 2.30, where β is the *spectral exponent*.

$$D_\beta = 1 + \frac{3 - \beta}{2} \tag{2.30}$$

The spectral exponent, β, is the slope of the power spectrum density of signal, or a stationary section of the signal, typically after it has been averaged to remove any Spanish moss. The spectral fractal dimension is applicable to self-affine time series, unlike other fractal dimensions which are for signals that are self-similar. This measure is further described in Section 3.2.3.

2.4 ICA Experiments

The following section presents experiments to test the effectiveness of ICA with deterministic and stochastic signals of varying features. All experiments were conducted using the open source FastICA algorithm developed for Matlab [Hyva99]. Simple deterministic signals are used, such as sine and square waves, as well as more complex signals like the Weierstrass function. Experiments are also performed with stochastic signals, including one with a normally-distributed source signal.

ICA was successful at separating the signal mixtures in nearly all cases. Notable ex-

ceptions were the sine waves with phase differences and the Weierstrass functions. In the former, the phases of the source signals were not recoverable but the shapes and frequencies of the signals were preserved. In the case of the Weierstrass functions, ICA was only able to separate a maximum of one Weierstrass signal successfully at a time. Matlab code to compare the extracted signals to the source signals is included in Appendix A.

2.4.1 Deterministic Signals

The first ICA experiments shown here were conducted using determinist waves: sine waves, square waves, and Weiestrass functions. The sine wave experiments examine the effect of phase, frequency, and amplitude differences in the source signals on the performance of ICA. The square and Weierstrass waves were also included to investigate the effect of singularities on ICA, as well as to demonstrate ICA with signals of differing complexity.

2.4.1.1 Two In-Phase Sine Waves

The first experiment to test the FastICA algorithm used two sine waves. Both signals had the same phase and amplitude, but differed in frequency as defined in Equations 2.31a and 2.31b where $t = [0 : 0.01 : 100 - 0.01]$. Figure 2.1 shows the two signals graphically.

$$s_1(t) = 2\sin(2\pi \times 0.1t) \tag{2.31a}$$

$$s_2(t) = 2\sin(2\pi \times 0.03t) \tag{2.31b}$$

The signals were then combined using the mixing matrix $\mathbf{A}$, defined in Equation 2.32 The result of this mixing process is shown in Figure 2.2. Neither mixed signal, x_1 or x_2,

resembles a sine wave any longer.

$$A = \begin{bmatrix} 2 & 3 \\ 2 & 1 \end{bmatrix} \tag{2.32}$$

After the FastICA algorithm was applied to $\mathbf{x}$, the original signals were recovered with only a small error. A visual comparison of the recovered and original signals is given in Figure 2.3. The RMSE for y_1 and s_1 was $RMSE_1 = 4.1145 \times 10^{-15}$ and the for y_2 and s_2 it was $RMSE_2 = 5.1159 \times 10^{-15}$. Both signals were successfully recovered by the FastICA algorithm, and the phase of each was preserved. In both cases the amplitude was adjusted as described in Section 2.3.6 in order to properly compare the extracted and original signals.

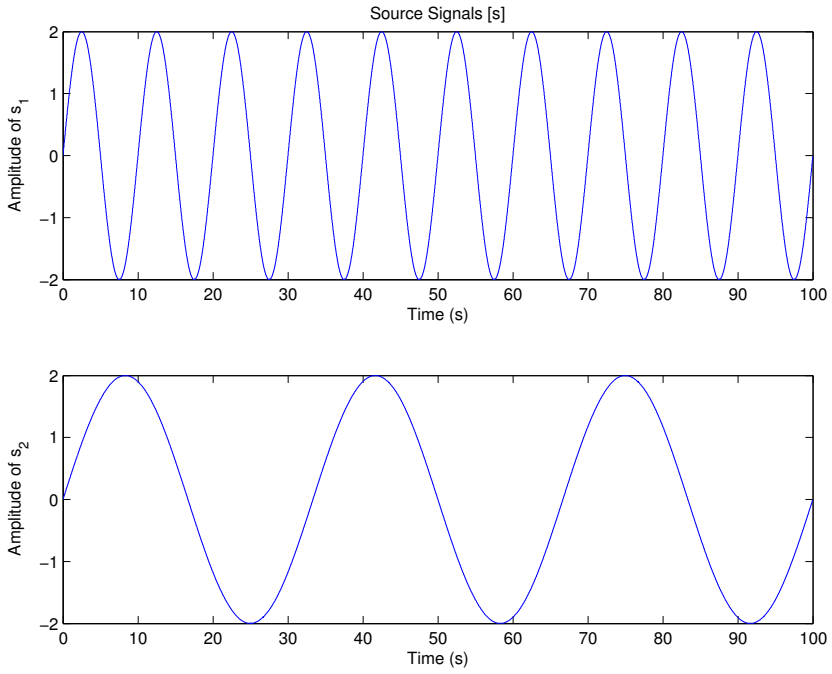

Fig. 2.1: The two sine waves defined in Equations 2.31a and 2.31b before mixing.

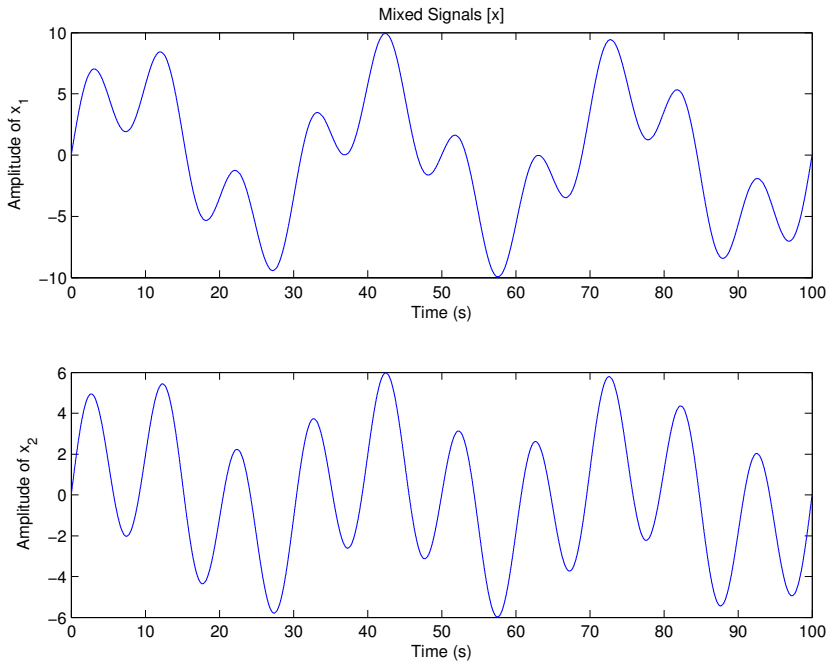

Fig. 2.2: The two sine waves after mixing.

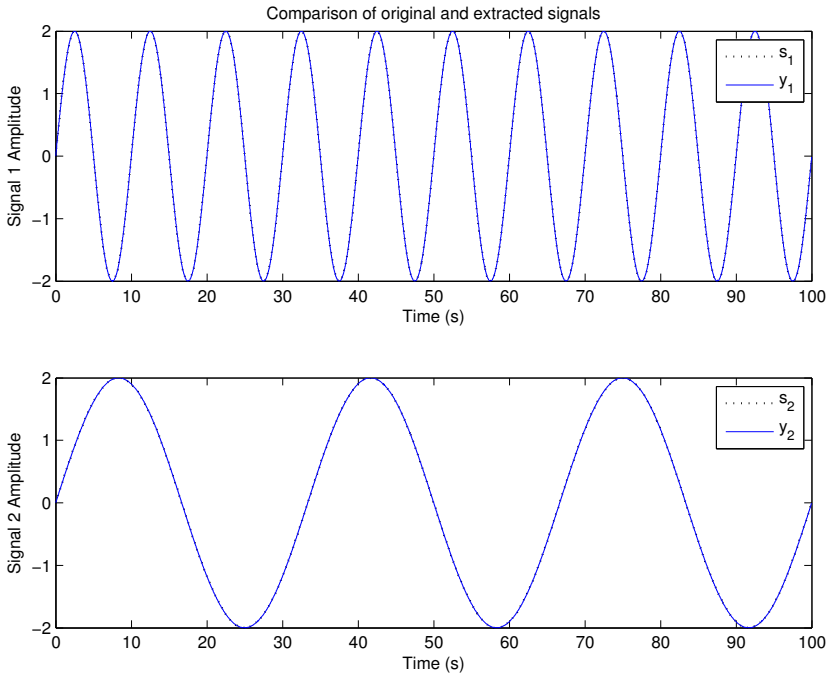

Fig. 2.3: The two sine waves after separation by the FastICA algorithm compared to the source signals.

Note: one of the signals was recovered up-side-down, and both were multiplied by a factor of $\sqrt{0.5}$, as shown in the estimated mixing matrix:

$$estA = \begin{bmatrix} -2.8284 & 4.2426 \\ -2.8284 & 1.4142 \end{bmatrix} \tag{2.33}$$

The change in entropy for the first signal was -1.4365×10^{-6}, and there was no change in entropy for the second signal.

The success of ICA in this example is also shown in the histograms of $\mathbf{s}$, $\mathbf{x}$, and $\mathbf{y}$. Both s_1 and s_2 are clearly non-Gaussian, but when mixed into x_1 and x_2 the distributions have a more Gaussian shape. The independent components y_1 and y_2 have the same distributions as their respective source signals.

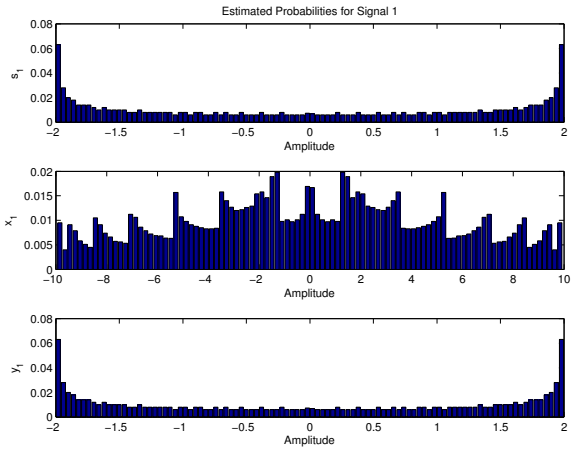

Fig. 2.4: Estimated probabilities of s_1, x_1, and y_1 for the first of the two sine waves.

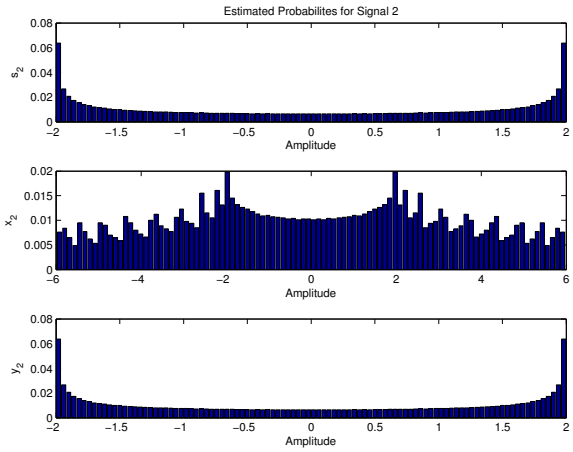

Fig. 2.5: Estimated probabilities of s_2, x_2, and y_2 for the second of two sine waves.

2.4.1.2 Square and Sine Waves

The next experiment shows how ICA performs with a square wave. The sharp corners and flat sections of the square wave, much different from the smooth and continuous sine wave, make the signal an interesting candidate for signal separation. In this experiment, a square wave was mixed with a sine wave using the same mixing matrix in Equation 2.32. The original signals are shown in Figure 2.6 and the mixed signals are shown in Figure 2.7. Neither mixed signal resembles either of the original signals, though both x_1 and x_2 exhibit smooth sections and sharp peaks.

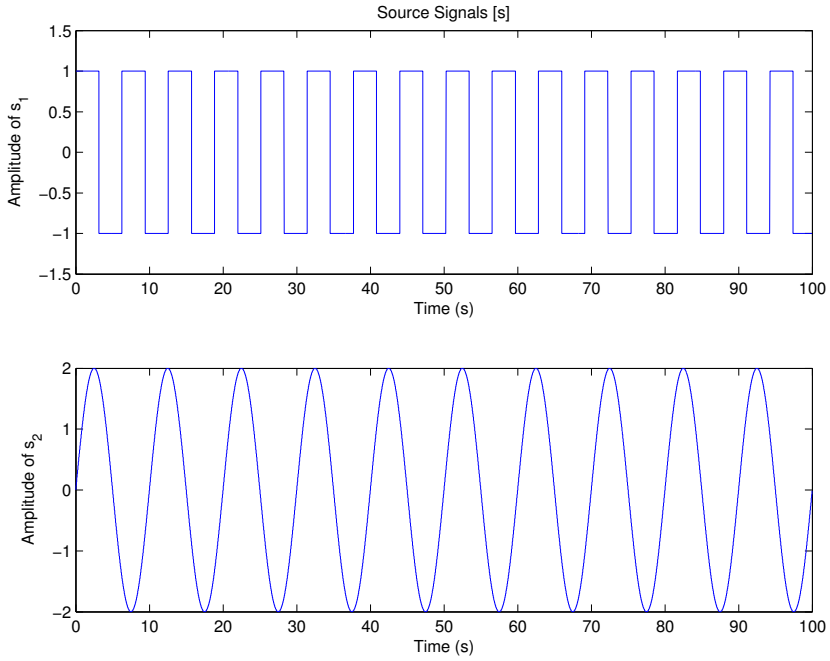

Fig. 2.6: The original square and sine waves before mixing.

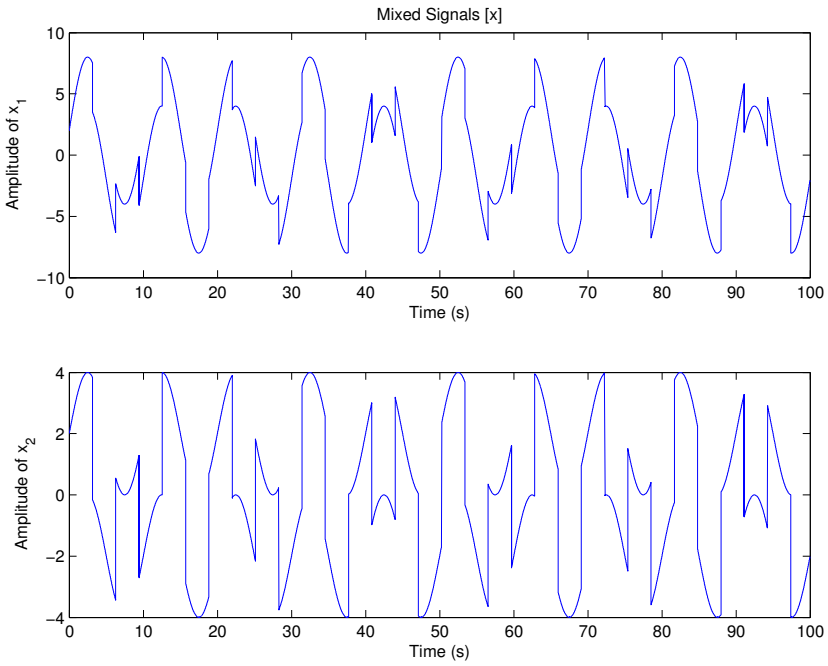

Fig. 2.7: The square and sine waves after mixing.

The signals extracted using the FastICA algorithm are shown in Figure 2.8. Both waves are recognizable, but in some trials the flat surfaces of the square wave were not preserved due to the statistical nature of ICA. Figure 2.9 shows a comparison of the extracted signals to their original counterparts when the flat surfaces were preserved. The RMSE was $rmse_1 = 4.6154 \times 10^{-12}$ between y_1 and s_1, while the error was $rmse_2 = 0.0173$ between y_2 and s_2. Compared to the previous experiment using only sine waves, the RMSE is 3-13 orders of magnitude greater here.

There was no change in entropy for the square wave, but the change in entropy for the sine wave was -0.034683.

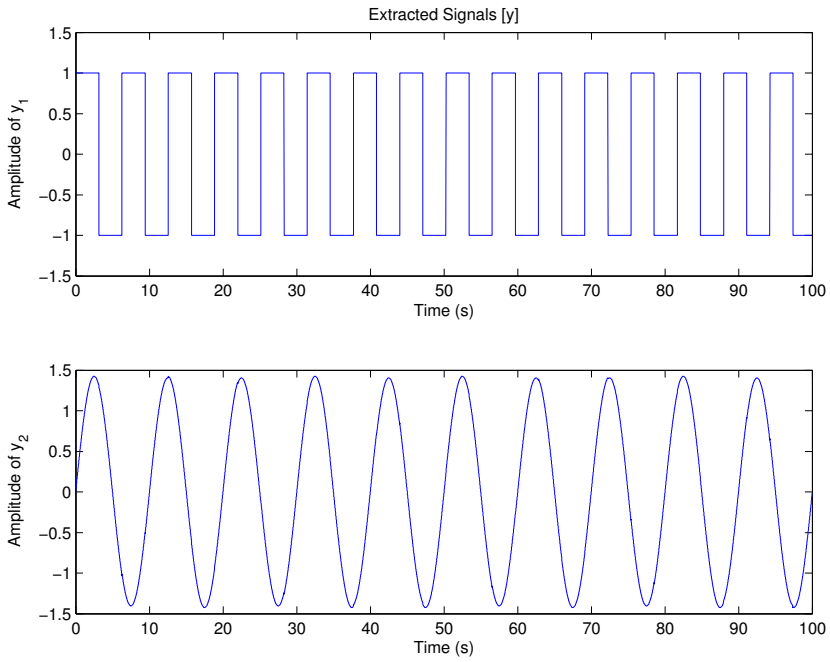

Fig. 2.8: The square wave, y_1, and sine wave, y_2, extracted by FastICA.

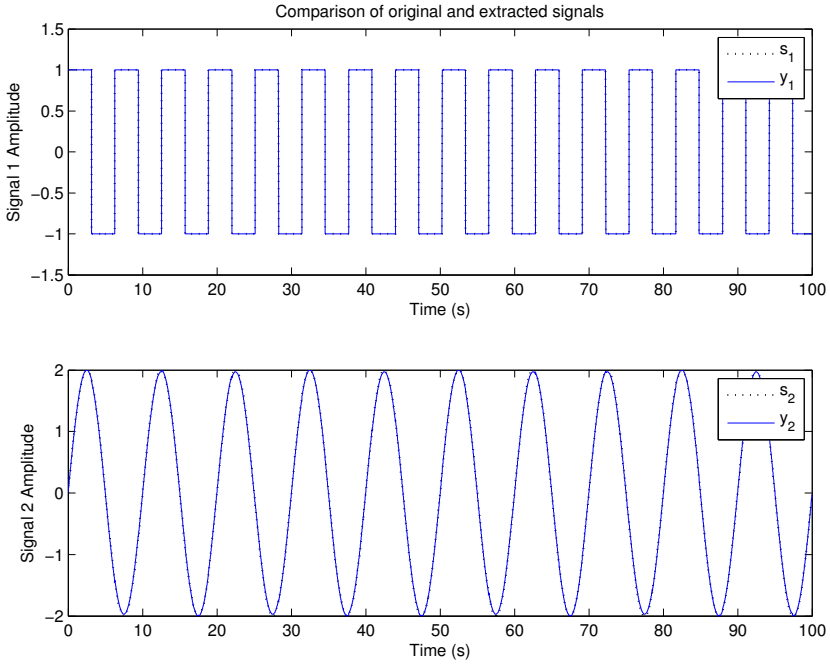

Fig. 2.9: Comparison of the original square and sine waves to the extracted waves.

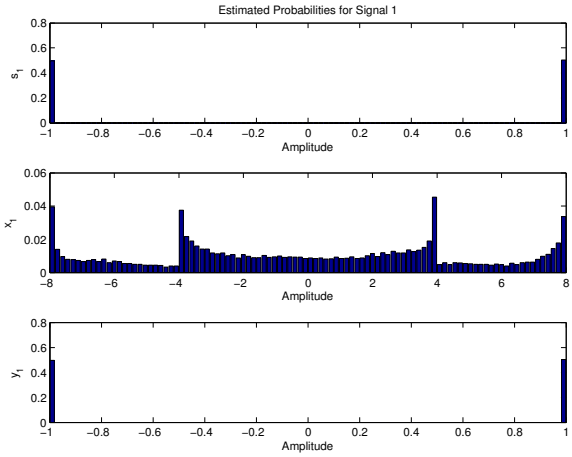

Fig. 2.10: Estimated probabilities of s_1, x_1, and y_1 for the square wave.

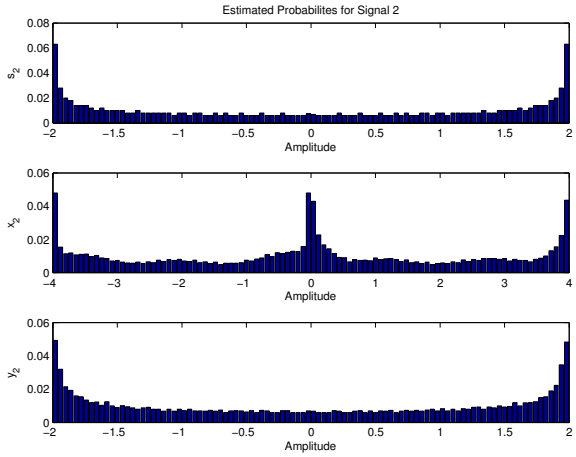

Fig. 2.11: Estimated probabilities of s_2, x_2, and y_2 for the sine wave.

2.4.1.3 Weierstrass and Sine Waves

As shown in the previous sections, ICA worked well for the smooth sine curves but was unable to fully recover the sharp corners and straight lines of the square wave. This observation led to the question of whether or not ICA is effective on other signals with singularities from rapid changes in amplitude. In this case, a *singularity* is defined as a point at which the function is not differentiable. This is an important question as many signals encountered in the physical world, including ECGs have sharp peaks and valleys that are essential to the accurate representation of the signal.

To investigate the effectiveness of ICA on waves with singularities, the *Weierstrass function* was used. The Weierstrass function is defined in such a way so that it is continuous and non-differentiable everywhere. There are multiple implementations of this function, including the original definition introduced by Karl Weierstrass in 1872 [Weie95] and shown in Equation 2.34.

$$W(t) = \sum_{n=0}^{\infty} a^n \cos(b^n \pi t) \quad \text{for:} \quad \begin{array}{l} 0 < a < 1 \\ b > 0 \\ ab > 1 + \frac{3}{2}\pi \end{array} \tag{2.34}$$

For this experiment, one Weierstrass signal was generated and mixed with one sine wave. As shown in Section 2.4.1.1, a mix of sine waves can be separated by ICA such that the separated mixtures are nearly identical to the original source signals. Therefore, the variable in this experiment is the Weierstrass signal.

The parameters for the Weierstrass signal were $a = 0.1$ and $b = 61$. This is one of the possible implementations such that the parameters satisfy the constraints of the Weierstrass function. The two signals were created as follows:

$$s_1(t) = \sum_{n=0}^{1000} 0.1^n \cos(61^n \pi t) \tag{2.35}$$

$$s_2(t) = \sin(2\pi t) \tag{2.36}$$

where $t = [-2 : 0.00001 : 2]$. Graphical representations of equations 2.35 and 2.36 are shown in Figure 2.12. The signals were then mixed as described in Equation 2.7 with mixing matrix $\mathbf{A}$, defined in Equation 2.32.

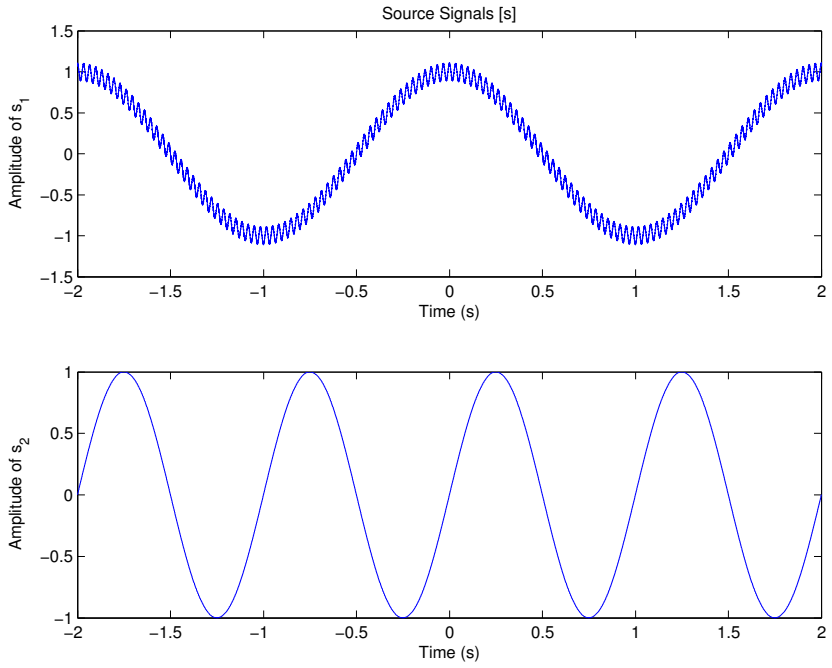

Fig. 2.12: The sine wave and Weierstrass function before mixing.

The resulting signals from the mixing process are shown in Figure 2.13. Neither of the mixed signals, x_1 or x_2, is recognizable as a Weierstrass or sine wave. However, when these mixed signals are supplied to the FastICA algorithm, the outputs y_1 and y_2 are nearly identical to the original signals, as shown in Figure 2.14.

The RMSE value for the comparison of the extracted signal, y_1, to the original signal, s_1, is $RMSE_1 = 1.1911 \times 10^{-12}$. Similarly, the RMSE value for the comparison of y_2 and s_2 is $RMSE_2 = 1.4062 \times 10^{-12}$. There was no change in entropy for the Weierstrass wave, and the change of entropy for the sine wave was 9.9008×10^{-7}.

In this experiment, despite the number of singularities in the Weierstrass wave and its lack of differentiability, ICA was still able to reliably separate the source signals.

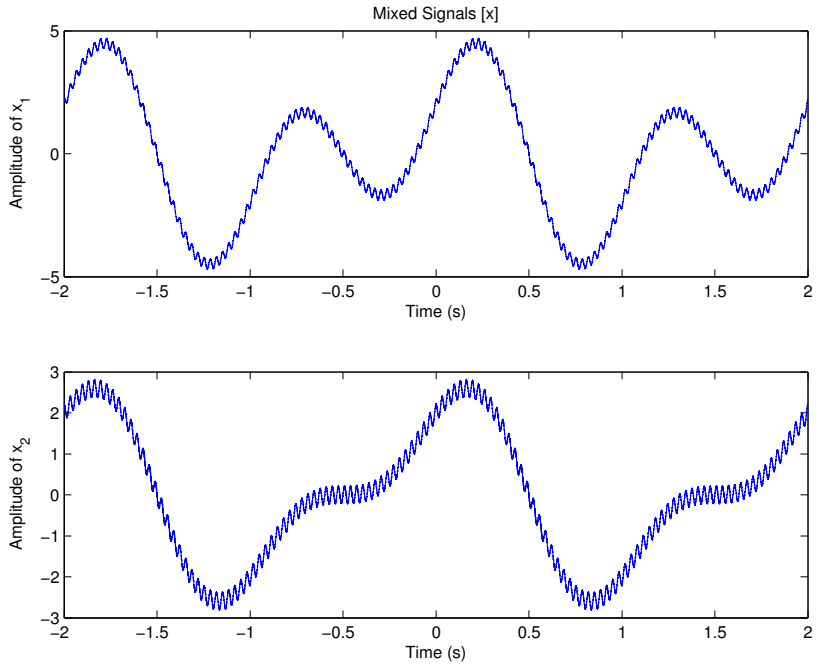

Fig. 2.13: Result of mixing the sine wave and Weierstrass function.

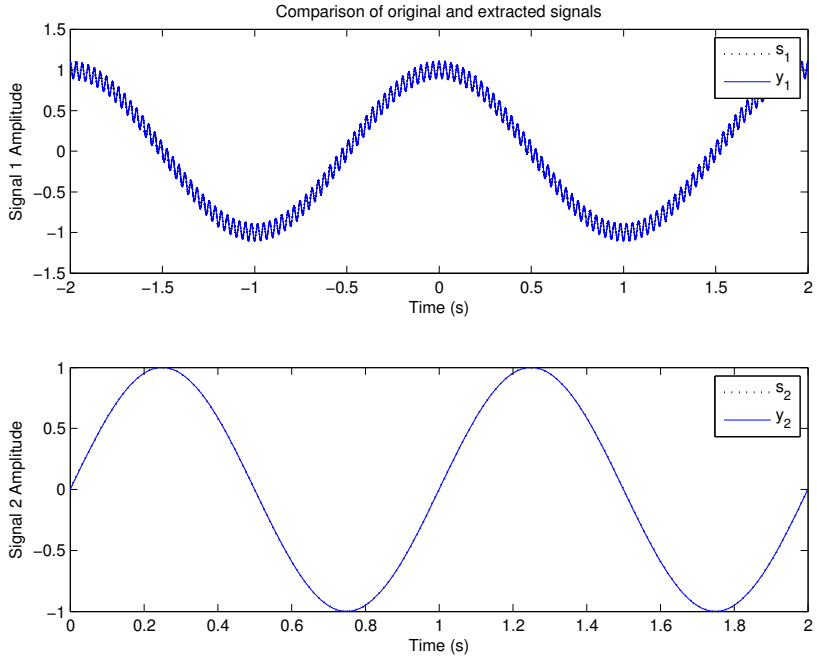

Fig. 2.14: ICA output shows the successful separation of the sine wave and Weierstrass function, although the Weierstrass function has been inverted and both outputs are scaled.

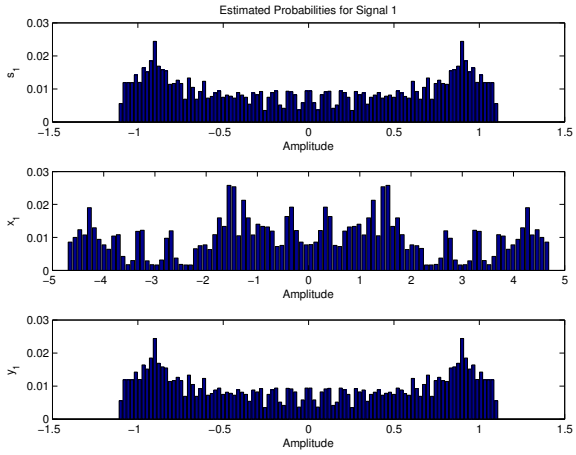

Fig. 2.15: Estimated probabilities of s_1, x_1, and y_1 for the Weierstrass wave.

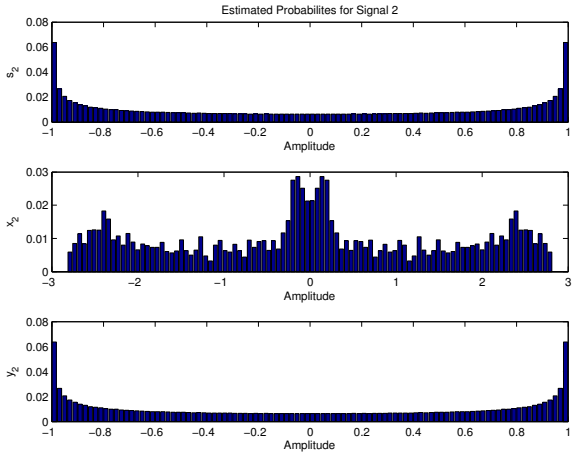

Fig. 2.16: Estimated probabilities of s_2, x_2, and y_2 for the sine wave.

2.4.1.4 Two Weierstrass Waves

In this next experiment, two Weierstrass signals with different a and b parameters were mixed. The two source signals were created as defined in Equations 2.37a and 2.37b, and graphed in Figure 2.17.

$$s_1(t) = \sum_{n=0}^{100} 0.1^n \cos(61^n \pi t) \tag{2.37a}$$

$$s_2(t) = \sum_{n=0}^{100} 0.2^n \cos(151^n \pi t) \tag{2.37b}$$

The results of mixing the signals are shown in Figure 2.18. They were mixed using the matrix **A** defined in Equation 2.32 using the standard ICA definition. Both of the mixed signals resemble the original source signals.

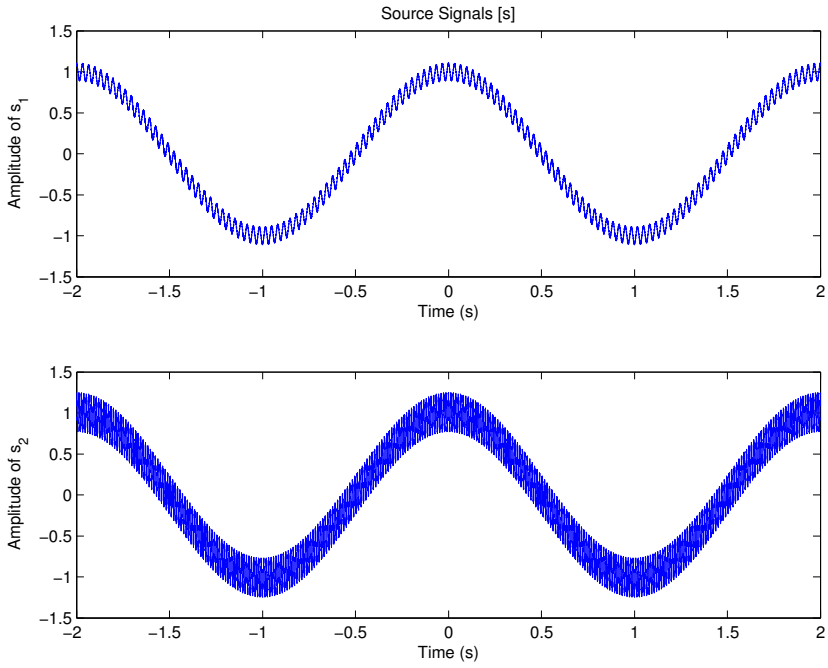

Fig. 2.17: The original two Weierstrass source signals.

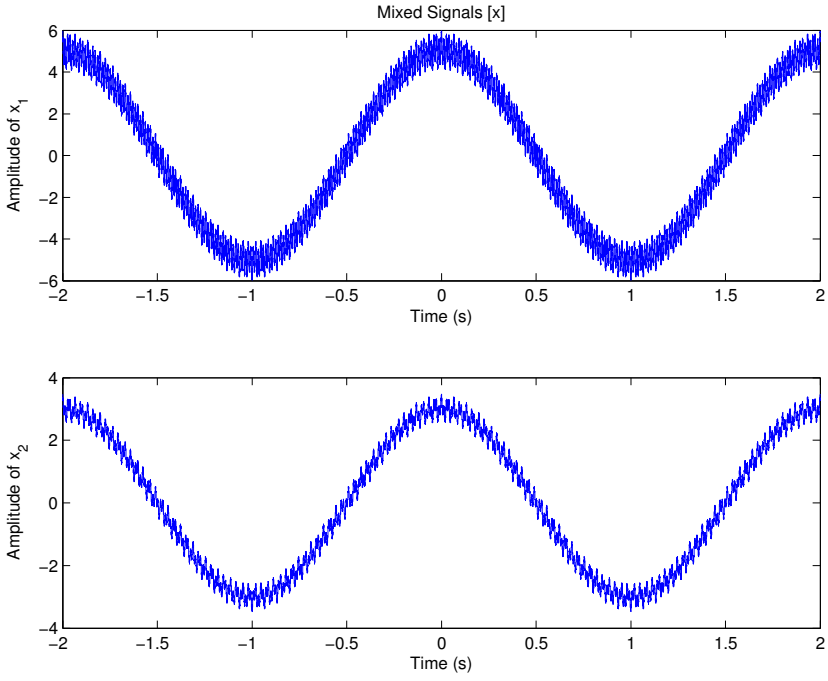

Fig. 2.18: The two Weierstrass signals after mixing.

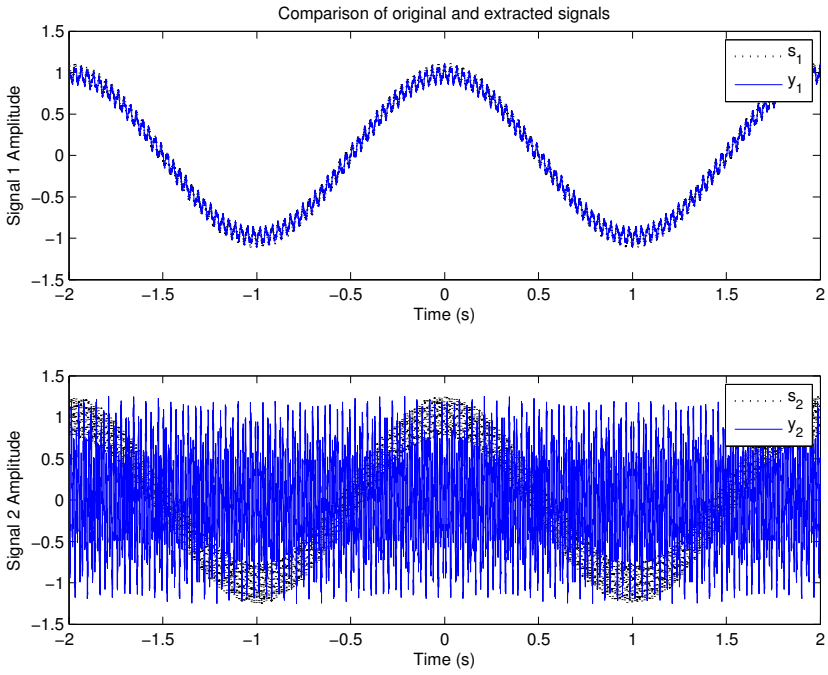

Fig. 2.19: The extracted Weierstrass signals.

The signals extracted by ICA are shown in Figure 2.19. After ICA is performed, only one of the extracted signals, y_1, resembles the original source. The RMSE values are $rmse_1 = 0.0354$ and $rmse_2 = 0.8348$. The change in entropy was 0.06682 for the first signal, and 0.081511 for the second signal.

This is particularly interesting because the amplitude of s_1 was less than that of s_2, yet it was recovered more accurately. The performance of ICA in this example may have been affected by the complexity of the Weierstrass signals.

2.4.1.5 Weierstrass and Square Waves

One final experiment was performed with the Weierstrass function. In this experiment, one Weierstrass signal (Equation 2.37a) was mixed with a square wave. The source signals are shown in Figure 2.20 and the mixed signals are shown in Figure 2.21.

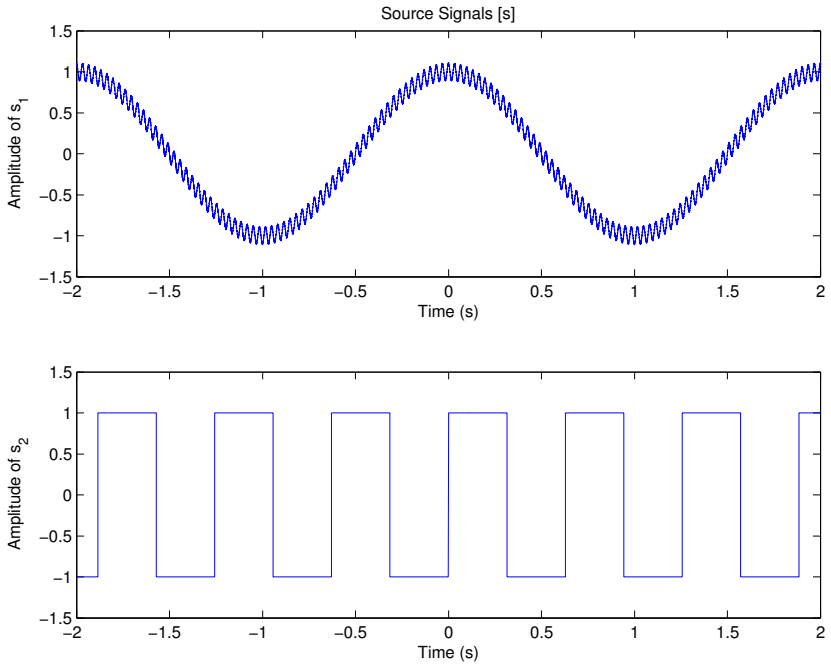

Fig. 2.20: The Weierstrass signal and square wave before mixing.

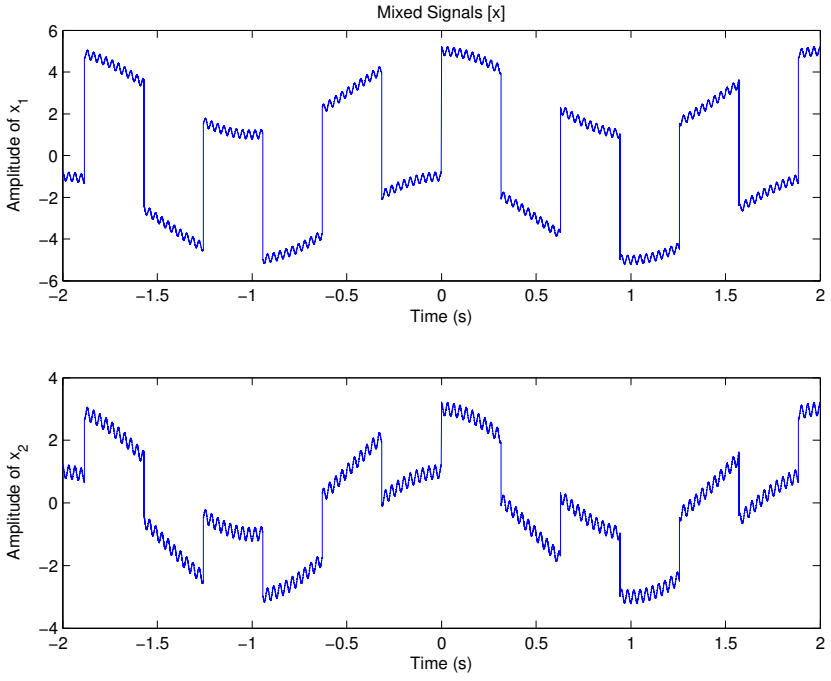

Fig. 2.21: The Weierstrass and square waves mixed together.

As shown in Figure 2.22, ICA was better able to separate the Weierstrass signal when combined with a square wave, compared to when it was mixed with another Weierstrass signal. The RMSE values were $rmse_1 = 3.2972 \times 10^{-5}$ and $rmse_2 = 7.6330 \times 10^{-15}$.

The change in entropy for the Weierstrass wave was 5.3477×10^{-5}, and there was no change in entropy for the square wave.

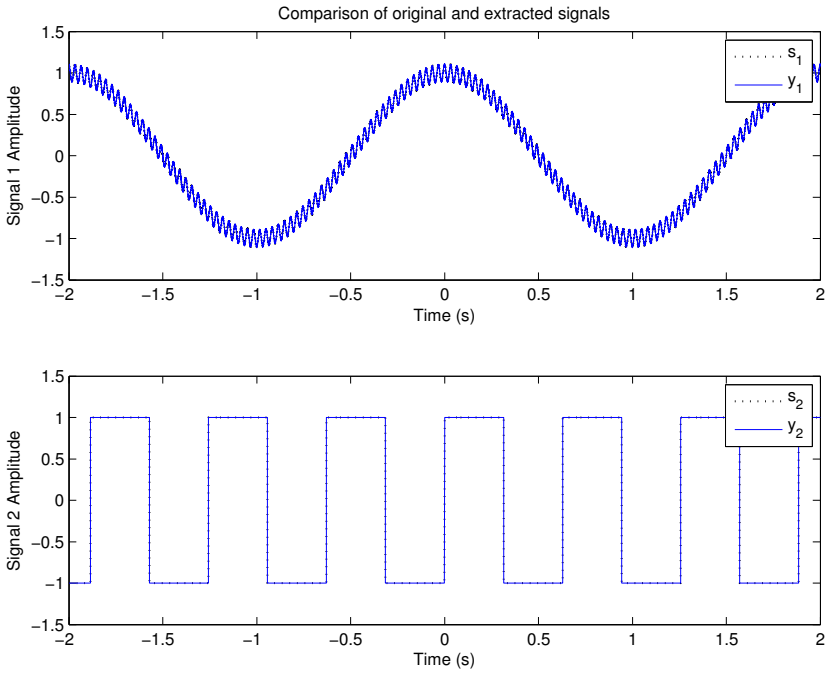

Fig. 2.22: The Weierstrass and square waves after ICA separation.

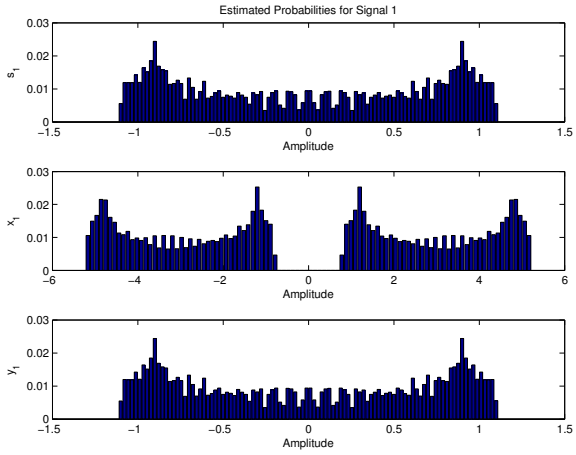

Fig. 2.23: Estimated probabilities of s_1, x_1, and y_1 for the Weierstrass wave.

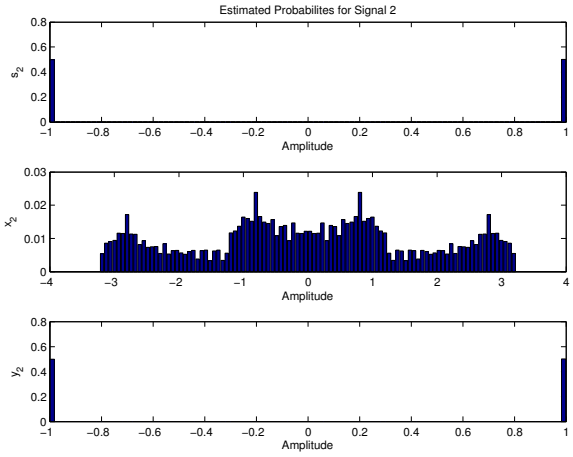

Fig. 2.24: Estimated probabilities of s_2, x_2, and y_2 for the square wave.

2.4.2 Effect of Amplitude Differences

The following experiment was performed to test ICA on signals with large differences in amplitude. The ability to extract small signals is an important characteristic of ISP techniques, and separates them from traditional SNR approaches. The two source signals used in the experiment are sine waves given in Equations 2.38a and 2.38b. In this experiment, the amplitufde of s_1 is 200 times greater than that of s_2, as shown in Figure 2.25.

$$s_1(t) = 200\sin(2\pi \times 01.t) \tag{2.38a}$$

$$s_2(t) = \sin(2\pi \times 0.03t) \tag{2.38b}$$

The source signals were mixed together using the ICA mixing model using mixing matrix **A** defined in Equation 2.32. The resulting signals x_1 and x_2 are shown in Figure 2.26.

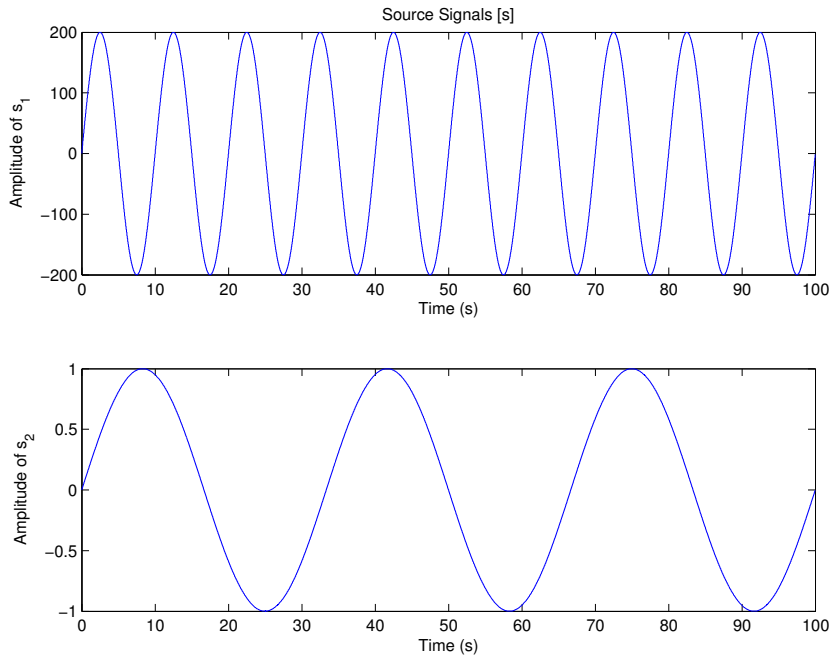

Fig. 2.25: Two source sine signals with amplitudes that differ by a factor of 200.

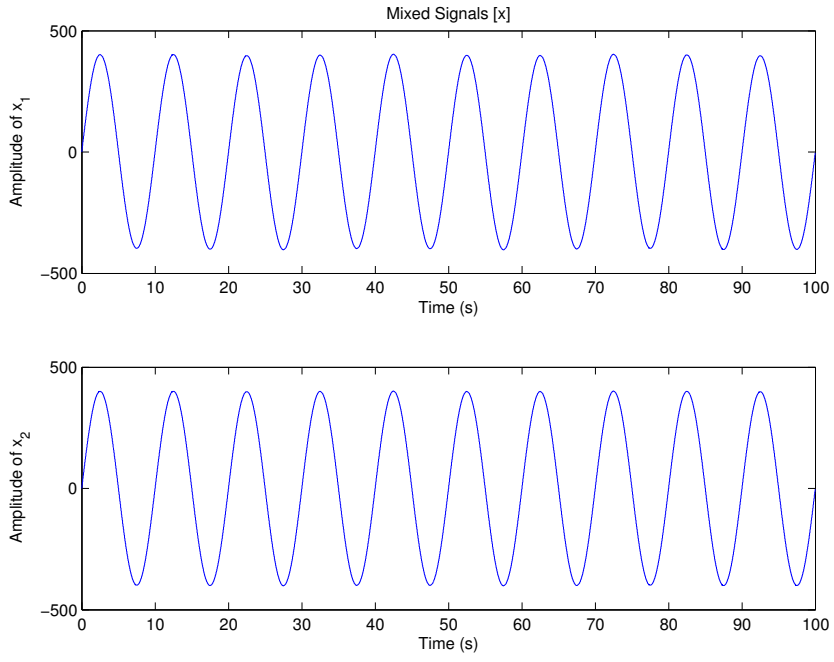

Fig. 2.26: The two signal mixtures created from sine waves of vastly different amplitudes.

After ICA, the separated signals match the original source signals as shown in Figure 2.27. The RMSE for the signals are very small, as $rmse_1 = 4.0016 \times 10^{-11}$ and $rmse_2 = 2.8929 \times 10^{-13}$. The normalized histograms in Figures 2.28 and 2.29 also show that the extracted signals match the source signals. This experiment shows that amplitude alone is not a significant factor in ICA, and that a small signal embedded in relatively much larger signal can be successfully recovered.

The change in entropy was -4.2864×10^{-6} for the first signal, and there was no change in entropy for the second signal.

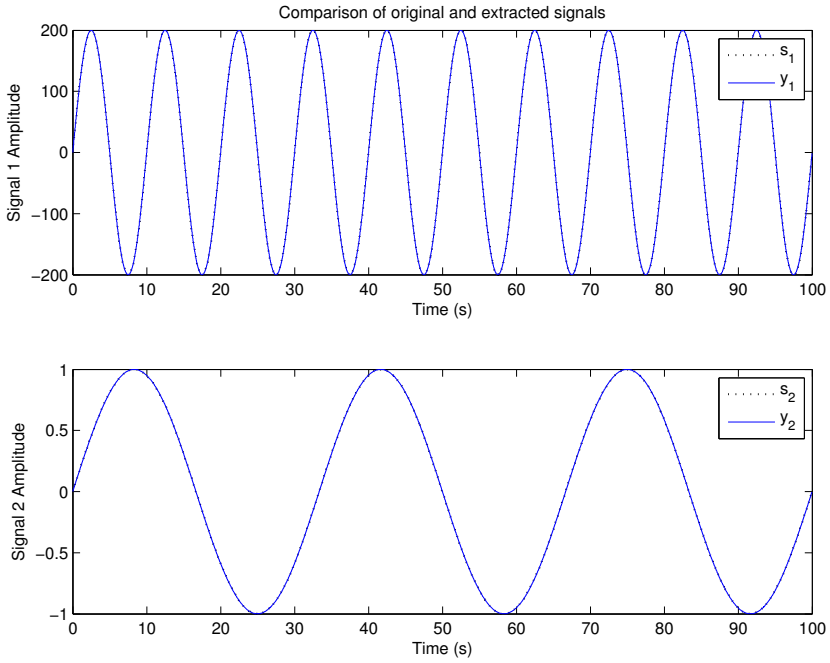

Fig. 2.27: Comparison of the extracted signals to the original signals.

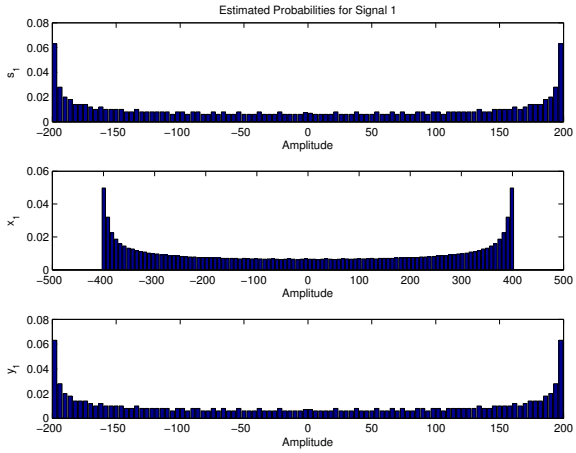

Fig. 2.28: Estimated probabilities histograms of s_1, x_1, and y_1 for the first signal.

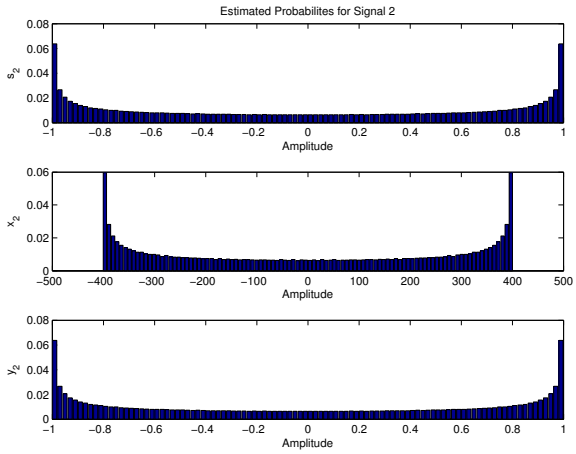

Fig. 2.29: Estimated probabilities of s_2, x_2, and y_2 for the second signal.

2.4.3 Effect of Phase Differences

This section presents two experiments that examine the effect of phase differences on the performance of ICA. In the maternal-fetal ECG separation problem, the relative position of the maternal and fetal heartbeats is unknown so any separation technique must be immune to phase differences.

2.4.3.1 Different Phase with Same Frequency

In this experiment the source signals have different phases but the same frequency, as define in Equations 2.39a and 2.39b. The source signals are shown graphically in Figure 2.30.

$$s_1(t) = 2\sin(2\pi \times 01.t) \tag{2.39a}$$

$$s_2(t) = \sin(2\pi \times 0.1t + 30) \tag{2.39b}$$

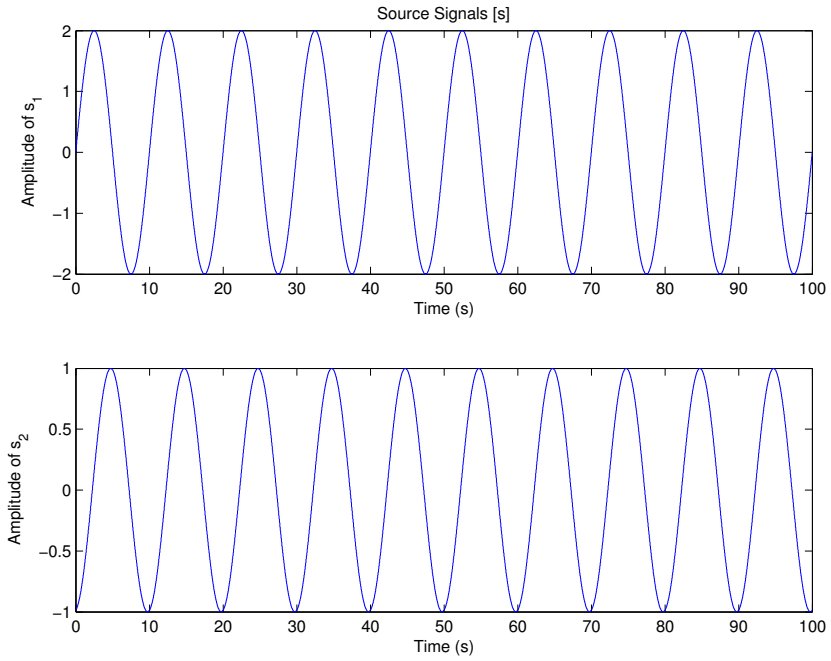

Fig. 2.30: Two source sine signals with different phases.

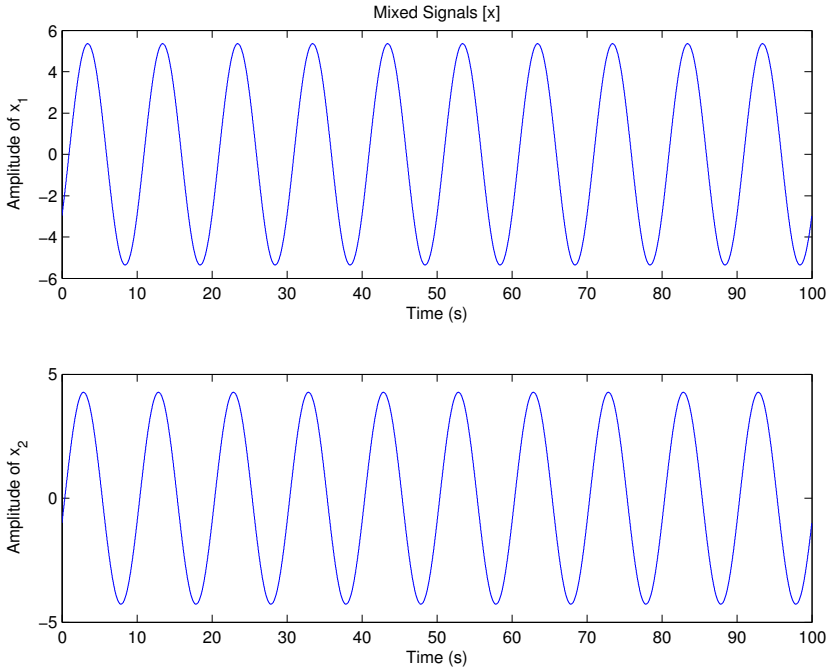

Fig. 2.31: The two signal mixtures created from sine waves with the same frequency but different phases.

When the signals are mixed together with mixing matrix **A** (Equation 2.32), the resulting signals x_1 and x_2 are both sine waves, as shown in Figure 2.31.

The results of extracting the signals using ICA are shown in Figure 2.32. It can be seen that both the shape and the frequency of the source signals are preserved, and this is corroborated in the normalized histograms in Figures 2.33 and 2.34. However, the phases of the original sources are not recovered correctly for either signal. The RMSE values for y_1 and y_2 are $rmse_1 = 0.2340$ and $rmse_2 = 0.2257$, respectively.

The change in entropy for the first signal was -0.0017293, and the change in entropy was 1.0658×10^{-14} for the second signal.

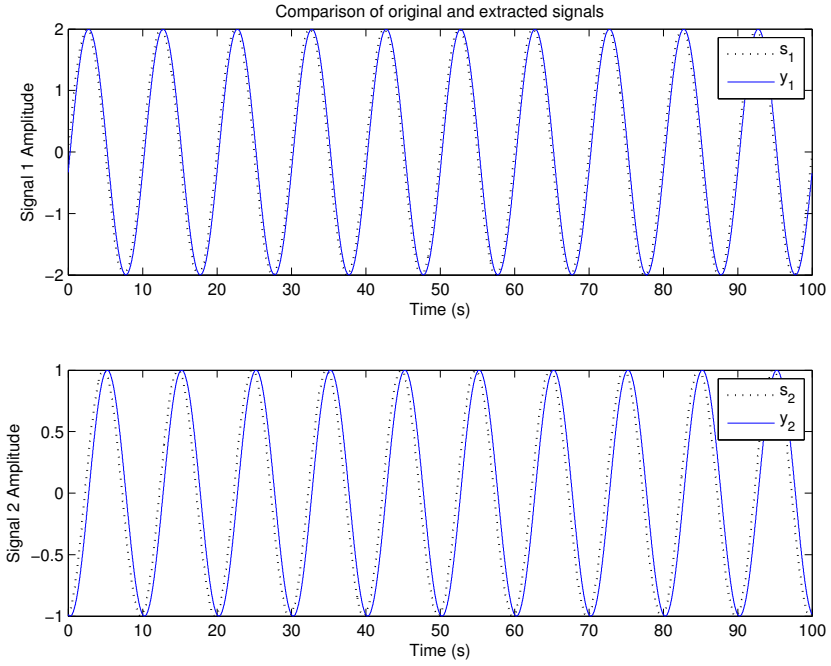

Fig. 2.32: Comparison of the extracted signals to the original signals.

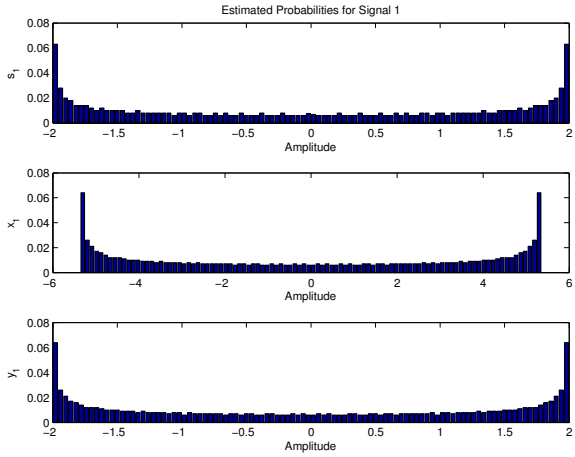

Fig. 2.33: Estimated probabilities of s_1, x_1, and y_1 for the first sine wave.

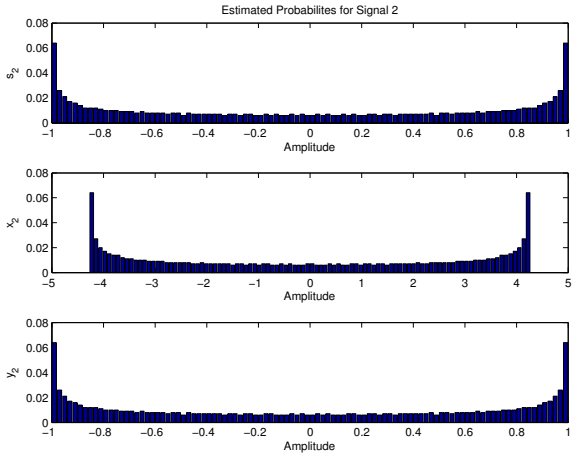

Fig. 2.34: Estimated probabilities of s_2, x_2, and y_2 for the second sine wave.

2.4.3.2 Different Phase with Different Frequency

In this experiment the source signals differ in both phase and frequency, as defined in Equations 2.40a and 2.40b. The source signals are also graphed in Figure 2.35.

$$s_1(t) = 2\sin(2\pi \times 01.t) \tag{2.40a}$$

$$s_2(t) = \sin(2\pi \times 0.03t + 30) \tag{2.40b}$$

The results of mixing the source signals with mixing matrix $\mathbf{A}$, defined in Equation 2.32, are shown in Figure 2.36. Since the frequencies of the source signals are different the mixtures are not sine waves, unlike in the previous example of source signals of the same frequency.

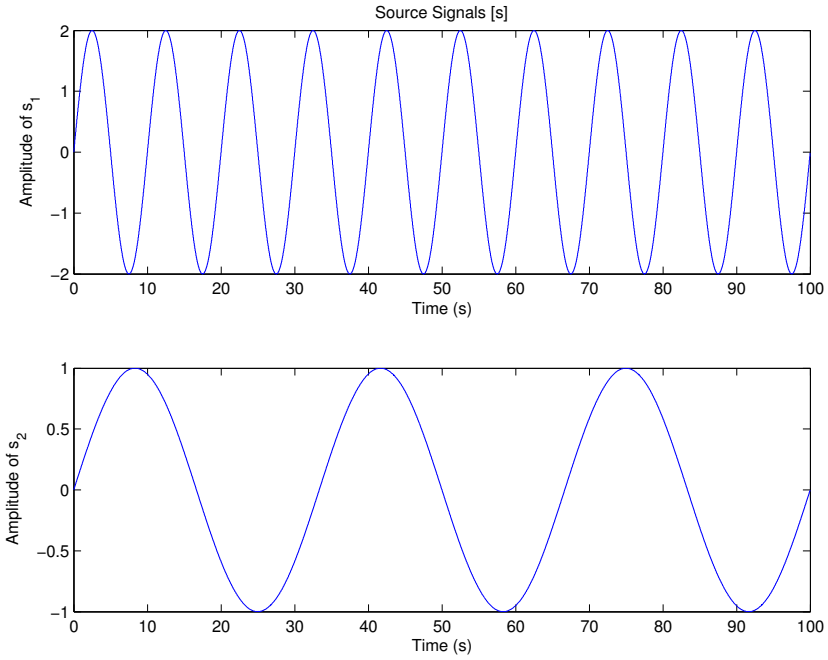

Fig. 2.35: Two source sine signals with different phases and frequencies.

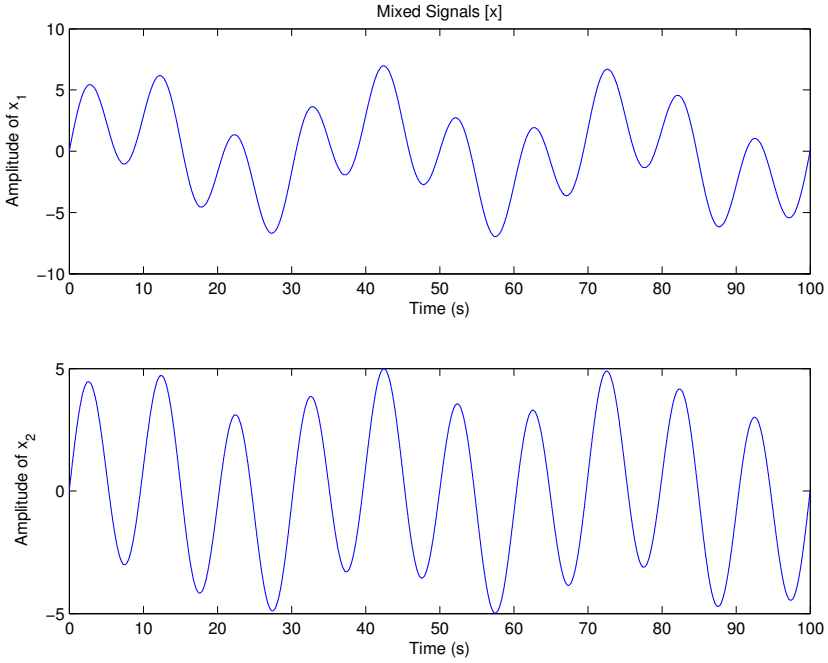

Fig. 2.36: The two signal mixtures created from sine waves with different phases and frequencies.

In this case, ICA was able to successfully separate the signal mixtures, as shown in Figure 2.37. This is further confirmed by the histograms in Figures 2.38 and 2.39. The RMSE values for the extracted signals are $rmse_1 = 2.4760 \times 10^{-15}$ and $rmse_2 = 2.0133 \times 10^{-15}$, which indicate that the extracted signals match the source signals much more closely than in the previous example.

The change in entropy was 1.1432×10^{-5} for the first signal, and there was no change in entropy for the second signal.

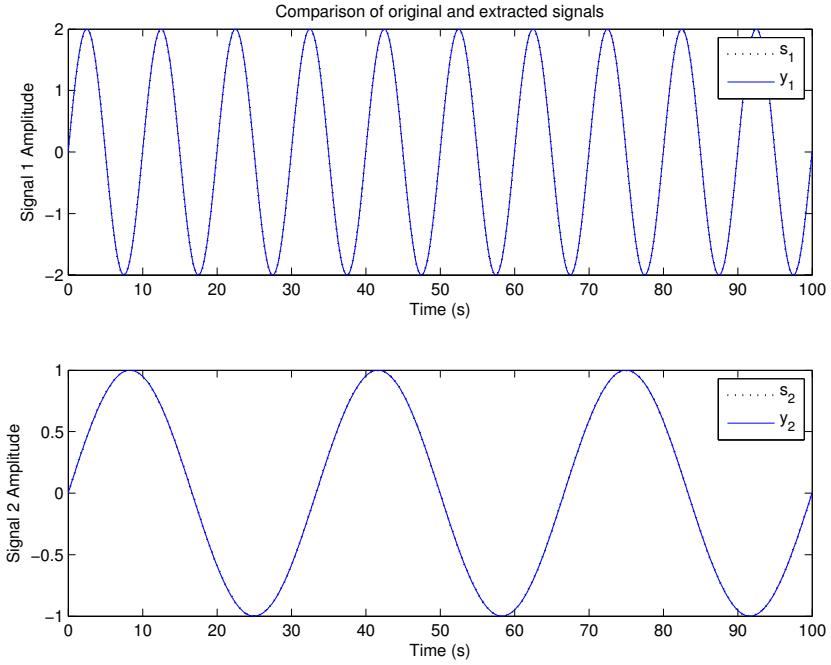

Fig. 2.37: Comparison of the extracted signals to the source signals when the phase and frequency of the source signals differ.

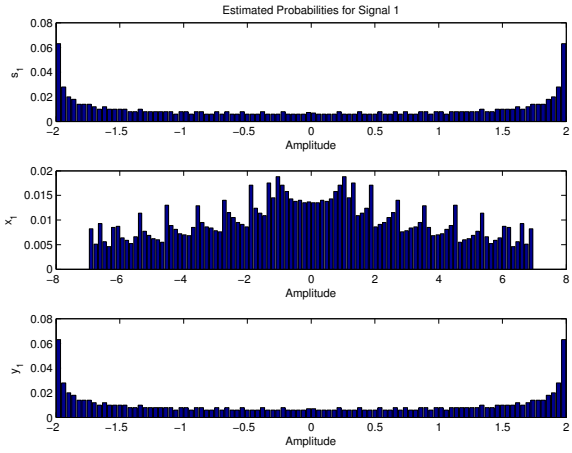

Fig. 2.38: Estimated probabilities of s_1, x_1, and y_1 for the first signal.

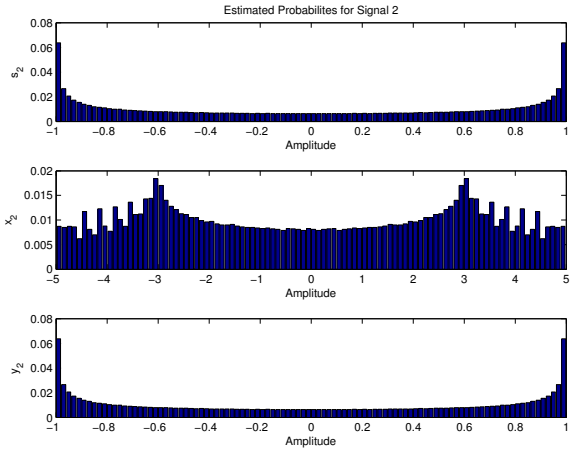

Fig. 2.39: Estimated probabilities of s_2, x_2, and y_2 for the second signal.

2.4.4 Stochastic Signals

The previous section demonstrated ICA using only deterministic signals such as the sine wave, square wave, and Weierstrass wave. In this section ICA is applied to mixtures of random signals, including an example with one Gaussian wave.

2.4.4.1 Sine Wave and Uniformly Distributed Random Wave

In this experiment a sine wave was mixed with a uniformly distributed random wave. The random wave was created using the random number generator in Matlab and has a fairly uniform distribution as shown in the top histogram of Figure 2.43. The sine wave was defined as $s_2 = \sin(2\pi \times 0.03t)$, and both signals are shown graphically in Figure 2.40.

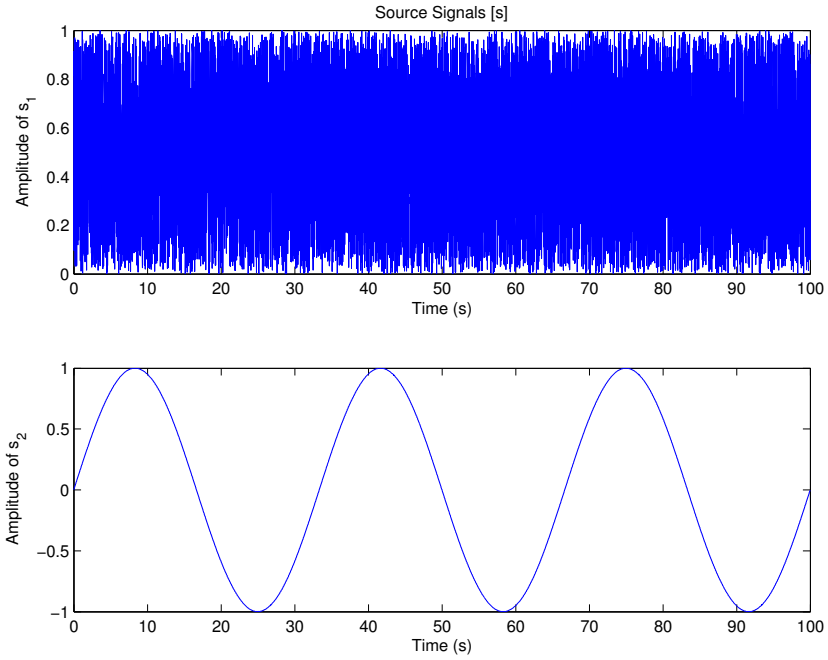

Fig. 2.40: The two source signals were a uniform random wave and a sine wave.

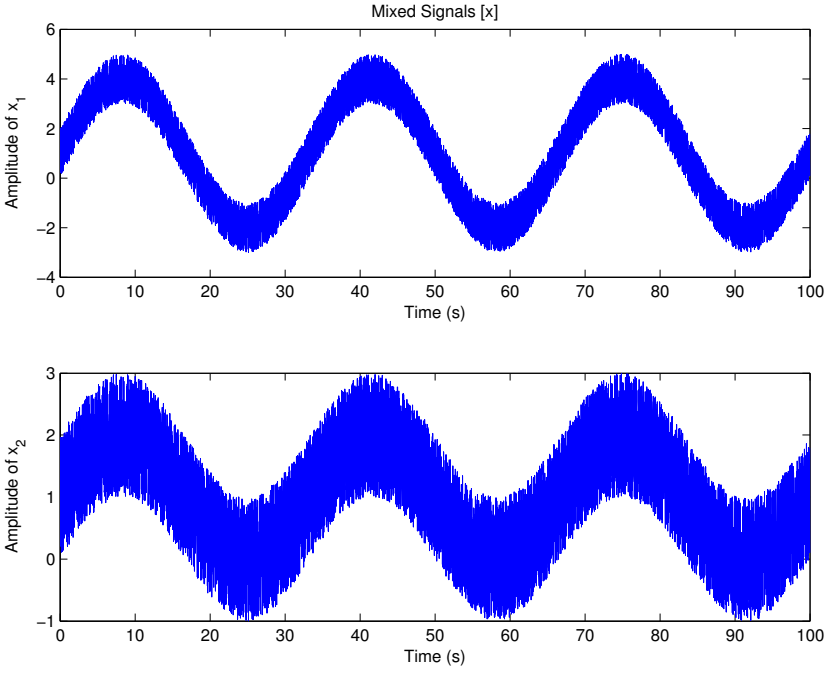

Fig. 2.41: The two signal mixtures created from one sine wave and one uniformly distributed random variable.

The source signals were mixed using the ICA model with mixing matrix **A** (Equation 2.32. As shown in Figure 2.41, the resulting mixtures look like noisy sine waves.

A comparison of the signals extracted by ICA with the source signals is shown in Figure 2.42. Both the clean, simple sine wave and random wave were successfully recovered by ICA, which can be seen more clearly when by zooming in on the random wave. However, the full signal is shown in Figure 2.42 in order to preserve the scale of both signals in the presentation of the example.

The performance of ICA is further shown in the normalized histograms of the two signals, in Figures 2.43 and 2.44. The RMSE values of the extracted signals are $rmse_1 = 5.5442 \times 10^{-4}$ and $rmse_2 = 0.0013$. The results were similar when this experiment was repeated using two uniform variables as the source signals. The change in entropy was 1.8987×10^{-4} for the first signal and -2.1107×10^{-3} for the second signal.

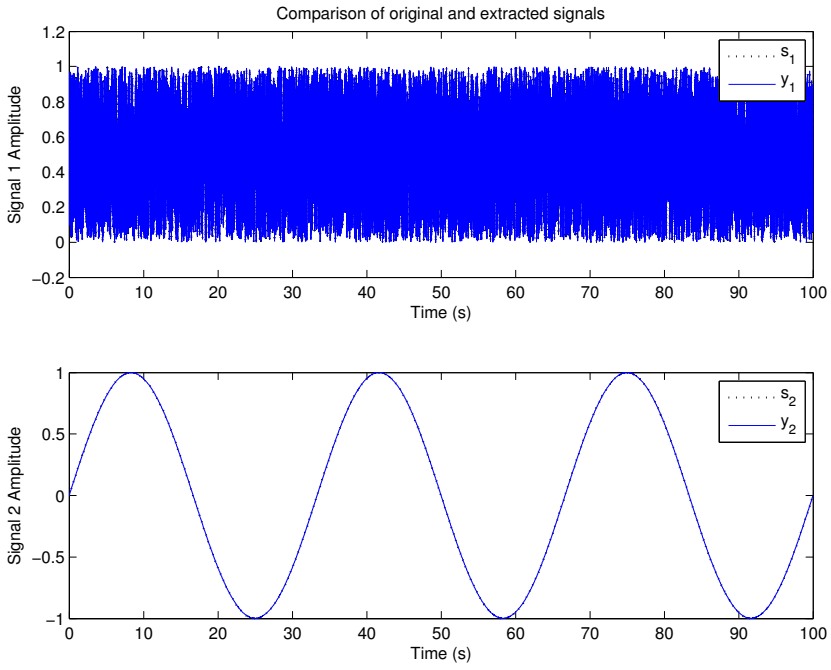

Fig. 2.42: Comparison of the extracted signals to the original signals.

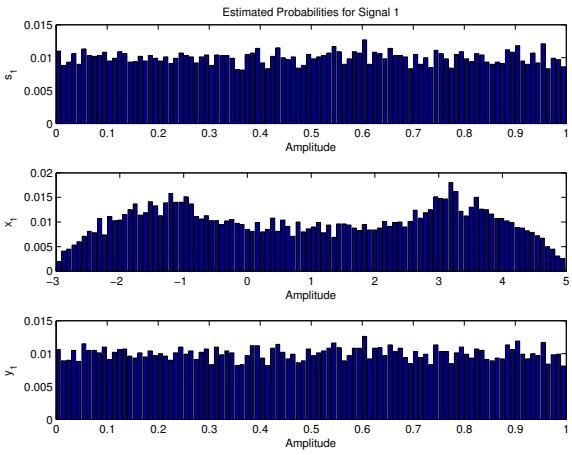

Fig. 2.43: Estimated probabilities of s_1, x_1, and y_1 for the uniform random wave.

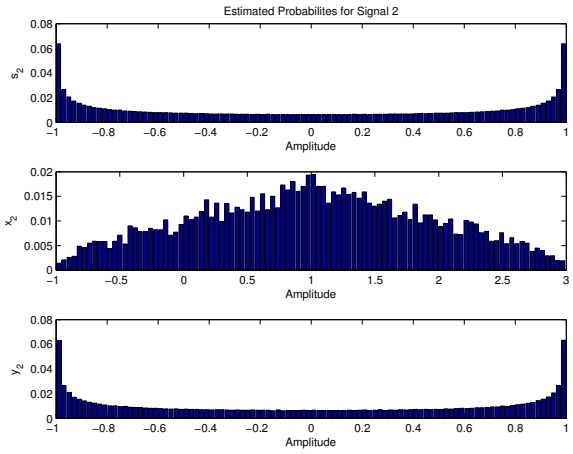

Fig. 2.44: Estimated probabilities of s_2, x_2, and y_2 for the sinew wave.

2.4.4.2 ICA with One Gaussian Wave

While ICA typically can not separate signals with Gaussian distributions, it is expected to work when a maximum of one of the source signals is normally-distributed. To demonstrate this feature of ICA, this experiment uses one Gaussian random variable and one sine wave as the source signals. The sine wave was defined to be $s_2 = \sin(2\pi \times 0.1t)$, and both source signals are shown in Figure 2.45.

The signal mixtures resemble noisy sine waves, as shown in Figure 2.46. They were created using mixing matrix $\mathbf{A}$ defined in Equation 2.32 using the standard ICA mixing model.

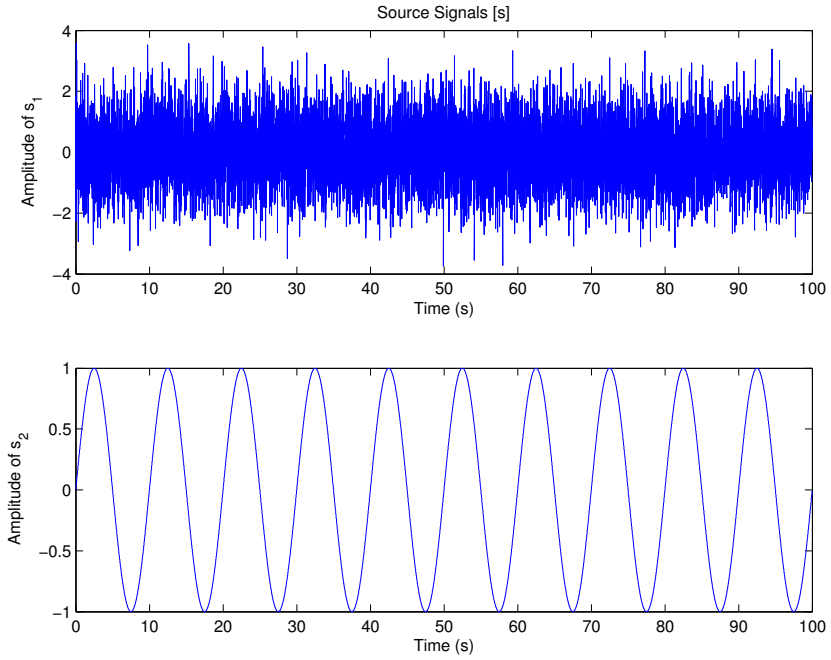

Fig. 2.45: The two source signals: one deterministic sine wave, and one Gaussian signal.

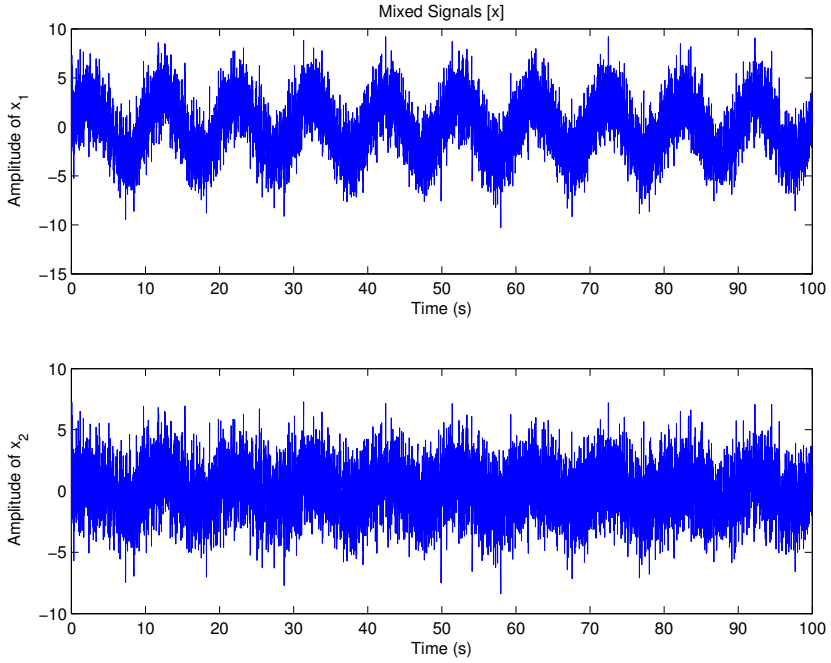

Fig. 2.46: The two signal mixtures created from a Gaussian random variable and a sine wave.

As shown in both Figure 2.47 and the normalized histograms of Figures 2.48 and 2.49, the source signals were successfully separated using ICA. Therefore, ICA is a feasible signal separation technique when a maximum of one of the source signals has a Gaussian distribution. The RMSE values of the extracted signals are $rmse_1 = 0.0031$ and $rmse_2 = 0.0019$.

The change in entropy was -3.9947×10^{-4} for the Gaussian signal, and -1.2764×10^{-2} for the sine wave.

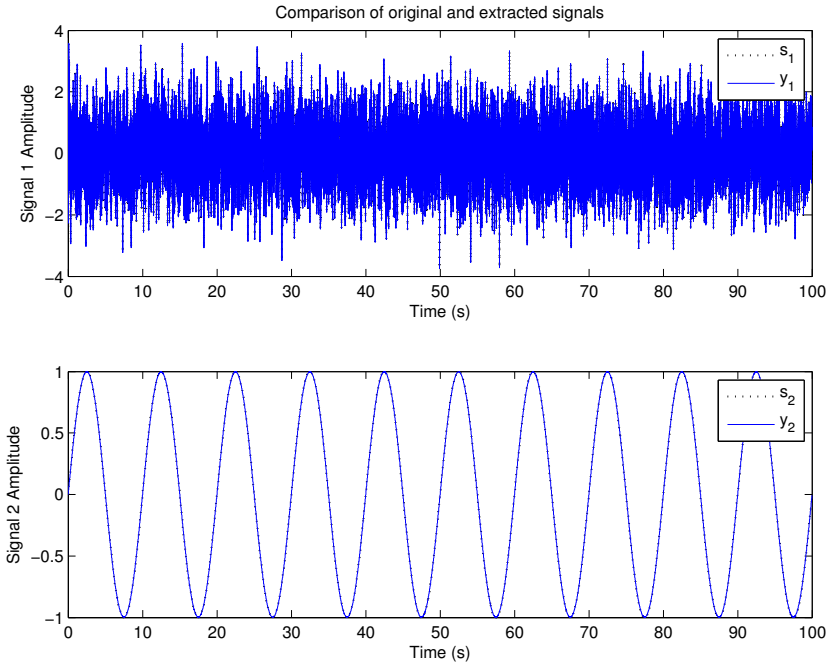

Fig. 2.47: Comparison of the extracted signals to the source signals, in the case when one is Gaussian.

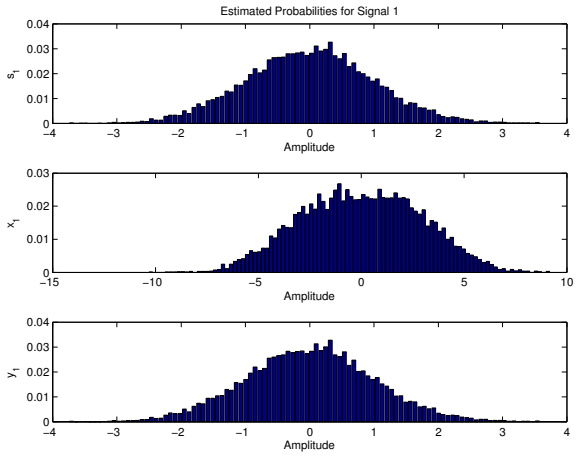

Fig. 2.48: Estimated probabilities of s_1, x_1, and y_1 for the Gaussian wave.

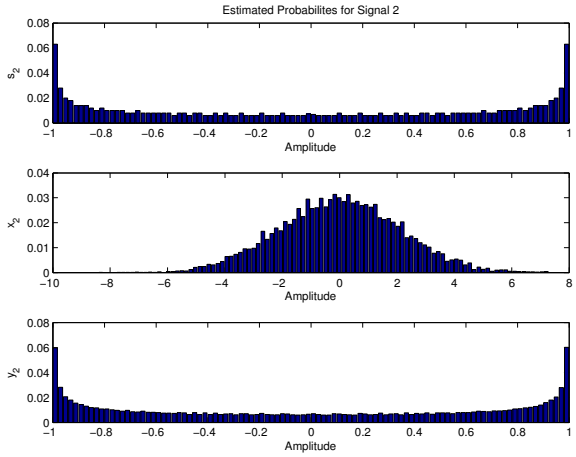

Fig. 2.49: Estimated probabilities of s_2, x_2, and y_2 for the sine wave.

2.5 Summary

This chapter introduced the technique of independent component analysis for blind source separation problems. The background and theory of ICA were discussed, including its fundamental principles and constraints. Algorithms to implement ICA were presented with a focus on FastICA, which was used to test ICA with a variety of source signal conditions.

Chapter 3

Electrocardiograms (ECGs)

Electrocardiograms are recordings of the electrical activity of the heart that may also be referred to as ECGs or EKGs. They are a non-invasive way to provide insight into how the heart is working and give diagnostic information on heart health. This chapter serves as an introduction to the use of ECG waveforms as a diagnostic tool, including an examination of the characteristic properties of these waveforms. An study of the benefits and limitations of current ECG simulation models is also presented, along with an introduction to the ECG recording repository databases on PhysioNet.

3.1 Introduction to ECGs

In order to pump blood throughout the body heart muscles are continuously activated in a repeating pattern. While each person is unique, their heart muscles fire in the same pattern to pump blood. The heart is composed of a special type of muscle called *cardiac muscle* that is striated, involuntary, and self-excitatory [Mous13]. Heart activity can be measured indirectly, and non-invasively, using ECGs.

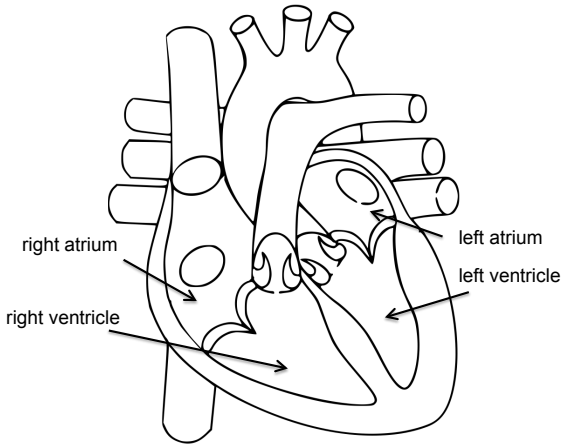

Fig. 3.1: The heart contains two chambers that receive blood (atria) and two chambers that pump blood (ventricles). [Wiki06].

As shown in Figure 3.1, there are four chambers in the heart that work together to direct de-oxygenated blood to the lungs and distribute freshly oxygen-rich blood to the rest of the body. The chambers that receive blood, one from the lungs and one from the body, are called atria while the chambers that pump blood are called ventricles.

The expansion and contraction of the heart to pump blood is governed by a series of natural pacemakers that ensure the heart beats steadily at nearly periodic intervals [Mous13]. For a healthy adult, the frequency of these beats is on average 60-100 *beats per minute* (bpm). The variations in the frequency of the heart's pumping motion will be discussed further in Section 3.1.2. This average frequency is maintained by the S-A node, which is the primary pacemaker of the heart. In the event the S-A node is compromised, the A-V node acts as a secondary pacemaker at 40-60 bpm, and the ventricles themselves are tertiary pacemakers that can pump the heart at 20-40 bpm.

3.1.1 The ECG waveform

The electrical signal recorded as an ECG is a superposition of the electrical activity of the various muscles of the heart as it expands and contracts to receive and pump blood. For adults, the ECG can be measured non-invasively with surface electrodes arranged over the chest and limbs [Mous13]. This arrangement produces potential differences between the electrodes, which are called *leads*. Typically a 12-lead arrangement is used in clinical practice.

An idealized version of an ECG waveform is shown in Figure 3.2. The distinct peaks and troughs in the wave correspond to specific heart muscle activity and have standard labels. The relative amplitudes of the peaks, as well as the time intervals between them, are used to evaluate heart health.

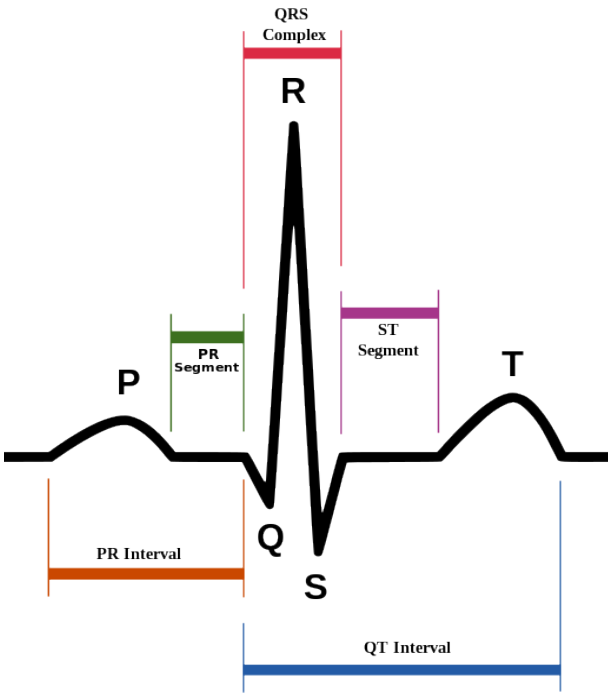

Fig. 3.2: Idealized ECG waveform with standard labels. [Wiki06a].

The correspondence between the main sections of the ECG waveform and heart muscle activity is as follows:

- **P-wave:** This is considered the start of a new heart beat and represents the depolarization of the receiving chambers, or atria.
- **QRS complex:** Consisting of the Q, R, and S peaks and troughs, the QRS complex is the largest and most easily identifiable part of the ECG waveform, even in the presence of electrical noise. It represents the rapid depolarization of the receiving chambers, or ventricles, and typically lasts 100ms.
- **T-wave:** Typically the final stage of the ECG waveform, the T-wave occurs when the ventricles are repolarized.

A final peak, called the U-wave, may also be present after the T-wave. This peak is typically much smaller than the rest of the ECG waveform, although a pronounced U-wave may be indicative of health issues.

3.1.2 Heart Rate Variability

The collection of peaks and troughs that make up an ECG waveform, as shown in Figure 3.2, are repeated each time the heart goes through a contraction-expansion cycle to pump blood. Heart rate frequency, or the number of beats per minute, is measured using the prominent R-waves by calculating the *RR-interval* as shown in Figure 3.3.

While heartbeats are repetitive, the time between two successive R-waves is not constant. This leads to an important characteristic of the cardiac cycle called the *heart rate variability* (HRV), which is the variation in RR-intervals [Kawa97]. HRV is a natural and healthy characteristic of ECGs, although the amount of variability can be a useful indica-

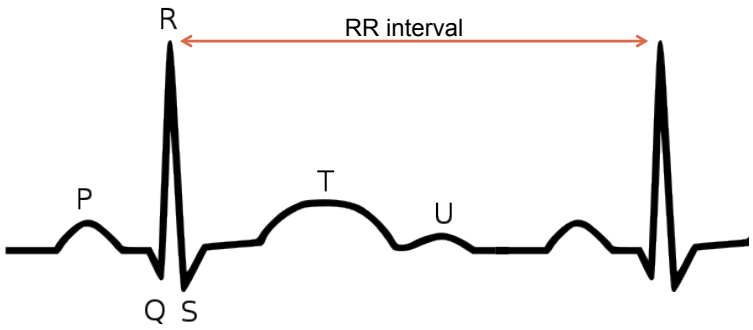

Fig. 3.3: Visualization of an RR-interval [Wiki07].

tor for the detection of some cardiovascular and non-cardiovascular diseases. In particular, lower HRV is associated with certain diseases [Klei87].

It is possible that the HRV could be affected by the sampling rate of the ECG signal. Since HRV is measured using the R-wave peaks, accurate R-wave peak readings are essential to correctly calculate the RR-interval. However, the entire QRS complex occurs over a timespan of about 100ms, so the sampled R-wave peak may not capture the true maximum. The true peak could be interpolated from the sampled ECG, however in practice it is assumed that the sampled peak is sufficient. Calculating the true peak is outside the scope of this thesis and the sampled peak will be used as it is the common practice in the literature.

3.2 Properties of ECGs

In this section the properties of ECGs that are relevant to signal analysis are presented. These properties can also serve as metrics when comparing ECG models to recorded signals. In particular, the three properties of ECGs that are discussed are the stationarity of both the ECG signal and HRV, the probability distribution of the time-varying signal, and the fractal dimension of the HRV.

Other measures used to analyze ECG signals in medical practice include:

- **RMSSD:** The *root mean square of the successive differences* (RMSSD) is the root mean square of consecutive RR-intervals as defined in Equation 3.1:

$$RMSSD = \sqrt{\frac{1}{n}\left(\sum_{i=2}^{n}(RR_i - RR_{i-1})^2\right)}, \tag{3.1}$$

  where n is the total number of RR-intervals, and RR_i is the i^{th} RR-interval [**?**]. The RMSSD is used in HRV analysis.

- **pNN50:** This is the number of consecutive RR-intervals that differ by more than 50ms divided by the total number of RR-intervals. The measure derives its name from the term *NN-intervals*, where the "NN" designation means "normal" RR-intervals [Miet02].

- **pNN20:** This measure is the same as pNN50, except with a threshold of 20ms. Both the pNN20 and pNN50 are used in heart health diagnosis.

3.2.1 ECG and HRV Stationarity

Like many biological signals, the ECG and HRV are *non-stationary* processes [Gao13] [Mous13]. A *stationary process* is one in which the "joint distribution of any subset of the sequence [of data points] is invariant with respect to shifts in the time index" [Cove06]. This means that the statistical properties of ECG and HRV signals, such as their mean and variance, are not constant over time and signal processing techniques that assume stationarity can not be applied.

Fortunately, the issue of non-stationarity with ECGs and the HRV can be overcome by segmenting the signals into shorter, stationary windows. This is the principle behind the short -time Fourier transform (STFT), for example. ECG signals are particularly well-suited

to this type of analysis, as these signals can often be considered stationary for minutes at a time [Zgal13].

3.2.2 Probability Distribution of ECGs

An important consideration for ICA is the probability distribution of the source signals, of which no more than one can be Gaussian for ICA to work properly. This constraint is satisfied when the source signals are ECG recordings. As shown in Figure 3.4, the normalized histogram of an ECG recording is distinctly non-Gaussian and demonstrates significant positive skew. The normalized histogram was produced using data from the MIT-BIH Normal Sinus Rhythm Database (nsrdb) on PhysioNet [Gold00]. The normalized ECG histogram more closely resembles the Pareto distribution, which is common for biological signals [Kins13].

While the HRV would not be a source signal for ICA separation, it is worth noting that its distribution is also non-Gaussian as shown in Figure 3.5. This normalized histogram was created from the same ECG recording and also demonstrates a positive skew.

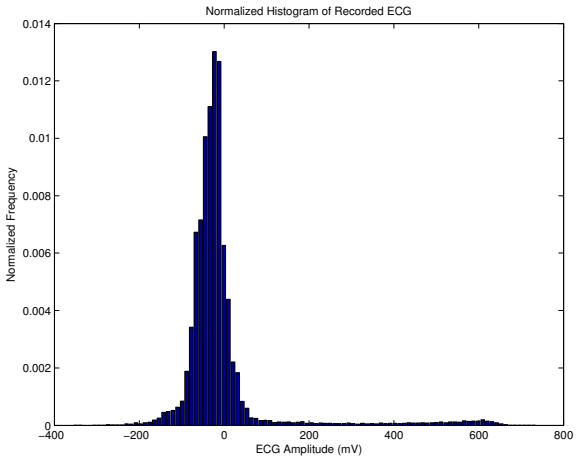

Fig. 3.4: Normalized histogram of a recorded ECG signal.

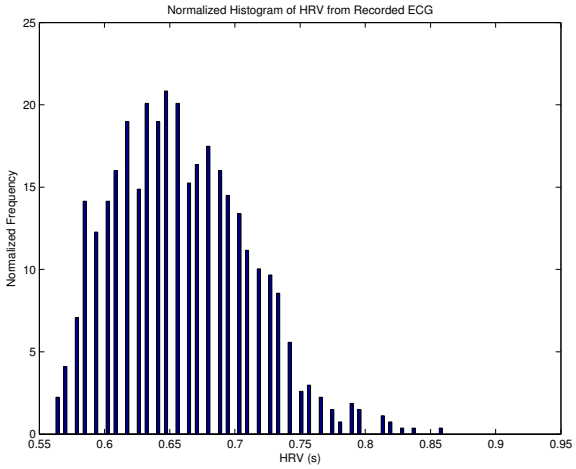

Fig. 3.5: Normalized histogram of the HRV from a recorded ECG signal.

3.2.3 Fractal Dimension as a Measure of HRV

HRV is a statistically *self-affine* process. That is, when the HRV is scaled along the x-axis (without changing the y-axis scaling) the resulting signal statistically resembles the original [Kins13]. This concept is shown in Figure 3.6, where the HRV from a recorded ECG is shown. The inset graph contains a segment of the HRV at 4x magnification, which statistically resembles the original signal.

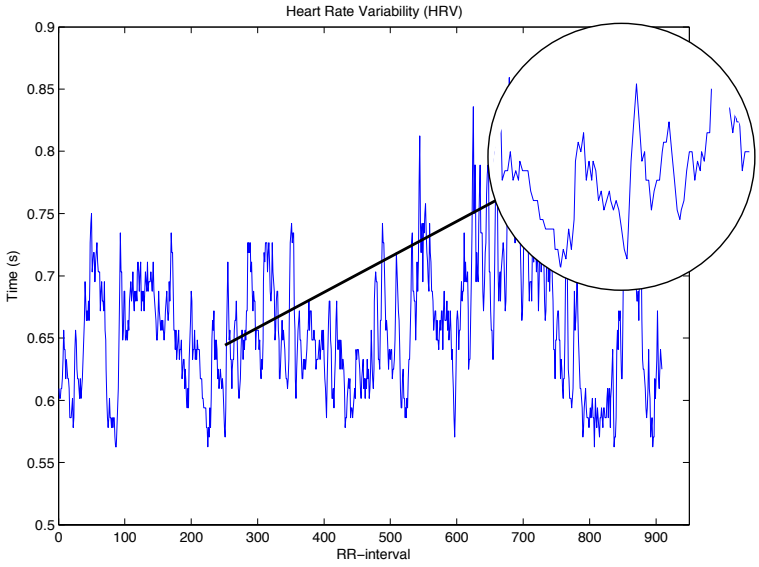

Fig. 3.6: HRV is a self-affine process.

Since the HRV is self-affine, its fractal dimension can be calculated. The fractal dimension can be considering a way of measuring the complexity of a signal. For a time series the appropriate fractal dimension is the *spectral dimension*, which is denoted D_β. As mentioned in Section 2.3.6.4, the spectral dimension is defined as

$$D_\beta = E + \frac{3-\beta}{2}, \tag{3.2}$$

where β is the *spectral component* obtained by finding the slope of the log-log power spec-

trum density of the signal in question. The term E is the number of independent variables in the signal, which for time series such as HRV is $E = 1$.

3.3 ECG Modelling

Due to the prevalence of ECGs in medicine, ECG modelling is a subject of interest in that field. ECG models can be used to understand the activity of the heart in a noise-free environment. One way to approach an electrical model of the heart is to build it from the ground up by simulating electrical currents at the cellular level and combining them into a composite signal [Sach98]. However, this method is computationally intensive so nonlinear dynamical models are more typically used to simulate ECGs [Pott08].

Three of the most common nonlinear dynamical ECG simulators are ECGwaveGen, ECGsyn, and ECGfm. Created by Floyd Harriott, ECGwaveGen is a simplistic ECG simulator with user-settable parameters [Harr11]. It does not take any variability or randomness into account, but does model each of the major components of ECG waves. A signal generated by ECGwaveGen is shown in Figure 3.7, with a sampling frequency of 1kHz, 60 bpm, and amplitude of 1V.

Since the ECGwaveGen model is completely deterministic, multiple trials with the same input parameters will always produce the same waveform. It is a good tool to graph the general shape of an ECG, but does not capture any of the subtleties of the waveform so this model will not be considered any further in this work. The other two models will be described in more detail in the following sections.

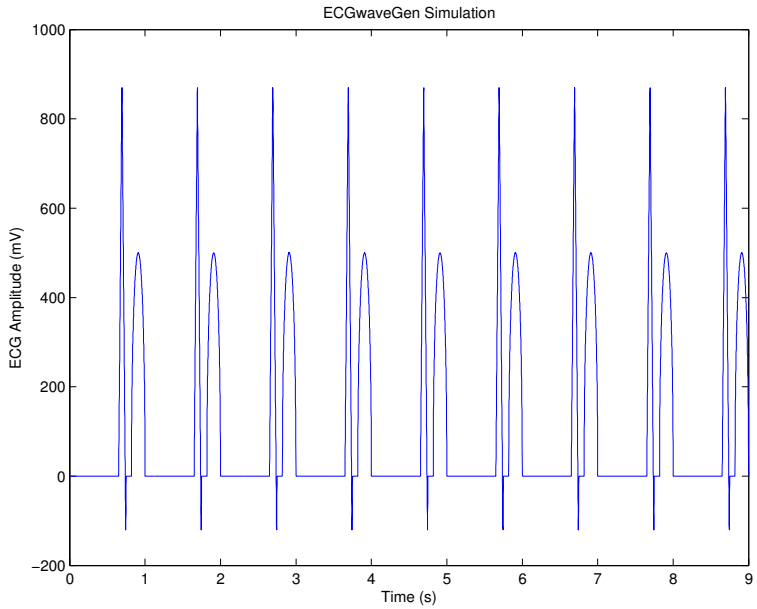

Fig. 3.7: Deterministic output from ECGwaveGen.

3.3.1 ECGsyn Model

One of the leading ECG waveform generators is the ECGsyn model developed by Patrick McSharry from the University of Oxford and Gari Clifford from MIT [McSh03][McSh12]. The ECGsyn model is freely available for Matlab, C, and Java on PhysioNet [Gold00]. By using a combination of parameters that may be adjusted by the user, along with random variables, the model successfully mimics many features of the cardiac cycle. The user is able to set the sampling frequency, signal length, and signal mean, as well as adjust each of the P, Q, R, S, and T waves individually.

In contrast to ECGwaveGen, the waveforms generated by ECGsyn are distinct and non-deterministic because some of the characteristics are governed by random processes in the model. The use of random numbers produces differences in the RR-intervals lengths, as well as the R-wave amplitudes. An example waveform generated by ECGsyn with a sampling frequency of 1kHz is shown in Figure 3.8. As can be seen, this model produces

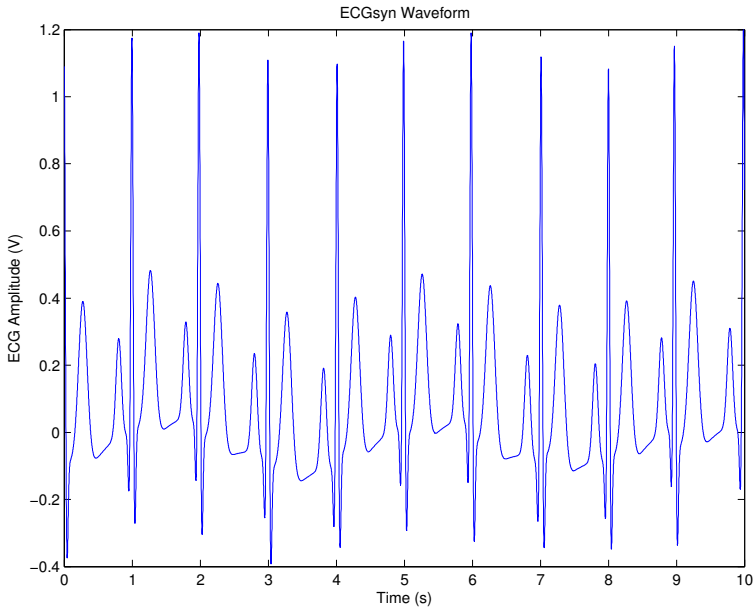

Fig. 3.8: Example output from ECGsyn.

visually realistic heartbeats.

The main weakness in the ECGsyn model is its inability to accurately simulate HRV. The RR-interval lengths in ECGsyn have a uniform distribution. As shown in Figure 3.5, the probability distribution of HRV in recorded ECG signals is distinctly non-uniform. Therefore, while ECGsyn is a valuable tool to model the cardiac cycle, it can not reliably be used in cases where the HRV is important.

3.3.2 ECGfm Model

In response to the HRV deficiencies of the ECGsyn model, a new ECG waveform generator was developed by Michael Potter [Pott08]. This model, ECGfm, builds on ECGsyn but takes RR-intervals from recorded ECG signals as input and uses them to build the ECG signal. This is a significant improvement over the uniformly distributed HRV in ECGsyn [Pott09]. In this way, the HRV of the ECGfm waveforms is guaranteed to have

the same properties as ECG recordings.

The ECGfm algorithm is written in Matlab and maintains the same functionality of user-settable parameters as ECGsyn. However, unlike ECGsyn which is widely accessible on PhysioNet, the source code for ECGfm is not freely available so it is not used in ECG research.

3.4 Recorded ECG Databases

In addition to modelled ECG waveforms, it is useful to perform experiments on recorded ECG data. One way to do this is through databases of ECG recordings that have been compiled for academic purposes. The largest collection of such databases is on PhysioNet, which stores recorded bioelectrical signals in its PhysioBank archives [Gold00]. All of the signals in the PhysioBank databases are free to use for academic purposes.

PhysioBank contains a number of databases dedicated to ECG records, including recordings from healthy subjects as well as databases dedicated to particular arrhythmias. There are also databases with maternal-fetal ECG recordings. However, these databases differ in both size and quality, with some containing under ten records and others containing over one hundred. The length of the recordings also differs between the databases, with some as short as one minute.

As well, there is no current standard for recording ECG signals, so the sampling frequency varies greatly between the databases. In some the signals are sampled at a rate as low as 128 samples/second, or up to 10,000 samples/second. The higher sampling frequency is used for shorter recordings that are typically one minute long. While even a lower sampling frequency will capture the general morphology of an ECG signal, including all of the waveforms depicted in Figure 3.2, it is possible that more subtle features of the

ECG are missed. It is also possible that the true R-wave peak may not be sampled, and a nearby point is captured instead, which would lead to changes in the HRV as discussed in Section 3.1.2. In practice, 128 samples/second is most typically seen, but it has been suggested that the sampling frequency for an ECG should be at least 200-500 samples/second [Rudi07]. Due to all of these differences in the recordings, not all of the databases are suitable for all purposes.

3.5 Electrocardiograms Summary

This chapter introduced the adult electrocardiogram waveform with a focus on its characteristics and properties including heart rate variability and estimated probability distribution. The three heart models of ECGwaveGen, ECGsyn, and ECGfm were discussed, outlining their strengths and limitations at simulating ECG waveforms. Finally, the databases on PhysioNet of ECG recordings were introduced as a valuable tool for ECG signal experiments.

Chapter 4

Maternal-Fetal ECG Separation Problem

In this chapter, the problem of maternal-fetal ECG separation is presented. Historical and current methods of monitoring the fetal heart are discussed, including the benefits and concerns with the standard practice of continuous fetal heart monitoring. The characteristics of fetal ECGs are compared to adult heart activity, and the properties of maternal-fetal ECG mixtures are introduced. These properties influence the signal separation techniques that are appropriate, including ICA which was introduced in Chapter 2. Several representations of maternal-fetal ECG recordings that can be used to test signal separation techniques are also presented.

4.1 Background on Fetal Heart Monitoring

The problem of assessing the health of a fetus has long been of interest to the medical community. This section summarizes the evolution of fetal monitoring techniques, with

emphasis on fetal heart monitoring and ways to capture the fetal ECG. Current standard practices in North America are discussed, including their limitations and the positive correlation between continuous fetal heart monitoring and the rise in Caesarean section deliveries.

4.1.1 Evolution of Fetal Heart Monitoring Techniques

Fetal heart monitoring in various forms has been in use for over a century, including attempts to record fetal ECGs by Cremer as early as 1906 [Lark58]. However, the recording equipment at that time did not have sufficient amplification to detect the small fetal heart beat signal consistently. Insufficient amplification devices continued to be a roadblock to fetal heart monitoring using electrical signals for at least the following 10-15 years [Skem58].

4.1.1.1 Early Fetal Heart Monitoring

By the 1950's, fetal ECG monitoring technology had improved enough to create a resurgence of interest in the topic. Initially, fetal heart monitoring was used as a fourth "positive sign" of pregnancy at approximately 20 weeks gestation. That is, simply the presence of a fetal heartbeat was sought to show that the fetus was still alive, without any attempts to assess the fetal heart health yet. The other three positive signs of pregnancy at the time were: listening for the fetal heartbeat; feeling for fetal movement; and taking x-rays of the fetus [Skem58]. Around 1960, fetal ECGs were explored as a tool to help reduce the number of perinatal deaths, nearly all of which occurred in the period before labour to within 48 hours after birth [Hare60]. It was estimated that such deaths accounted for 10% of all U.S.A. fatalities.

One of the dangers to fetal health that was addressed at this time with fetal heart

monitoring in the form of ECGs was the ability to determine fetal position [Lark58]. When the fetus is in vertex presentation, or head down, the QRS complexes of the fetal ECG point in the opposite direction from the mother's. However, if the fetus is in breech position, or feet down, the QRS complexes of both fetus and mother point in the same direction. This is an important characteristic, as breech births can present additional complications. Electrical fetal heart monitoring was also found to be useful in detecting multiple pregnancies, such as twins or triplets [Lark58]. The first successful detection of twins using this method was in 1938 [Novo59]. At that time there was no standard for electrode placement to measure fetal ECGs. While some set-ups used non-invasive surface electrodes, others required that the electrodes be placed in uncomfortable or invasive positions for the mother, or "applied with pressure so as to bring the electrodes closer to the fetal heart" [Lark58].

In addition to fetal heart monitoring by ECGs, fetal phonocardiography was explored during the twentieth century [Hare60]. This technique uses a microphone to record fetal heart sounds, and was first successfully implemented in 1923. However, it required further improvements to eliminate extraneous noise and to represent the sounds graphically in order to record the data and make it more accessible to physicians.

4.1.1.2 Current Fetal Heart Monitoring

For the most part, the tests used as positive signs of pregnancy continue to be in use today with the exception of x-rays. Due to concern over the effects of ionizing radiation on a developing fetus, and inconclusive limits on safe levels during pregnancy, x-rays are now used only in exceptional cases [Bren89]. In their stead, ultrasounds have become the preferred method for monitoring the development and position of the fetus, and have become standard practice during pregnancy in North America.

Continuous electronic fetal heart monitoring has also become standard practice during

labour and delivery in North America [Wool90][Davi93][Meni01]. It augments intermittent auscultation, which is the process of periodically listening to the fetal heartbeat. When the uterine contractions are measured in addition to the fetal ECG readings it is called *cardiotocography*. There are two main types of fetal heart monitoring, external and internal. External monitoring consists of sensors placed on the mother's abdomen which capture mixtures of the maternal and fetal ECGs as shown in the visual abstract. In contrast, internal monitoring uses a scalp electrode on the baby's head and is more accurate because it does not capture noise such as the maternal heartbeat. However, unlike external monitoring it requires the use of a needle electrode which can be a concern if infections are a risk factor.

4.1.2 Motivation: Current state of problem and consequences

Continuous fetal heart monitoring during labour and delivery is used to assess fetal distress. This type of monitoring is preferred by hospitals because it produces a permanent record of fetal heart activity, unlike intermittent auscultation. As well, some women may feel safer and reassured with continuous monitoring.

4.1.2.1 Concerns with Continuous Fetal Heart Monitoring

Despite the rise in use of continuous fetal heart monitoring, the output from non-invasive external fetal health monitors can be difficult to interpret due to the relative strong maternal ECG and noise interference from other biological signals [Albe93][Rode08]. In practice it is not always possible to reliably detect the features of the fetal heart wave, including even the strongest component, the QRS complex [Donk93]. Accurate detection of the QRS complex is necessary to calculate the fetal heart rate, which is critical in assessing fetal distress. More subtle features of the fetal ECG are even more difficult to extract,

and standards for interpreting data from non-invasive fetal heart monitors vary between hospitals [Davi93].

Invasive internal fetal health monitors are the current clinical solution, but present many complications of their own. These monitors can only be used during the latter stages of labour, unlike external monitoring which is possible throughout the pregnancy. Since internal monitoring uses a needle electrode there are risks of infection for the baby, particularly if the mother has any blood-born pathogens [Ferr01]. As well, some women are not comfortable with using needle electrodes for routine monitoring, as there are still concerns about fetal pain and even long-term consequences of this method.

4.1.2.2 Relationship Between Fetal Heart Monitoring and Caesarean Section Rates

Fetal ECGs and heart rate have become an influential factor in the decision of whether or not to perform an emergency Caesarean section, so reliable data is critical [Decl06][Clar07]. Since the introduction of continuous fetal heart monitoring, the number of unscheduled Caesarean sections deliveries has been on the rise [Decl06][Matt03]. By 2002, one-quarter of first-time mothers in the United States delivered with a Caesarean section. However, there has not been a corresponding rise in the number of high-risk pregnancies. Studies have attributed part of the increase in these operations to decisions made using unreliable fetal ECG data [Clar07][Anan13]. It is speculated that when the fetal ECG data is inconclusive, emergency Caesarean sections are performed out of caution. However, any false positives from the continuous fetal heart monitoring have consequences, as this is a major surgery with additional risks and healing time for the mother.

4.2 Characteristics of Maternal-Fetal ECGs

The maternal-fetal ECG recorded by non-invasive fetal heart monitors is a mixture of (i) the fetal ECG, (ii) the maternal ECG, and (iii) other electrical noise. Even though the fetal ECG is of interest, it is the much stronger maternal component that dominates the mixture. However, certain characteristics are known about both the fetal heartbeat itself and how the maternal-fetal mixture is created. This knowledge can be used to select techniques to extract the fetal ECG from the mixture.

4.2.1 Fetal ECG Characteristics

The fetal ECG signal closely resembles adult ECG recordings and can typically be detected near the six-to-seven week gestation point [Skem58]. Both adult and fetal ECGs contain the same repetitive, morphological pattern depicted in Figure 3.2 with a dominant QRS complex. As with adults, the fetal heart rate is measured as the average number of R-wave peaks per minute.

However, there are two main features that differ between the fetal and adult ECGs: the amplitude and the heartbeat frequency. The physically smaller fetal heart produces a weaker ECG signal than the much stronger maternal heart. Depending on sensor position and the natural variation in heart strength among individuals, the maternal heartbeat may be on the order of 10 to 1000 times stronger than the fetal heartbeat on non-invasive ECG recordings [Buxt63][Gupt07].

The fetal heart also beats at a higher frequency than an adult heart, and fetal heart rate has been closely linked to overall fetal heath. It is considered normal to have a fetal heart rate of 100-120 bpm at six to seven weeks gestation, increasing to 120-160 bpm at 30 weeks and levelling out at term to 110-150 bpm [Oudi04]. This is compared to a typical adult

heart rate of 60-100 bpm. A slower than average heart rate, also known as *bradycardia*, is detrimental to the fetus particularly in the first trimester [Doub95]. In one study the instances of first trimester fetal fatalities increased from 2.5% of those with normal heart rates to 15.6% of those with bradycardia [Ozte09]. A less common, but also dangerous, condition for the fetus is an abnormally fast heart rate or *tachycardia* [Oudi04].

4.2.2 Combined Maternal-Fetal Signals Characteristics

While non-invasive recording devices produce a combined maternal-fetal ECG signal, it is important to remember that this signal is produced by two separate, physical sources. This simple fact has two important consequences for methods of maternal-fetal ECG separation:

1. **Independence:** Since the fetal and maternal heartbeats are produced by two distinct, physical sources they can be assumed to be statistically independent from one another. Despite being a stronger muscle, the maternal heart does not influence the physiological processes that produce the fetal heartbeat.

2. **Linear Mixing:** The combined maternal-fetal ECG can be considered to be the weighted sum of the individual maternal and fetal ECGs. The weighting is dependent on the placement of the non-invasive sensors, which cause each of the these two source signals to be scaled by some factor in the recording.

These two characteristics of maternal-fetal ECGs, independence and linear mixing, make maternal-fetal ECG mixtures good candidates for blind source signal separation techniques to be applied to extract the source signals. ICA is particularly well-suited in this case, as neither the order nor amplification of the extracted signals is critical to the maternal-fetal ECG separation problem.

4.3 Maternal-Fetal ECG Separation Techniques

The task of separating the maternal and fetal heartbeats from a combined ECG recording has been one of the benchmark problems in signal processing since the 1970's, due to the complicated nature of the signals and their real-world clinical relevance [Widr75][Vand87]. However, since the maternal component of the signal is much stronger than the desired fetal heartbeat, traditional signal processing techniques that rely on the signal to noise ratio (SNR) are not applicable. There is currently no standard for isolating these signals and, as discussed in Section 4.1, there are concerns with current clinical techniques.

4.3.1 ICA for Maternal-Fetal ECG Separation

The maternal-fetal ECG separation problem is an application that is well-suited for ICA because it is predominately characterized by two independent, non-Gaussian signals: the heartbeat of the mother and that of the fetus. One of the first proposals to use ICA for the materal-fetal ECG separation problem was by De Lathauwer in 2000 [DeLa00]. Since then, different ICA algorithms have been applied to the problem, including FastICA, JADE, and AMUSE [Lee04][Naja06][Tara11]. These ICA experiments have produced promising results, but not the complete elimination of the maternal heartbeat and noise from the fetal ECG signal.

ICA has also been modified or used in conjunction with other methods to try to improve the speed and efficacy of maternal-fetal ECG separation [Xu12]. One proposed method uses both ICA and wavelets to circumvent the requirement of having the same number of recordings as source signals that ICA imposes [Pour06]. However, this method was not as successful as using ICA with multiple recordings. Another method used a modified version of ICA called *Multidimensional Independent Component Analysis (MICA)* [Cama11]. MICA assumes that the source signals can be separated into statistically independent groups,

while allowing that signals in the same group may be dependent. However, it was found that MICA did not produce consistently reliable results for maternal-fetal ECG separation.

The majority of the work in ICA separation of maternal-fetal ECG signals has two missing components. The first is that in the majority of cases the stationarity of the maternal and fetal ECGs is not taken into account, or they are assumed to be stationary [Sabr01]. As described in Chapter 3, it is well-known that ECG signals are non-stationary over long periods of time. However, they can be considered stationary over shorter intervals. Therefore it is important that the length of the maternal-fetal ECG recordings is considered before applying ICA. The second missing component is that there is no standardized method of quantifying how successful ICA was at maternal-fetal ECG separation. In many cases the extracted signals are visually verified to see if they look like typical ECG waveforms, or energy-based measures such as mean squared error are used [Pour06][Gupt07][Subh14][Sugu14]. These verification methods do not take the complexity of the signals into account.

4.3.2 Other Maternal-Fetal Separation Techniques

As maternal-fetal ECG separation is a benchmark problem in signal processing, as well as a problem of practical interest in medicine, numerous techniques other than ICA have also been attempted to varying degrees of success. The most common fetal ECG extraction techniques include:

- **Adaptive noise cancelling:** One of the first techniques to isolate the fetal ECG signal was to simultaneously record the maternal ECG separately, then use this to cancel out the strong maternal signal from the combined signal [Widr75][Gupt07][Liu11]. Adaptive filters are used for blind source separation problems because their parameters can be adjusted, unlike fixed filters. This technique is also called *adaptive noise filtering* because in this context the maternal ECG is considered to be noise.

Adaptive filtering can also be used in conjunction with other techniques, such as ICA [Gupt07][Khal13].

- **Component Analysis:** In addition to ICA, there are other component analysis techniques such as PCA and *periodic component analysis* (πCA). These techniques have been applied to maternal-fetal ECG separation with limited success [Llam13].
- **Singular value decomposition:** Based on matrix factorization, this technique has been used with some success on the maternal-fetal separation problem before the invention of ICA [Vand87]. It has also been used in conjunction with ICA [Gao03].

Other techniques used for maternal-fetal ECG separation include wavelets [Jafa05] [Lee04] [Pour06], the Hilbert transform method [Wils08], local linear projection [Schr96], artificial neural networks [Ye07], Kalman filtering [Venn14], and adaptive neuro-fuzzy inference systems [Nasi11][Subh14]. These methods have been applied with limited success, though none have fully solved the blind source separation problem. As well, some of these methods seek only to detect the fetal QRS complex, rather than to extract the full fetal ECG signal.

4.4 Experimental Maternal-Fetal ECG Signals

Ideally, to measure the effectiveness of a signal separation technique, the separated signals would be compared to the source signals. However, an essential feature of BSS problems is that the individual source signals are unattainable. In the case of maternal-fetal ECG separation, it is not possible to non-invasively record the fetal heartbeat directly. If there were, then by definition the problem would be solved.

In the absence of these direct signals, there are other ways to test maternal-fetal ECG separation techniques. Source signals can be simulated or synthesized from existing ECG

recordings. Maternal-fetal ECG mixtures can also be used to test signal separation techniques, although the results can not be explicitly verified.

4.4.1 Simulated Maternal-Fetal ECG Signals

Modelled signals, whether they are deterministic or probabilistic, can be useful to mimic physical signals because they can be made completely free of noise. For BSS problems, modelled signals have an additional benefit because they can be used as known source signals that are mixed and then separated again. The separated signals can then be quantitatively compared to the original, modelled source signals. By their very nature, it is not possible to directly obtain the source signals from BSS problems of physical processes, so the models provide a way to measure how well a BSS algorithm works.

The maternal-fetal ECG mixture can be modelled with the ECGsyn algorithm, which was introduced in Section 3.3. One of the advantages to this model is that it allows the user to set the average heart rate of the ECG signal in bpm. Figure 4.1 shows the simulated fetal ECG signal produced by ECGsyn in Matlab with an average heart rate of 120 bpm. The amplitude of the simulated fetal ECG has been reduced by a factor of ten.

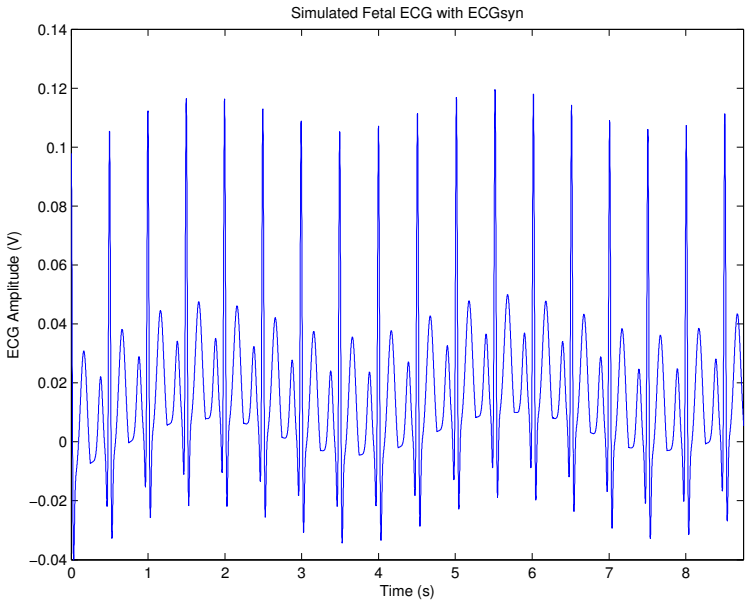

Fig. 4.1: Simulated fetal ECG with ECGsyn.

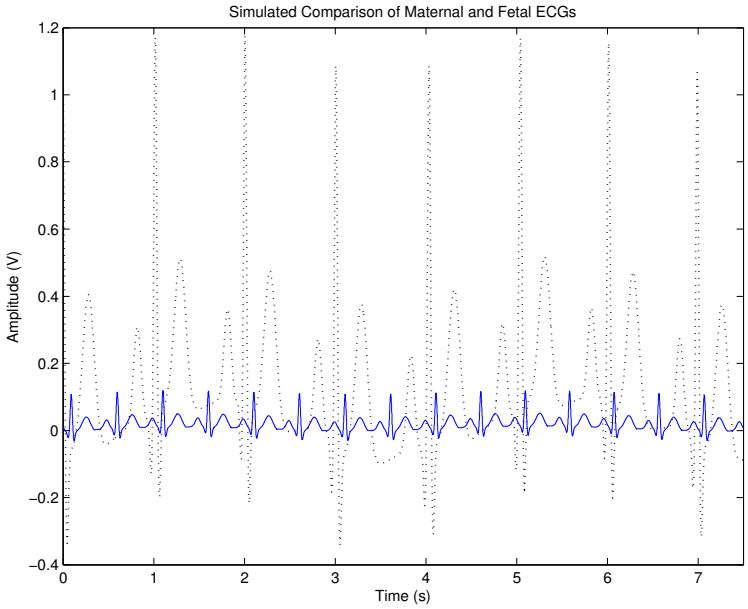

Fig. 4.2: Comparison of simulated maternal and fetal ECGs with ECGsyn.

The maternal ECG can be simulated in ECGsyn to have a different heart rate from the fetal signal. In Figure 4.2 the maternal ECG is simulated with an average heart rate of 60

bpm and shown in comparison with the simulated fetal ECG. The maternal ECG is drawn in a dotted line, while the fetal ECG is drawn in a solid line. The largest part of the fetal ECG, the QRS complex, is still smaller than the minor waves of the maternal ECG such as the P wave.

The resulting mixture from linearly combining the simulated fetal and maternal ECG signals is shown in Figure 4.3. The two simulated signals were added together to produce this noise-free model of what a non-invasive maternal-fetal ECG recording would look like. Only the QRS complex from the fetal ECG is visible amidst the much stronger maternal signal, and even this is aided by the fact that in this case they happen to occur in between the major maternal ECG peaks.

While the modelled signals do not fully capture all of the intricacies of the physical ECG signals, they provide a good substitute to test BSS techniques in a noise-free and fully defined environment. Furthermore, it is possible to model noise and add it to the maternal-fetal ECG model as well.

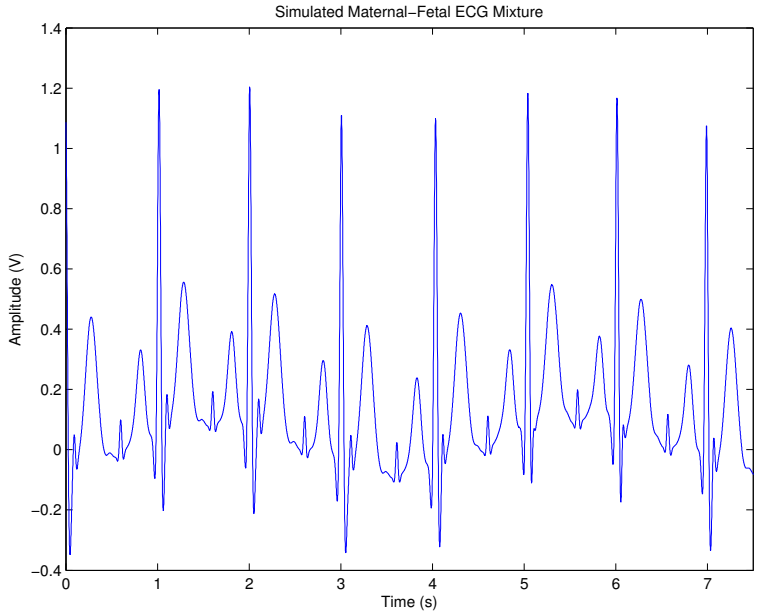

Fig. 4.3: Mixture combining the maternal and fetal ECGs simulations from ECGsyn.

4.4.2 Maternal and Fetal ECG Recordings from PhysioNet

In addition to the simulated signals produced by heart wave models, it is useful to perform experiments using recorded ECG data. As described in Section 3.4, the online resource PhysioNet hosts one of the largest open access repositories of electrical recordings of the heart, including healthy heart recordings and ones from patients with arrhythmias [Gold00].

4.4.2.1 Synthetic Maternal-Fetal ECG Signals

The *MIT-BIH Normal Sinus Rhythm Database* (nsrdb) on PhysioNet contains eighteen long-term ECG recordings from adult subjects that were assessed to have no significant arrhythmias, or irregular heartbeats. The recordings can be imported into Matlab for study and analysis using the PhysioToolKit provided by PhysioNet. Using recordings from this database, fetal heart signals and maternal-fetal recordings can be synthesized in order to produce mixtures in which the source signals are known. The synthesized mixtures can then be separated with BSS techniques such as ICA, and the results compared to the original source signals.

While the synthesized maternal-fetal ECG signals in this approach are created with two adult recordings, they can be scaled to represent fetal heart signals. Figure 4.4 shows a comparison of two nsrdb recordings scaled to represent the maternal and fetal ECGs. The maternal ECG is plotted as a dotted line and was not scaled. The fetal ECG was scaled in the time domain to have a heart rate of approximately 120 bpm, and scaled in amplitude to be ten times smaller than the original adult recording from which it was derived. The synthetic fetal ECG is plotted as a solid line.

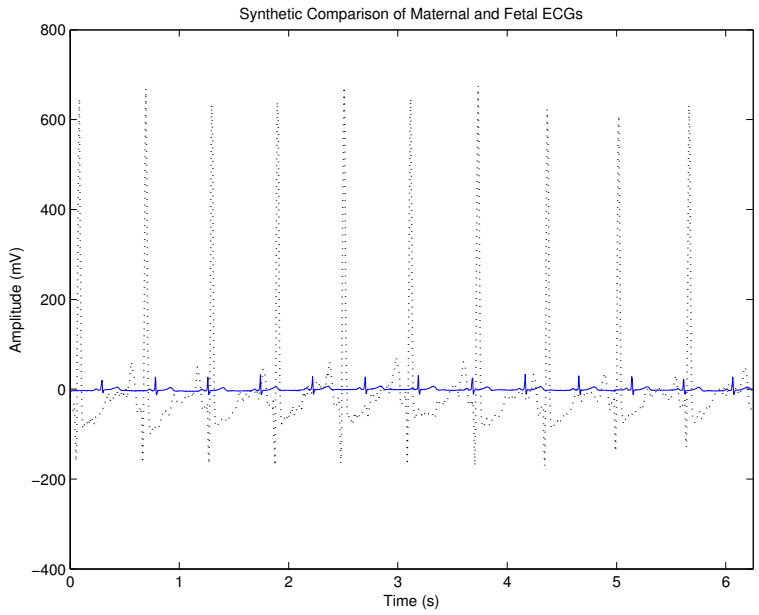

Fig. 4.4: Comparison of adult ECG recording and synthetic fetal ECG.

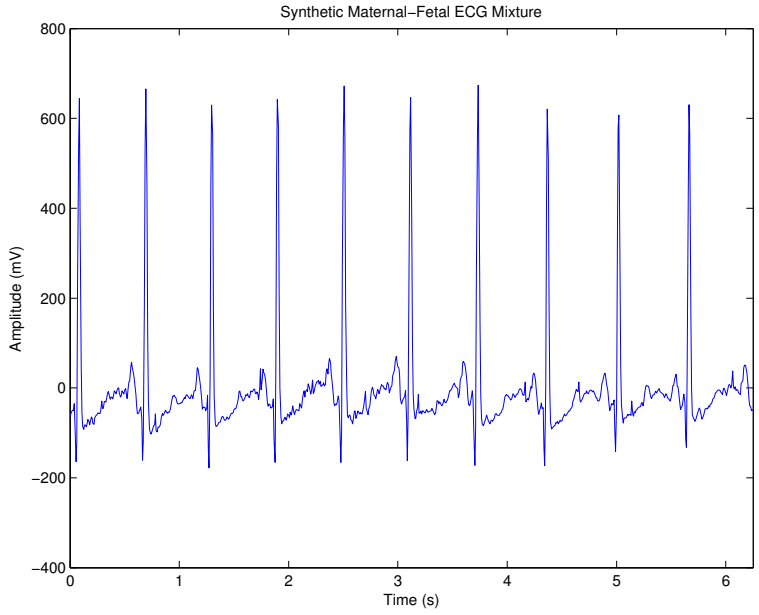

Fig. 4.5: Synthetic maternal-fetal ECG signal produced from nsrdb recordings.

To demonstrate a synthetic maternal-fetal ECG recording, the two signals in Figure 4.4

were added together to produce the synthetic mixture signal in Figure 4.5. In this mixture signal even the prominent QRS complex of the the fetal heartbeat is nearly undetectable by eye.

The two recordings used to create the synthetic fetal ECG and maternal-fetal ECG mixture depicted in Figures 4.4 and 4.5 were both measured from women of child-bearing age (ages 28 and 20, respectively).

4.4.2.2 Open Source Maternal-Fetal Recordings

In addition to adult ECG recordings, PhysioNet also contains a small number of non-invasive maternal-fetal ECG recordings that are available for public use. Each record in the *Non-Invasive Fetal ECG Database* (nifecgdb) contains three non-invasive maternal-fetal ECG abdominal recordings and two thoracic recordings of the mother's heartbeat. All five signals in each record were captured simultaneously, so they can be used to explore BSS methods such as ICA. One of the abdominal recordings is shown in Figure 4.6.

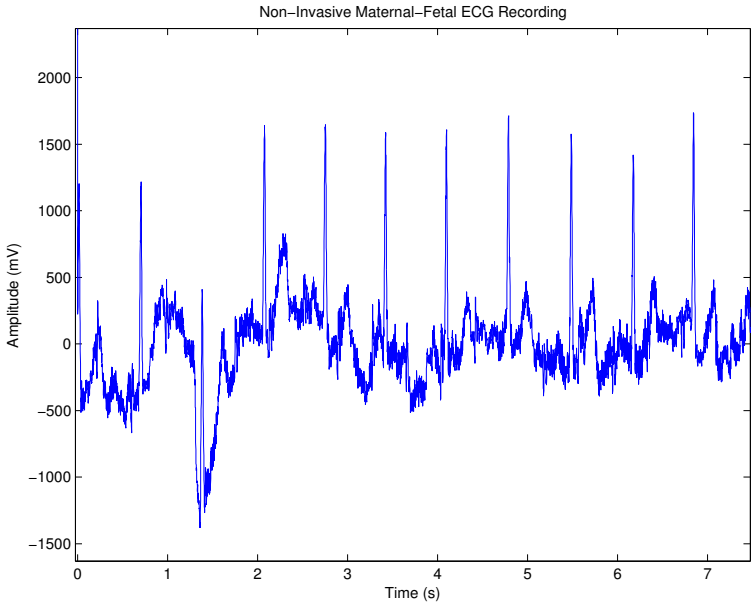

Fig. 4.6: Recorded non-invasive abdominal maternal-fetal ECG.

In the non-invasive maternal-fetal ECG recording in Figure 4.6 the fetal component of the signal is difficult to distinguish by eye. As well, there is more noise in this recording than in both the simulated and the synthetic ECG signals, which presents an added difficulty for maternal-fetal signal separation. Electrode placement and movement can contribute to the recorded noise. The need to use more sensitive equipment to capture the faster and weaker fetal heartbeats, compared to recording only an adult ECG, may also be a factor in the increased noise.

4.5 Maternal-Fetal ECG Separation Problem Summary

This chapter defined the problem of maternal-fetal ECG separation. The past and current fetal heart monitoring techniques were reviewed, including invasive and non-invasive continuous monitoring. The characteristics of fetal heart waveforms relative to adult heart waveforms were explored, noting that they can be considered scaled versions of the same signal archetype. Finally, methods of simulating and synthesizing maternal-fetal mixtures were presented in order to design experiments that recreate non-invasive maternal-fetal mixtures with known source signals.

Chapter 5

Case Study: ICA for Maternal-Fetal ECG Separation

In this chapter, the experimental results of ICA separation on simulated and recorded ECG signals are presented. First, ICA is used to separate mixtures of two adult ECG waveforms, in order to test its performance on this type of signal when amplitude differences are not a significant factor. Maternal-fetal ECG signals are modelled next using both simulated and synthetic data, then these signals are separated by ICA. To test how well ICA recovered the original source signals the energy, information theoretic, and fractal based measures described in Section 2.3.6 were calculated for each experiment. The heart rate of the source and extracted signals was also calculated and compared as a fifth measure. The fetal heart rate is one of the most critical components in determining fetal distress, so accurate recovery of the heart rate is essential to the success of any signal separation technique used for this purpose. These measures are used in the discussion in Chapter 6 to quantify the performance of the results.

ICA was also tested on non-invasive maternal-fetal ECG recordings. In this experiment

the original source signals are unknown, so the same success measures could not be used directly. However, they will be used indirectly in Chapter 6 to justify how well ICA worked by comparing the characteristics of the extracted signals to known averages.

5.1 Experiment: Two Adult ECGs (model)

In this ICA experiment, two adult ECG source signals modelled using ECGsyn were mixed and separated. As shown in Figure 5.1, the first signal had an average heart rate of 60 bpm while the second was 80 bpm. Both signals were created for a duration of two minutes, with a sampling frequency of 500 Hz and all other default values of the model. Only the first 20 seconds of each signal are shown in Figure 5.1, as well as in the following figures, in order to more clearly show the ECG waveforms.

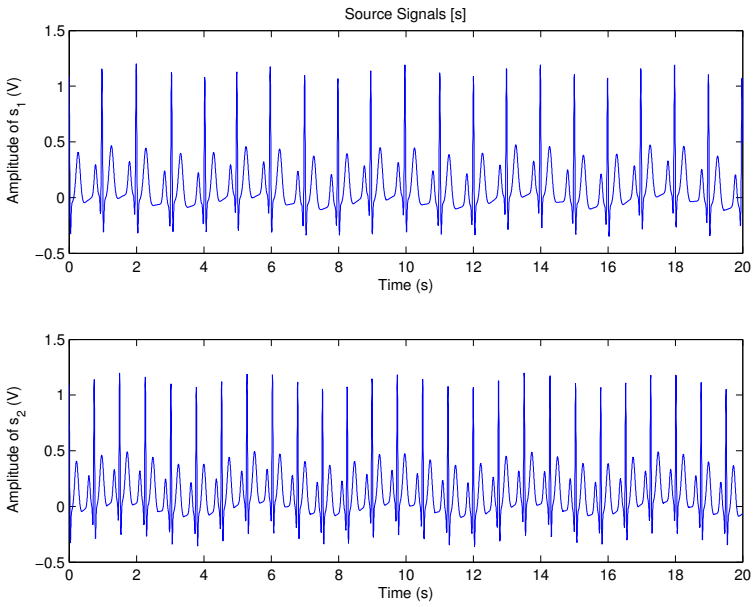

Fig. 5.1: Two simulated adult ECG source signals prior to mixing.

The two mixed signals, designated by vector $\mathbf{x}$, were created using the ICA mixing model, $\mathbf{x} = \mathbf{As}$, where $\mathbf{A}$ is the mixing matrix and $\mathbf{s}$ is a vector of the source signals. The mixing matrix was defined as:

$$\mathbf{A} = \begin{bmatrix} 2 & 3 \\ 2 & 1 \end{bmatrix} \tag{5.1}$$

so that each of the original source signals was stronger in one of the mixed signals.

The resulting mixed signals are shown in Figure 5.2, and the signals separated by ICA are shown in Figure 5.3.

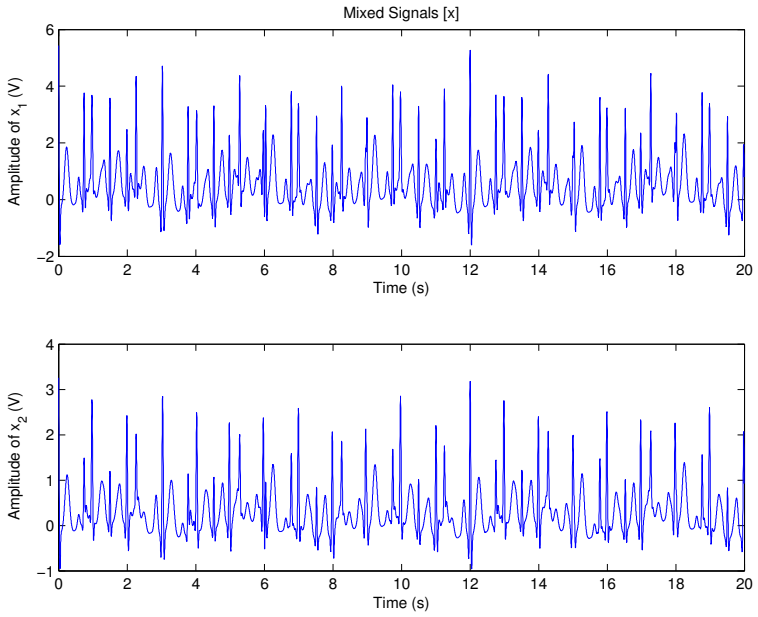

Fig. 5.2: Mixtures produced from two simulated adult ECG signals.

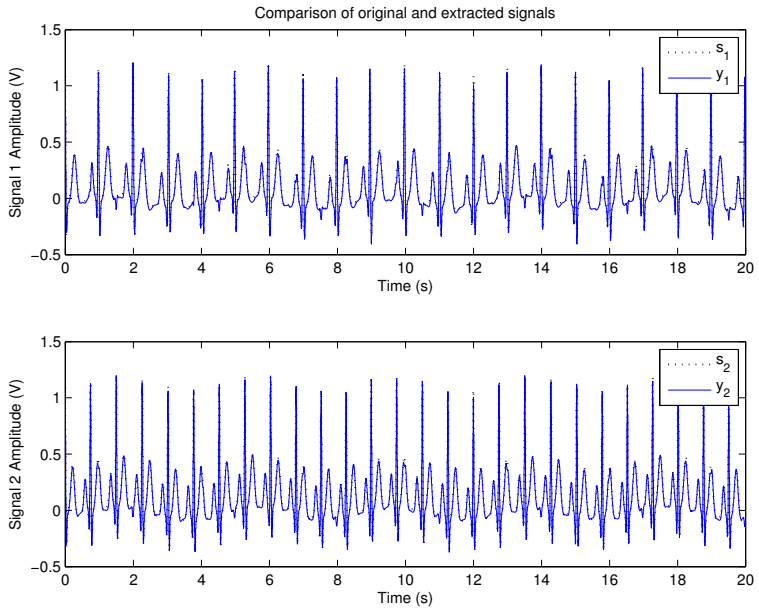

Fig. 5.3: The source and extracted simulated adult ECG signals are visually the same.

Figures 5.4 and 5.5 show a comparison of the estimated probabilities of the source, mixed, and extracted signals.

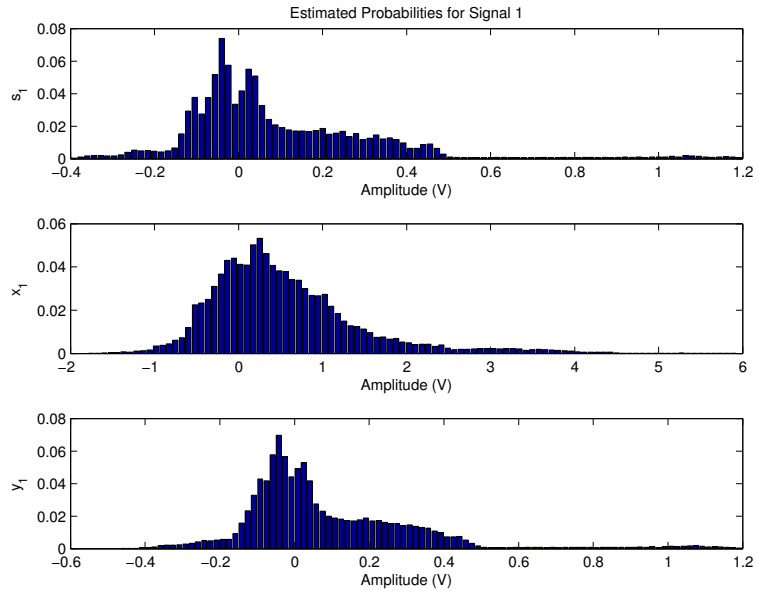

Fig. 5.4: Two simulated adult ECGs: estimated probabilities of s_1, x_1, and y_1.

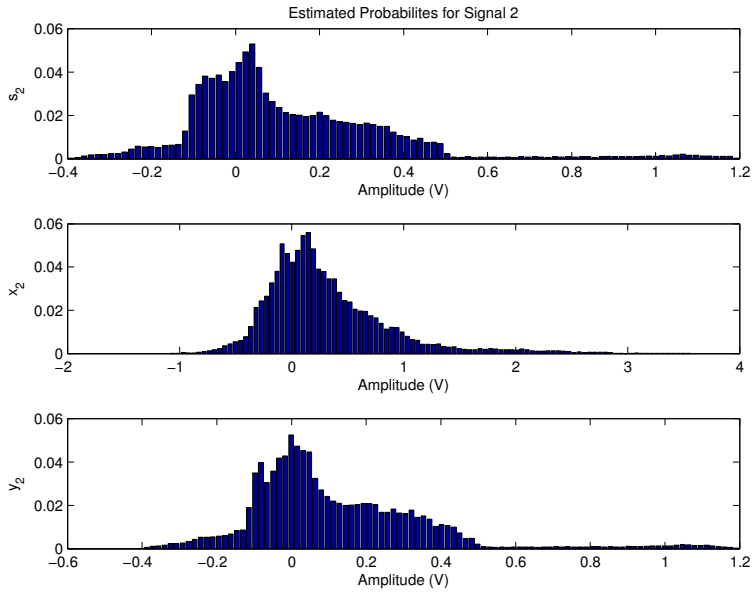

Fig. 5.5: Two simulated adult ECGs: estimated probabilities of s_2, x_2, and y_2.

In all of the examples in this section, the results are presented with the full precision returned by Matlab. The root mean square error (RMSE) between the first source signal and the corresponding extracted signal was $RMSE_1 = 0.020993$ V. Similarly, the error for the second set of signals was $RMSE_2 = 0.016738$ V.

The entropy of the first source signal was $H_{s_1} = 3.7908$, while the entropy of the corresponding extracted signal was $H_{y_1} = 3.76$. The change of entropy between the two signals was $\Delta H_1 = 0.030769$. For the second set of signals, the entropy of the source was $H_{s_2} = 3.8858$ and the entropy of the extracted signal was $H_{y_2} = 3.8779$ for a difference of $\Delta H_2 = 0.0078185$.

The average heart rate of the source signals was preserved in the extracted signals in both cases. For the first signal the average heart rate was 60.495 bpm and for the second signal it was 80.5234 bpm.

The spectral fractal dimension of the first source signal was $D_{\beta s_1} = 1.9624$, and

$D_{\beta y_1} = 1.918$ for the corresponding extracted signal. Therefore the change in spectral fractal dimension for the first signal was $\Delta D_{\beta 1} = -0.0444$. For the second signal, the source had a spectral fractal dimension of $D_{\beta s_2} = 1.7533$ and the for the extracted signal it was $D_{\beta y_2} = 1.87$. The change in fractal dimension for the second signal was $\Delta D_{\beta 2} = 0.1167$.

5.2 Experiment: Two Adult ECGs (recorded data)

The above experiment was repeated using recorded data from the MIT-BIH Normal Sinus Rhythm Database (nsrdb) on PhysioNet [Gold00]. The two source signals were two-minute ECG recordings from healthy female subjects aged 20 and 28. Due to their short duration, the ECG signals could be assumed to be stationary. As described in Section 4.3, this is standard practice in the ECG separation literature. All signals in the nsrdb have a sampling frequency of 128 Hz.

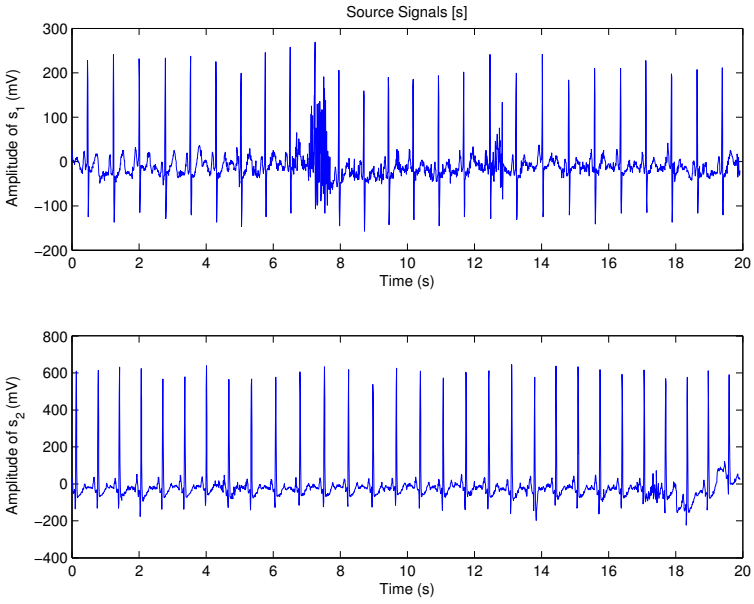

Fig. 5.6: Two recorded adult ECG source signals prior to mixing.

The first 20 seconds of each of the source signals are shown in Figure 5.6. Unlike the previous experiment with simulated data, the recorded source signals are inherently noisy and have more natural variation. Compared to the simulated source signals, the recorded signals show more variation in the ECG morphology and also exhibit artefacts.

The source signals were mixed using the same mixing matrix, defined in Equation 5.1, as the previous experiment. Figure 5.7 shows the resulting mixtures.

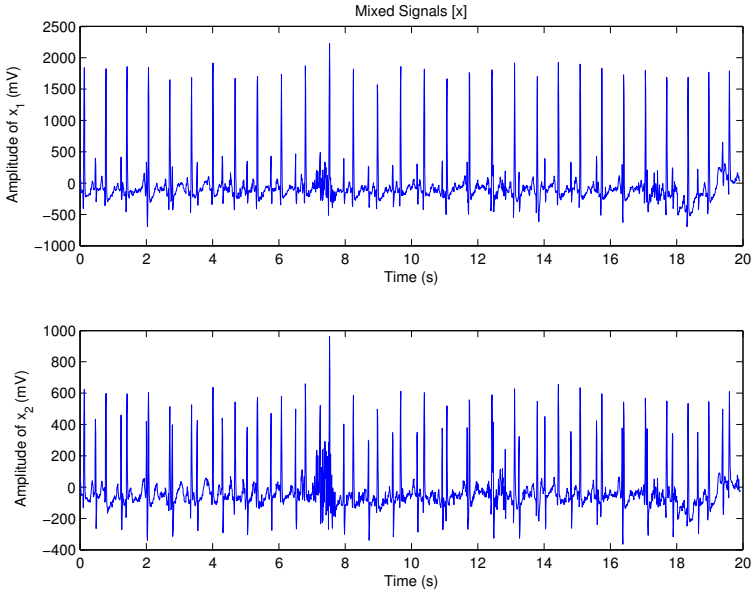

Fig. 5.7: Mixtures produced from two recorded adult ECG signals.

Figure 5.8 then shows a visual comparison of the source signals to the extracted signals after ICA separation. The extracted signals match the source signals so completely, that they overlap and are indistinguishable by eye. As can be seen in the comparison, even the irregularities of the source signals were extracted by ICA. In particular, the first signal exhibits a large artifact around the seven second mark that is uncharacteristic of the rest of the signal but is recovered by ICA.

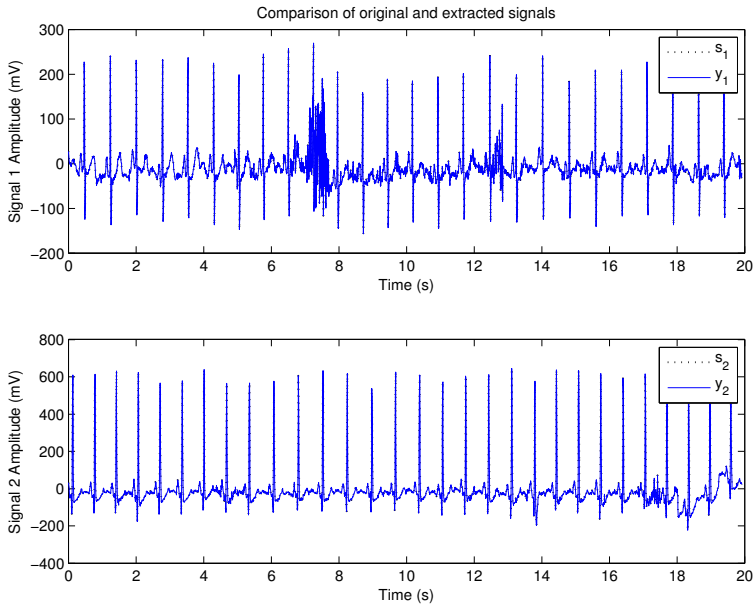

Fig. 5.8: The source and extracted recorded adult ECG signals are visually the same.

The estimated probabilities for the source, mixed, and extracted signals were also calculated. Figure 5.9 shows the estimated probabilities for the first signal, including s_1, x_1, and y_1. The estimated probabilities for the second signal, s_2, x_2, and y_2, are given in Figure 5.10. In both figures the estimated probabilities for the source and extracted signals look similar and skewed to the right, while the mixed signals appear to be more normally-distributed.

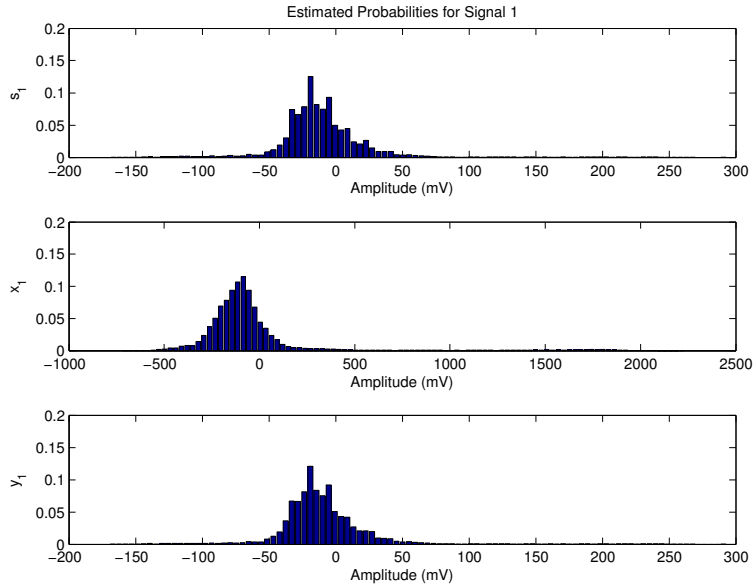

Fig. 5.9: Estimated probabilities of s_1, x_1, and y_1 from two recorded adult ECGs.

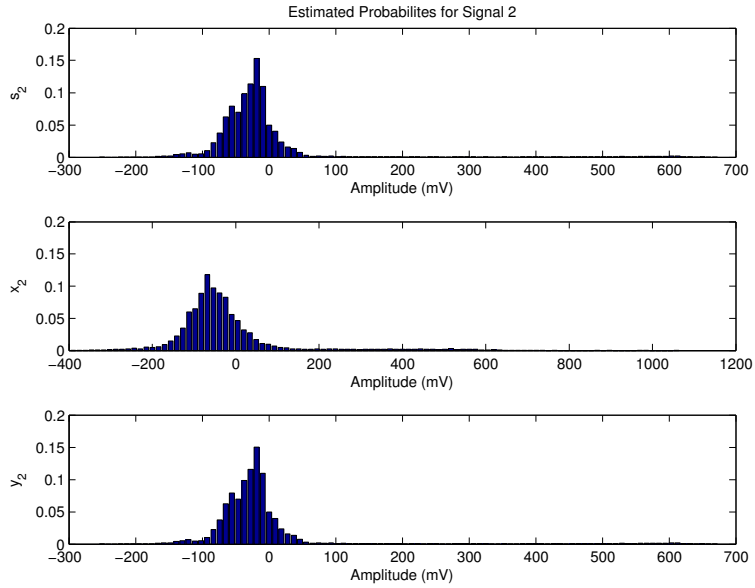

Fig. 5.10: Estimated probabilities of s_2, x_2, and y_2 from two recorded adult ECGs.

The root mean square error (RMSE) between the first source signal, s_1 and first extracted signal, y_1 was $RMSE_1 = 0.43758$ mV. For the second signal, the RMSE was $RMSE_2 = 0.054264$ mV.

The entropies of each of the source signals and extracted signals were calculated and compared. The entropy of the first source signal was $H_{s_1} = 3.1568$ and the entropy of the corresponding extracted signal was $H_{y_1} = 3.1681$. The change in entropy for the first signal was $\Delta H_1 = -0.011307$. For the second signal, the entropy of the source was $H_{s_2} = 2.9602$ and the entropy of the corresponding extracted signal was $H_{y_2} = 2.9508$. So, $\Delta H_2 = 0.0094196$ was the change in entropy for the second signal.

The heart rate was also calculated for the source and extracted signals. The signals recovered by ICA had the same heart rate as the original sources in both cases. The heart rate was 73.3342 bpm for the first signal, and 92.8082 bpm for the second signal.

The spectral fractal dimension for the first source signal was $D_{\beta s_1} = 1.64935$ and for the corresponding extracted signal it was $D_{\beta y_1} = 1.6494$, for a change of $\Delta D_{\beta 1} = 0.00005$. The second signal had a spectral fractal dimension of $D_{\beta 2} = 1.3321$ for both the source signal and extracted signal.

5.3 Experiment: Maternal-Fetal ECGs (model)

To simulate a non-invasive maternal-fetal ECG mixture, two separate signals were created with the ECGsyn model. The input average heart rate to the model was 60 bpm for the maternal signal and 110 bpm for the fetal signal. As in the previous experiment with two simulated adult ECGs, in this experiment the signals were approximately two minutes long, the sampling frequency was 500 Hz, and all of the other model parameters had their default values. The two source are displayed in Figure 5.11. In order to clearly see the ECG morphology, only the first 20 seconds of each signal is shown.

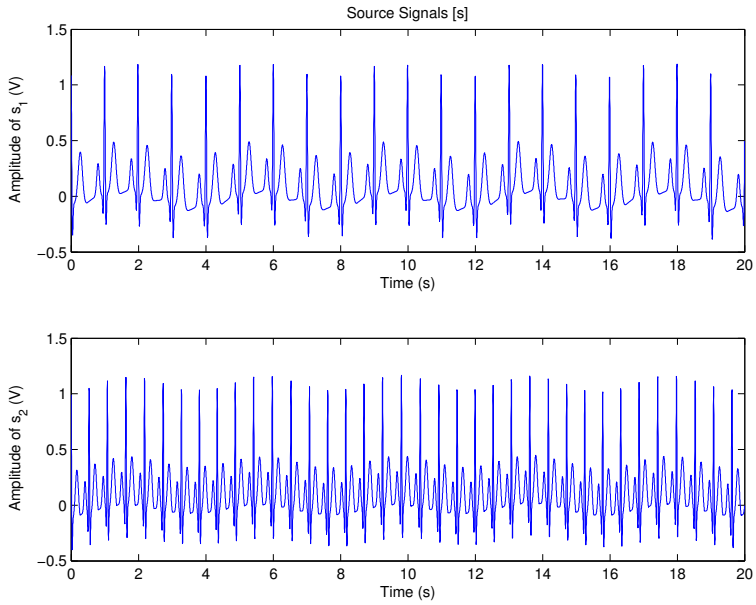

Fig. 5.11: Simulated maternal and fetal ECG source signals prior to mixing.

To represent the disparity in amplitudes between the maternal and fetal ECGs, the following mixing matrix was defined:

$$\mathbf{A} = \begin{bmatrix} 1000 & 1 \\ 1010 & 1 \end{bmatrix}. \tag{5.2}$$

With this mixing matrix, the maternal component was at least 1000 times stronger than then fetal component in both signal mixtures. As discussed in Chapter 4, the fetal component can be this small relative to the maternal heartbeat. This mixing matrix was also chosen to create two similar mixtures, under the assumption that non-invasive ECG data recorded simultaneously with two different electrodes could be quite similar depending on how close the electrodes are physically placed on the mother's abdomen. The signal mixtures produced from this mixing matrix are given in Figure 5.12.

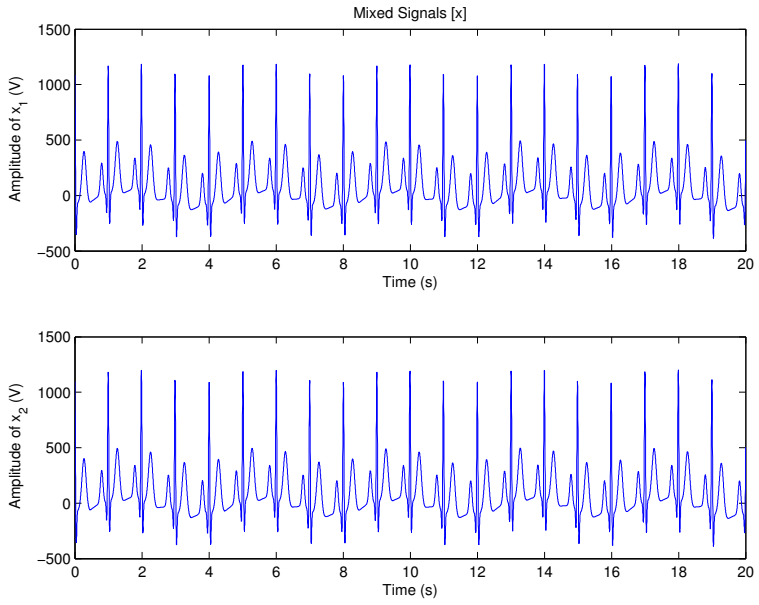

Fig. 5.12: Mixtures produced from the simulated maternal and fetal ECG signals.

As shown in Figure 5.12, the two mixtures are visually very similar. Further, the fetal component is so small that it is indistinguishable by eye, such that these mixtures appear to consist of only the maternal source signal.

The signals extracted by ICA are visually compared to the original source signals in Figure 5.13. The extracted signals overlap so completely with the source signals that they are indistinguishable in the figure.

Figure 5.14 shows the estimated probabilities for the source signal, s_1, mixed signal x_1, and extracted signal, y_1. Both the source and extracted signal have a significant positive skew, but the mixed signal has a more Gaussian distribution. The same observations can be made for the estimated probabilities of s_2, x_2, and y_2, which are shown in Figure 5.15

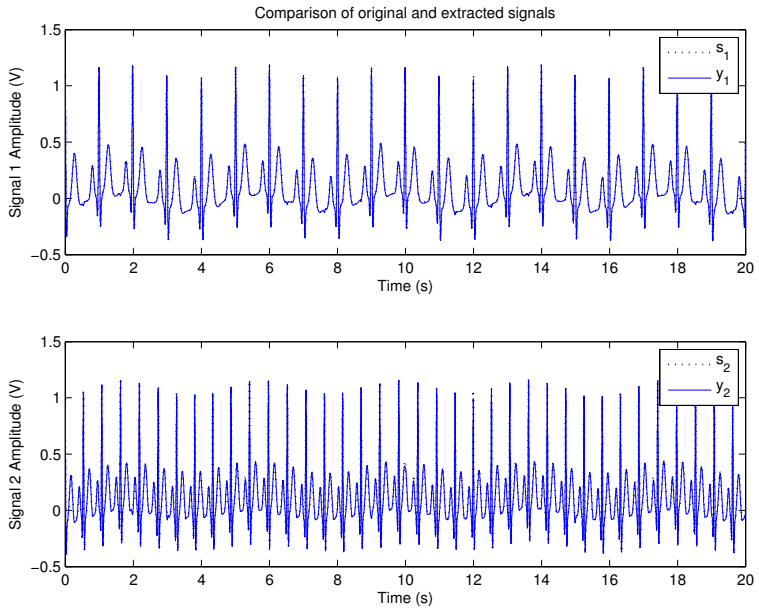

Fig. 5.13: The source and extracted simulated signals of the both maternal and fetal ECGs are visually the same.

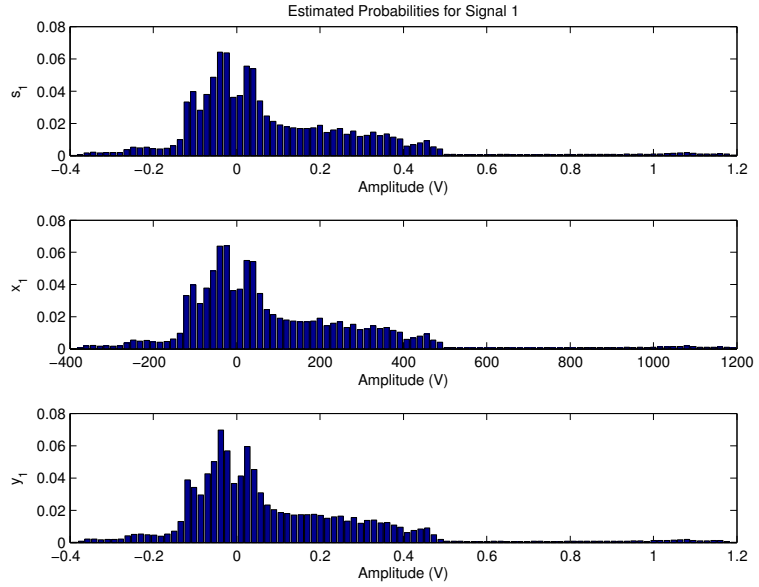

Fig. 5.14: Estimated probabilities of s_1, x_1, and y_1 from the simulated maternal ECG signal.

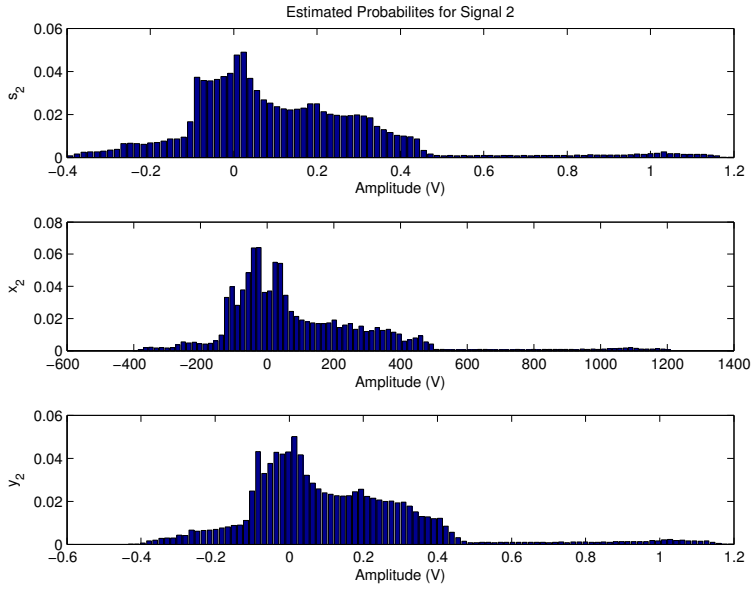

Fig. 5.15: Estimated probabilities of s_2, x_2, and y_2 from the simulated fetal ECG signal.

The root mean square error (RMSE) between the source and extracted signals of the maternal component was $RMSE_1 = 0.0055866$ V, and $RMSE_2 = 0.011359$ V between the synthetic fetal source signal and extracted signal.

The entropy of the maternal source signal was $H_{s_1} = 3.7893$, and it was $H_{y_1} = 3.7868$ for the maternal extracted signal. The difference in entropy before and after ICA for the maternal component was $\Delta H_1 = 0.0024542$. For the fetal signal, the entropy of the source was $H_{s_2} = 3.9275$, and the entropy of the extracted signal was $H_{y_2} = 3.903$. The change in entropy between the source and extracted signals was $\Delta H_2 = 0.024488$.

There was no change in the heart rate of the extracted signals from the source signals. The maternal component had a heart rate of 60.5235 bpm, while 110.52658 bpm was the heart rate of the second signal.

The maternal component had a spectral fractal dimension of $D_{\beta s_1} = 1.81925$ for the source signal, and $D_{\beta y_1} = 1.8173$ for the extracted signal; a change of $\Delta D_{\beta 1} = 0.00195$.

The spectral fractal dimension of the fetal component source signal was $D_{\beta s_2} = 1.81385$ and for the extracted signal $D_{\beta y_2} = 1.8062$. The change for the fetal component was $\Delta D_{\beta 2} = 0.00765$.

5.4 Experiment: Maternal-Fetal ECGs (synthetic data)

In order to mimic the non-invasive maternal-fetal ECG separation problem, while having full knowledge of the source signals for later comparison, synthetic signals were created from the nsrdb on PhysioNet [Gold00]. An ECG recording from a 28-year old woman was used to represent the maternal source signal. To synthesize the fetal source signal, another adult ECG recording was sped up to have a faster heart rate.

The source signals are shown in Figure 5.16, where the top subfigure is the maternal ECG and the bottom subfigure is the synthesized fetal ECG. Both source signals are two-minute long segments from the original recordings, but only 20 seconds are shown in the figure so the ECG waves can be clearly seen. The source signals representing maternal and fetal ECG components were mixed using the same mixing matrix as the previous experiment, given in Equation 5.2. Doing so produced mixtures where the maternal component was approximately 1000 times stronger than the fetal component, which is possible in non-invasive maternal-fetal ECG recordings. As shown in Figure 5.17, the mixtures both look like unaltered copies of the maternal ECG component because the synthetic fetal signal is so small.

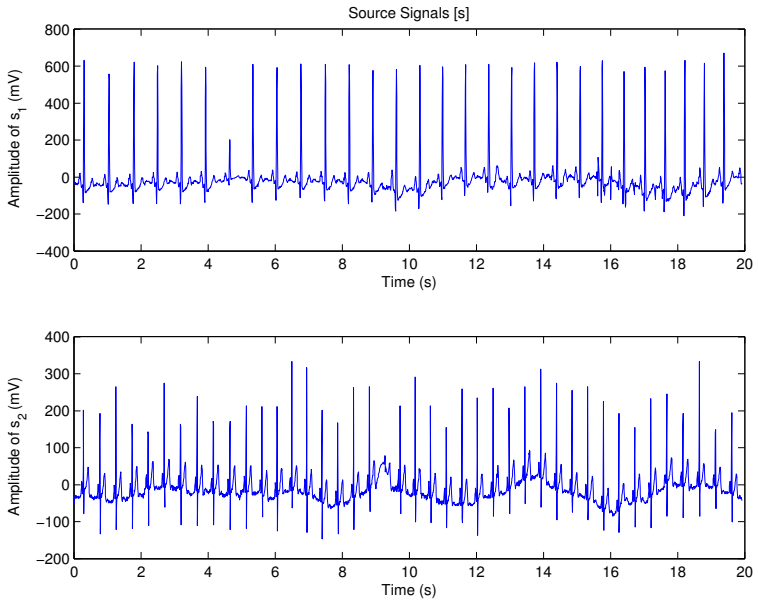

Fig. 5.16: The top is a recorded adult ECG used to represent the maternal component. The bottom is the fetal ECG synthesized from an adult recording.

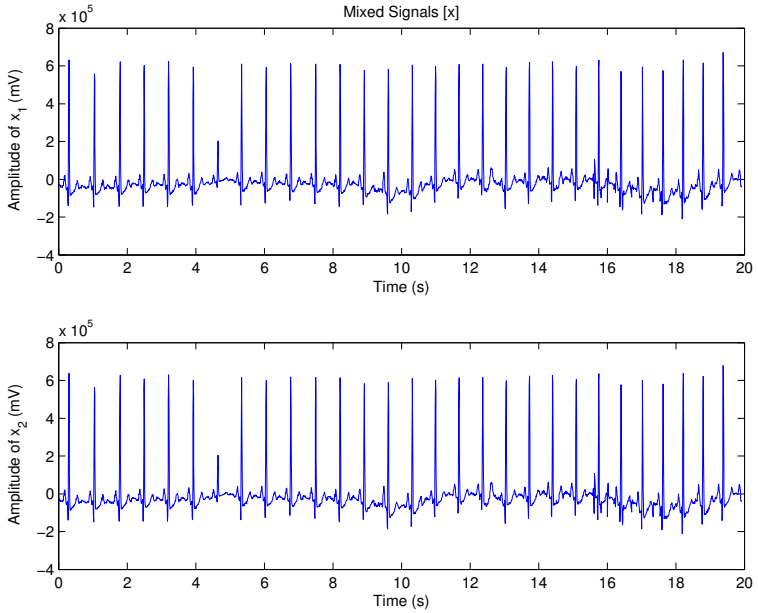

Fig. 5.17: The two mixtures produced from a recorded adult ECG signal and a synthetic fetal ECG signal.

The signals extracted by ICA are presented on the same graphs as their original source signals in Figure 5.18. Since the extracted signals so closely match the source signals, they nearly complete overlap on the graphs giving the appearance of only one signal on each. Even the extracted fetal component matches its source signal, despite being nearly invisible amidst the much larger maternal component in the signal mixtures.

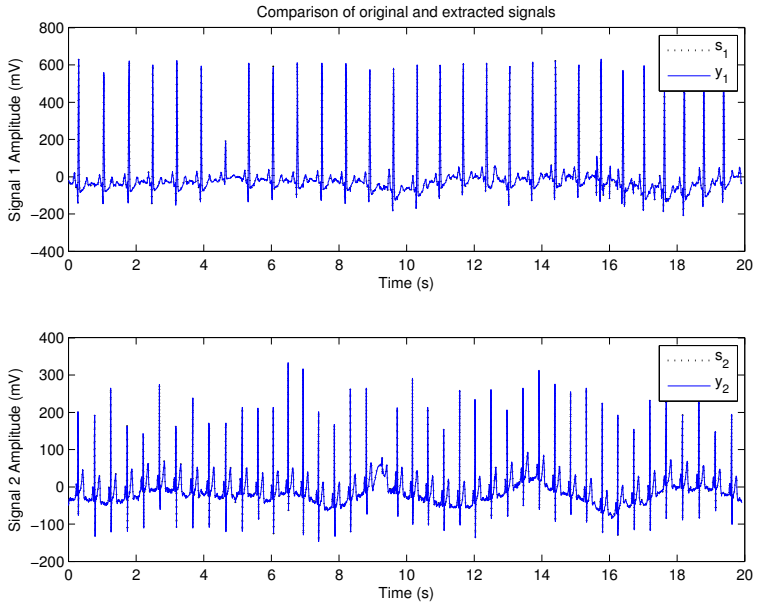

Fig. 5.18: The extracted signals representing the maternal and synthetic fetal components closely match the original source signals.

The estimated probabilities for the source signals, mixed signals, and extracted signals are given in Figures 5.19 and 5.20.

The root mean square error (RMSE) between the source and extracted signals of the maternal component was $RMSE_1 = 2.1487$ mV, and for the synthetic fetal component it was $RMSE_2 = 0.51913$ mV.

The entropy of the maternal source signal was $H_{s_1} = 3.1038$, and the entropy of the corresponding extracted signal was $H_{y_1} = 3.1054$. The change in entropy in the maternal component was $\Delta H_1 = -0.0016819$. For the synthetic fetal component, the source signal

entropy was $H_{s_2} = 3.0911$ and the extracted signal entropy was $H_{y_2} = 3.0964$. The change in entropy in the synthetic fetal component was $\Delta H_2 = -0.0053071$.

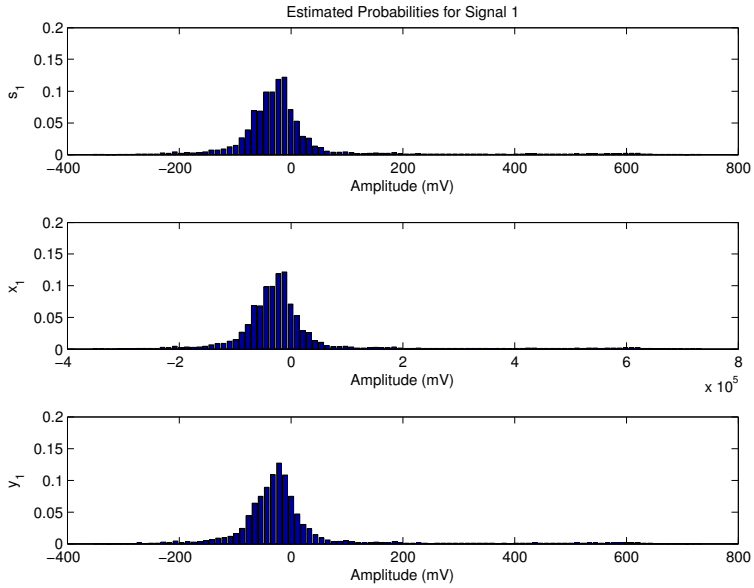

Fig. 5.19: Estimated probabilities of s_1, x_1, and y_1 from a recorded adult ECG.

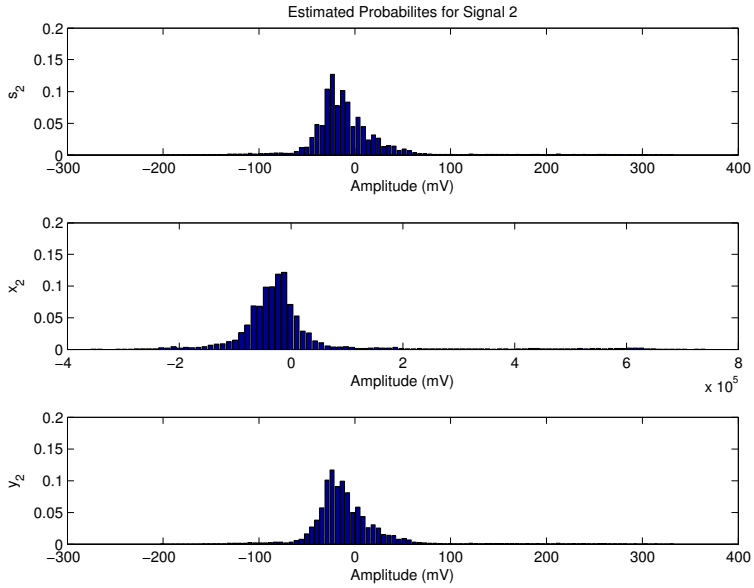

Fig. 5.20: Estimated probabilities of s_2, x_2, and y_2 from a synthetic fetal ECG.

The heart rate of the source signals and extracted signals was the same for both of the

maternal and fetal components. The maternal ECG had a heart rate of 91.9369 bpm, while the synthetic fetal heart rate was 134.86 bpm.

The spectral fractal dimension of the maternal source signal was $D_{\beta s_1} = 1.14895$ and for the corresponding extracted signal it was $D_{\beta y_1} = 1.15825$. The change in fractal dimension for the maternal component was $\Delta D_{\beta 1} = 0.0093$. For the synthetic fetal source signal, the spectral fractal dimensions were calculated to be $D_{\beta s_2} = 2.330085$ and $D_{\beta y_2} = 2.329855$ with $\Delta D_{\beta 2} = 0.00023$.

5.5 Experiment: Non-Invasive Abdominal Maternal-Fetal ECGs (recorded data)

In addition to the simulated and synthetic experiments above, the FastICA algorithm was tested using non-invasive maternal-fetal ECG recordings as well. These signals were obtained from the *Non-Invasive Fetal Electrocardiogram Database (ecgca)* on PhysioNet [Gold00]. Each record in the database consists of two thoracic signals and multiple abdominal signals that were recorded at the same time for one minute at a 1 kHz sampling frequency. There are no direct fetal ECG measurements in this database, so the signals extracted by ICA can not be directly compared to the original source signals.

In this experiment, two abdominal signals from the same record were used as input $\mathbf{x}$ to the ICA algorithm. Figure 5.21 shows the first 20 seconds of each mixture. Both the original source signals, $\mathbf{s}$, and the mixing matrix, $\mathbf{A}$, are unknown.

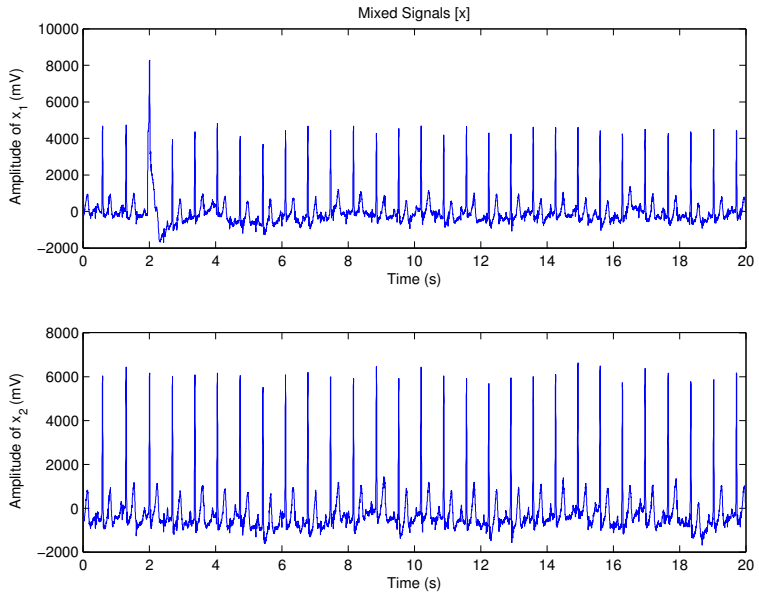

Fig. 5.21: Two abdominal maternal-fetal ECG recordings.

The extracted signals are shown in Figure 5.22. The top signal has periodic peaks but not the typical ECG waveform shape, while the bottom looks like an abdominal ECG.

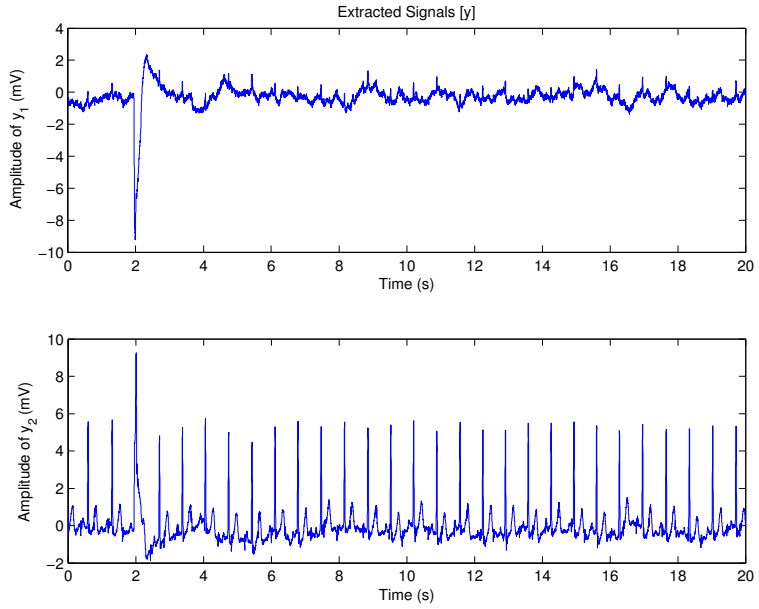

Fig. 5.22: Signals extracted by ICA from non-invasive abdominal maternal-fetal ECG recordings.

The estimated probabilities of the mixed and extracted signals are given in Figures 5.23, and 5.24, respectively.

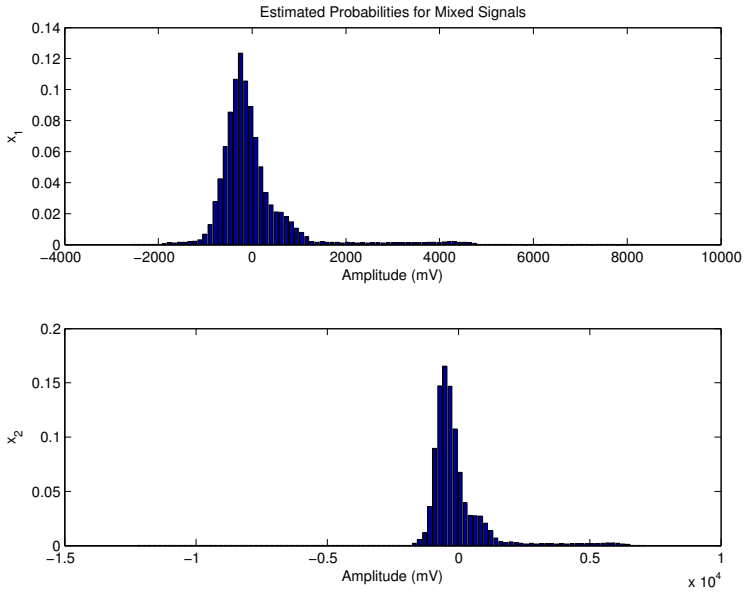

Fig. 5.23: Estimated probabilities for the non-invasive abdominal maternal-fetal ECG mixtures.

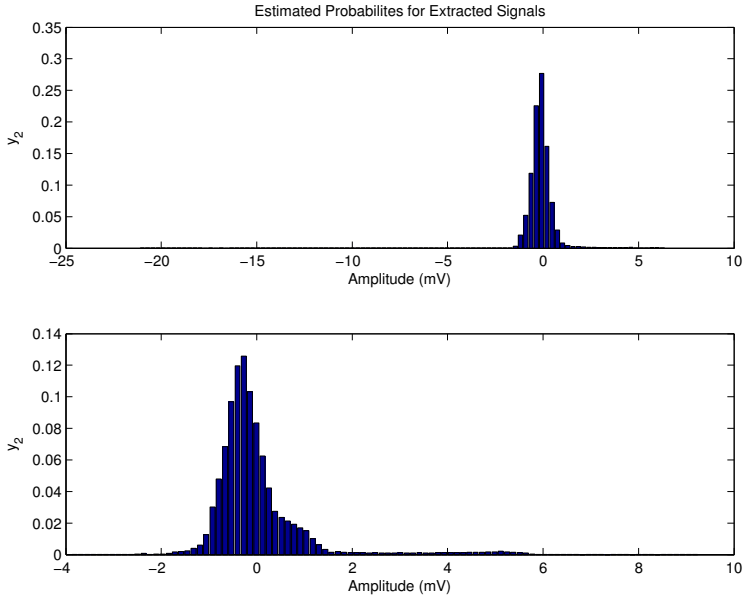

Fig. 5.24: Estimated probabilities for the signals extracted by ICA from the abdominal recordings.

The entropy of the first abdominal ECG recording was $H_{x_1} = 3.4988$, and $H_{x_2} = 3.3038$ for the second abdominal recording. The entropy of the first extracted signal was $H_{y_1} = 4.1347$, while the entropy of the second extracted signal was $H_{y_2} = 3.4625$.

The maternal heart rate obtained from the dominant R-waves in the abdominal ECG recordings was 88.3937 bpm. The first extracted signal did not have the ECG morphology, but did have peaks that occurred 121.4411 times per minute. The second extracted signal resembled the maternal ECG from the abdominal recordings and had a heart rate of 88.3937 bpm.

The spectral fractal dimensions of the abdominal mixture signals were $D_{\beta x_1} = 1.6808$ for the first signal and $D_{\beta x_2} = 1.659$ for the second signal. The extracted signals had spectral fractal dimensions of $D_{\beta y_1} = 1.651$ and $D_{\beta y_2} = 1.66455$.

5.6 Case Study Summary

This chapter presented experiments to test ICA on maternal-fetal ECG data. The first four experiments were designed to replicate the conditions of non-invasive maternal-fetal ECG mixtures from simulated and recorded data. In these experiments the source signals were known so they could be used for comparison with the signals extracted by ICA. Five effectiveness measures were used to evaluate the success of ICA at extracting the source signals: (i) visual comparison, (ii) RMSE, (iii) entropy, (iv) spectral fractal dimension, and (v) change in heart rate. The ECGsyn heart model was implemented to simulate ECG data, and recorded data was obtained from PhysioNet databases. The final experiment tested ICA on recorded non-invasive abdominal maternal-fetal ECG mixtures also obtained from PhysioNet. In this experiment the extracted signals could not be directly compared to the unknown source signals.

Chapter 6

Discussion of Experimental Results

In this chapter the results of the five experiments presented in Chapter 5 are discussed. The success of ICA at separating the mixed signals is justified with five qualitative and quantitative performance measures. Each of these measures is examined individually for the experiments with known source signals, to show that the signals extracted by ICA did compare favourably to the original signals. Differences between the experiments with simulated and recorded data are also highlighted. Finally, the subset of relevant performance measures are applied to the experiment with non-invasive maternal-fetal ECG recordings and found to be inconclusive. Possible reasons for the diminished performance of ICA in this experiment are also discusssed.

6.1 Discussion of Performance Measures

In this section the performance measures used for the first four experiments are discussed. All of these experiments had known source signals, so the ICA output could be compared directly to the original unmixed signals. The extracted signals were compared to the original sources (i) visually; (ii) with the energy-based RMSE metric; (iii) with the information-based Shannon entropy; (iv) by calculating the fractal dimension; and (v) by calculating the heart rate in each signal.

6.1.1 Visual Comparison

In Chapter 5 the signals extracted by ICA are shown on the same graphs as their corresponding source signals, giving a visual way to compare how closely they match. Visually comparing the source and extracted signals is a quick, intuitive way to evaluate the effectiveness of a separation technique. This is the predominant ICA verification technique for maternal-fetal ECG separation found in the literature.

The first four experiments in Chapter 5 all show an extremely close match between the source and extracted signals. In fact, the signals overlap so completely that they are indistinguishable on most of the graphs. Successful visual comparisons provide a quick check, and compelling argument, that ICA had worked on the known source signals. In the case of maternal-fetal ECG separation the extracted signals could be compared to ECG models, but the variability in what constitutes a healthy ECG make this impractical. There is also the risk of attempting to force an extracted signal to fit the model and overlooking true signal differences as noise.

In addition to comparing the signals as ECG time series, the histograms of the source and extracted signals can also be compared. In Chapter 5 the histograms of these signals

have been converted into estimated probabilities. In the first four experiments, the estimated probabilities of the source signals and extracted signals are visually similar. This provides additional evidence to support the success of the ICA separation technique. Furthermore, while the estimated probabilities of the source and extracted signal are similar, they are both different from the estimated probabilities of the mixture signals. Both the source and extracted signals have a distribution that is skewed to the right, as is common for electrophysiological signals. As expected, the mixture signals are more normally-distributed than the source signals. ICA exploits this difference to separate the mixture signals into less-Gaussian components.

Visual comparison of both the ECG time series and probability distributions can provide insight into the effectiveness of the ICA separation. However, these comparisons are qualitative and quantify the extent to which the signal separation was successful. By simply looking at the signals it is not possible to define how accurate an extracted signal is to the source, or to compare if one extracted signal is more or less accurate than another. A desire to quantitatively assess the signals separated by ICA leads to the other four methods used next: RMSE, entropy, fractal dimension, and heart rate.

6.1.2 Energy-Based Measure: RMSE

Another one of the more popular methods of evaluating the success of a signal separation technique such as ICA, is to compare the amplitudes of the source and extracted signals. That is, these methods compare the energies of the signals. The root mean square error (RMSE) is an energy-based metric to compare two signals. An advantage of the RMSE is that it assigns a value to the error, thus providing a base for comparison of how well two signals match. The ability of the RMSE to quantify error, or success, is an advantage of this metric over visual comparisons.

However, the RMSE is not necessarily a reliable measure for signal comparison. Like visual comparison, it is only possible to calculate the actual RMSE when the source signals are known. If the source signal was not known, the extracted signal could be compared to a model. However, any energy-based comparison between the extracted signal and a model is difficult when the original source signals are as naturally varied as ECGs.

Even when the source signals are known, there are drawbacks to the energy-based comparison methods. A small number of large amplitude differences can significantly affect any error measure based on energy, even if the rest of the signals match well. Furthermore, the error value produced by an energy-based measure such as the RMSE will be significantly affected by the amplitude range of the signals. As an example, a RMSE of 1mV may be insignificant if the source signal has a range of over 10V. However, that same RMSE may be very significant if the source signal has a range of only 5mV. Therefore, the size of the source signal must be taken into account if the RMSE is to be used for comparison purposes. This may not be possible if the source signal is not known and the extracted signal is being compared to a model.

Table 6.1: Summary of the experimental root mean square values

Two Adult ECGs	$RMSE_1$ (mV)	$RMSE_2$ (mV)
Simulated	20.993	16.738
Recorded	0.43758	0.054264

Maternal-Fetal ECGs	$RMSE_M$ (mV)	$RMSE_F$ (mV)
Simulated	5.5866	11.359
Synthetic	2.1487	0.51913

The values of the RMSE for the first four experiments presented in Chapter 5 are sum-

marized in Table 6.1. While all of the RMSE values are reasonably small, they vary widely for each of these experiments and do not provide very much insight into how successful ICA was at extracting the original signals.

6.1.3 Information-Based Measure: Entropy

Entropy can be thought of as a measure of the complexity of a signal. As with the RMSE, calculating the entropy produces a quantitative value that can be useful for comparisons. Entropy is calculated for each signal individually, so for BSS problems the entropy of the extracted signals can be calculated without prior knowledge of the source signals. However, if the entropy of source signals is known to be within a particular range, then the entropy of the extracted signals can provide an estimate of whether or not the results are reasonable.

Table 6.2 shows the Shannon entropy of the source and extracted signals from the experiments with known sources. The difference in entropy between the source signals and those extracted by ICA was minimal in all cases, supporting the assessment that source signals were recovered successfully.

Also of note, the entropies of the simulated signals were higher than the recorded signals on average. The simulated signals consistently had entropies in the range of $3.8 - 3.9$, while the entropies were closer to $3.0 - 3.1$ for the recorded signals. This suggests another difference between the ECGsyn model and recorded ECG signals.

Table 6.2: Summary of the experimental entropy values.

Two Adult ECGs	H_{s_1}	H_{y_1}	ΔH_1	H_{s_2}	H_{y_2}	ΔH_2
Simulated	3.7908	3.76	0.030769	3.8858	3.8779	0.0078185
Recorded	3.1568	3.1681	-0.011307	2.9602	2.9508	0.0094196

Maternal-Fetal ECGs	H_{s_M}	H_{y_M}	ΔH_M	H_{s_F}	H_{y_F}	ΔH_F
Simulated	3.7893	3.7868	0.0024542	3.9275	3.903	0.024488
Synthetic	3.1038	3.1054	-0.0016819	3.0911	3.0964	-0.0053071

6.1.4 Fractal-Based Measure: Spectral Fractal Dimension

Similar to the entropy, the fractal dimension of a signal can be used as a measure of its complexity. The spectral fractal dimension is appropriate for studying stationary, time-varying signals [Kins13]. Table 6.3 summarizes the spectral fractal dimensions calculated in each of the experiments with known source signals. As shown in the table, the spectral fractal dimension was robust under ICA. The extracted signals had nearly the same fractal dimension as their corresponding source signals, with D_β changing after ICA even less than the entropy.

As shown in Table 6.3, there is a difference in average fractal dimension between the simulated and recorded ECG signals. As was the case with entropy, the average fractal dimension of the simulated signals is higher than the recorded signals. This indicates another difference between the ECGsyn model and physical ECG signals.

The synthetic fetal ECG signals also had markedly different spectral fractal dimensions than either the simulated or recorded signals. This anomaly may be because the synthetic signals were created by compressing recorded ECG signals in the time domain, so the fractal dimension may not be meaningful when signals are manipulated in this way.

Table 6.3: Summary of the experimental calculated fractal dimensions.

Two Adult ECGs	$D_{\beta s_1}$	$D_{\beta y_1}$	$\Delta D_{\beta 1}$	$D_{\beta s_2}$	$D_{\beta y_2}$	$\Delta D_{\beta 2}$
Simulated	1.9624	1.918	-0.0444	1.7533	1.87	0.1167
Recorded	1.64935	1.6494	0.00005	1.3321	1.3321	0

Maternal-Fetal ECGs	$D_{\beta s_M}$	$D_{\beta y_M}$	$\Delta D_{\beta M}$	$D_{\beta s_F}$	$D_{\beta y_F}$	$\Delta D_{\beta F}$
Simulated	1.81925	1.8173	0.00195	1.81385	1.8062	0.00765
Synthetic	1.14895	1.15825	0.0093	2.330085	2.329855	0.00023

6.1.5 Heart Rate Preservation

Fetal ECG monitoring is predominately used to assess the fetal heart rate. The fetal heart rate has become a critical factor in determining fetal health and distress. Therefore, it is important to be able to accurately calculate the heart rate from the fetal signal extracted by ICA. In the four experiments with known source signals, peak detection methods were used to find the R-wave peaks and calculate the heart rates for each of the source and extracted ECG signals. Table 6.4 summarizes the heart rates of each signal in beats per minute.

Table 6.4: Summary of the derived experimental heart rates.

Two Adult ECGs	s_1 (bpm)	y_1 (bpm)	s_2 (bpm)	y_2 (bpm)
Simulated	60.495	60.495	80.5234	80.5234
Recorded	73.3342	73.3342	92.8082	92.8082

Maternal-Fetal ECGs	s_M (bpm)	y_M (bpm)	s_F (bpm)	y_F (bpm)
Simulated	60.5235	60.5235	110.52658	110.52658
Synthetic	91.9369	91.9369	134.86	134.86

As shown in the above table, despite any minor discrepancies between the source and extracted signals, the R-wave peaks were properly recovered for each signal by ICA so that the heart rate of each of the extracted signals was the same as its corresponding source signal. Therefore, ICA was successful at extracting signals with the same heart rates as the original signals.

6.2 Comparison Between Simulated and Synthetic Data

Between the trials with simulated ECG data and those with recorded signals, there were no statistically significant differences in the success of ICA based on the above measures. The RMSE was higher for the trials with simulated signals, however this is likely because the amplitude of the simulated signals was larger than that of the recorded signals. The changes in entropy and fractal dimension between the source and extracted signals were minimal for both types of data. However, as mentioned above, the simulated signals had larger entropies and fractal dimensions overall.

For both the simulated and synthetic data there was no significant difference in ICA performance between the trials with two adult ECGs and the trials with both maternal and fetal ECGs. This supports the claim that ICA is insensitive to amplitude differences between signals, such as maternal ECG signals and much smaller fetal ECG signals.

6.3 Discussion of Non-Invasive ECG Experiment Results

In the experiments with known ECGs as source signals, ICA was able to successfully recover the original signals as verified through the above measures. However, the experiment with unknown source signals was less conclusive. The signals in this experiment

came from non-invasive maternal-fetal ECG recordings, so the isolated source signals were inaccessible.

Visually, the first of the two signals extracted by ICA closely resembled the strong maternal ECG evident in the mixtures. The other extracted signal did not exhibit typical ECG wave patterns, however it did have peaks at periodic intervals which could be the fetal R-waves. Since the source signals were unknown the RMSE could not be calculated. The entropies of the extracted signals were slightly higher than in the first four experiments, which may be natural variation in the ECG signals or may indicate that the source ECGs were not fully recovered. The spectral fractal dimensions of the extracted signals were also on the high end of the range of what was found in the other experiments. However, the heart rates of the extracted signals showed promising results. The first extracted signal, that resembled the maternal ECG, had a heart rate of approximately 88 bpm which is within the healthy range for adults. For the second extracted signal, the periodic peaks were assumed to be R-waves so the heart rate could be calculated. This produced an approximate heart rate of 121 bpm, which is also within the healthy fetal range.

There are a number of factors that could have resulted in the less conclusive results of the experiment with non-invasive maternal-fetal ECGs. One difference between the non-invasive recorded ECGs and the signals used in the previous experiment was their length. The previous experiments used two-minute long signals, but the recordings in the *Non-Invasive Fetal ECG Database* (nifecgdb) were each one minute in length. The difference in duration may have affected the entropy and fractal dimension calculations. Finally, the noise levels in the recordings in the nifecgdb were unknown and may have been a contributing factor. This differs from the previous experiments with both simulated and recorded known source signals. There was no noise in the experiments with simulated ECG signals, and any noise in the obtained ECG recordings was considered part of the source signals. However, the non-invasive abdominal recordings likely included additive noise from other bioelectrical

activity such as muscle movements.

6.4 Discussion Summary

This chapter discussed the success of ICA at separating the ECG signals from the experiments presented in Chapter 5. All of the experiments with known source signals were considered successful in terms of the five effectiveness measures individually. However, when used together these measures present a fuller picture of how successfully the source signals were extracted by ICA. Differences between the experiments with simulated and synthesized data were also discussed. Finally, the results from the experiment with unknown source signals were discussed and found to be inconclusive.

Chapter 7

Conclusions

This thesis presented an assessment of ICA for the maternal-fetal ECG separation problem, including an investigation of its successes and current limitations. Chapter 2 began with an introduction to blind source separation problems in general and an introduction to the theory behind ICA. The signal separation technique of ICA is based on the critical concepts of minimum mutual information and maximum non-Gaussianity, without relying on a priori information about the mixed signals [Hyva00][Ston04]. Currently, there is no standard way to assess the success of signal separation methods, such as ICA, so four effectiveness measures were introduced: (i) visual comparison, (ii) energy-based comparison, (iii) information-based comparison, and (iv) fractal-based comparison. Also in Chapter 2, ICA was implemented in Matlab to perform experiments on both deterministic and stochastic signals [Hyva99].

The characteristics of ECGs were presented in Chapter 3, including their characteristic probability distribution to assess the signals suitability for ICA. Three heart models were also discussed: (i) ECGwaveGen, a deterministic model [Harr11]; (ii) ECGsyn, a stochastic model also known as the Oxford model [McSh03][McSh12]; and (iii) ECGfm, a stochastic

model created to more accurately capture heart rate variability [Pott08]. The characteristics of ECGs were used in Chapter 4 to discuss the maternal-fetal ECG separation problem. The social motivating factors of this problem were presented, including the current state of fetal heart monitoring techniques. The mixture models of maternal-fetal ECG were explored to determine that these signals are suitable for separation under ICA. Maternal-fetal signal mixtures were also simulated and synthesized from recorded data in this chapter.

Experimental results of ICA with ECG signals were presented in Chapter 5. Experiments were successfully performed on simulated and recorded ECG data, as well as on mixtures representing maternal-fetal ECGs. The four effectiveness measures previously introduced were used to assess the performance of ICA in these experiments with known source signals, along with heart rate as a fifth effectives measure. An experiment with unknown source signals using non-invasive maternal-fetal ECG recordings was also presented with inconclusive results. The results of these experiments were discussed in Chapter 6, as were the usefulness of the five effectiveness measures for ICA.

7.1 Thesis Conclusions

This thesis addressed a number of research questions related to blind source separation problems, ICA, ECG characteristics, and the maternal-fetal ECG separation problem. This section links back to the research questions first outlined in Section 1.2.3, to attempt to provide insight into them from the experimental results.

ICA was found to be insensitive to amplitude differences between source signals, both in the deterministic and stochastic experiments of Chatper 2 and the experiments with ECG data. However, other characteristics of the source signals were seen to have an effect on the success of ICA. When two Weierstrass signals were used as source signals, only one was recovered by ICA. Possible reasons for this include the complexity of the Weierstrass waves

or the discontinuous nature of the signals. It was also found that source signals with the same phase and frequency had different phases when extracted by ICA.

The success of ICA at separating signal mixtures into the original source signals was quantified using four measures, as well as heart rate for the ECG signals. The visual comparison were shown to be intuitive, but not quantitative. It was also found that an energy-based measure, such as the RMSE, was not sufficient on its own to indicate how well ICA worked. The entropy and spectral fractal dimension showed promise as additional verification measures. These two measures could work as a negative test, indicating when a signal extracted by ICA has significantly different characteristics from the source signals.

Probability distribution analysis of ECG signals showed that they are positively skewed, therefore having non-Gaussian distributions, so they are suitable for use with ICA. The fetal heartbeat was found to have the same characteristics as the maternal heartbeat, but with the waveforms occurring at more frequent intervals and a smaller amplitude. Therefore, the fetal ECG could be simulated using heart models created for adult ECGs by increasing their R-wave frequency of bpm. While the ECGsyn heart model simulates many of the characteristics of ECGs, the entropy and spectral fractal dimension calculations indicated systematic differences between the ECGsyn model and ECG recordings.

Research into current fetal heart monitoring indicated that it is performed either invasively or non-invasively. Invasive methods are non-ideal as they present additional discomfort and risks in terms of infection. However, recent trends suggest that current non-invasive monitoring increases the likelihood of false positives of fetal distress leading to increases in surgical intervention. Therefore, more accurate non-invasive fetal heart monitoring does have the potential to provide a tangible impact.

Yet, maternal-fetal ECG separation is still an unsolved problem. ICA was success-

ful in simulated and synthetic trials, as indicated by each of the five measures, but failed to perform as well on recorded maternal-fetal ECG mixtures. However, ICA did produce promising results in all experiments at extracting the R-waves for fetal heart rate calculation.

7.2 Contributions

This thesis contributes to the body of knowledge on ECG modelling, the maternal-fetal ECG separation problem, and signal separation effectiveness measures. The following are the main contributions:

(i) The thesis presents a comprehensive review on previous work in ICA and maternal-fetal ECG separation techniques.

(ii) Suggests and tests information-based and fractal-based measures to quantify and compare the success of ICA separation, in addition to the commonly used visual and energy-based measures. The experimental data also shows that energy-based measures such as the RMSE do not fully capture whether or not the signal separation was successful.

(iii) Identifies differences between the signals produced by the ECGsyn model and recorded ECGs. Both the entropy and spectral fractal dimension are shown to be consistently higher in the simulated signals when compared to the recorded signals.

(iv) Adds further test cases to the body of knowledge on maternal-fetal ECG separation to show that these signal mixtures are appropriate for separation with ICA.

(v) Identifies a gap between theory and practice for ICA separation of maternal-fetal ECG signals. Separation of simulated, synthetic, and recorded ECGs was successful in the cases where the source signals were known. However, the results of the experiment

with non-invasive maternal-fetal ECGs were inconclusive although the periodic peaks in the extracted signal were a promising indication that the fetal heart rate could be recovered.

7.3 Limitations and Future Work

Reliable maternal-fetal ECG separation, and blind separation of signals in general, is still an unsolved problem. In this work and similar studies, there are still a number of limitations and unanswered questions that provide rich ground for further research.

The main limitation of ICA in the context of maternal-fetal ECG separation is the gap between theory and practice. ICA separation was successful on the simulated and synthesized maternal-fetal ECG mixtures, but failed to completely separate the non-invasive recorded mixtures. Noise is likely a factor, and the type of electrical noise that is endemic to these recordings should be studied to better understand how it might affect signal separation attempts. The experiments with simulated data assumed a noiseless system, and in the experiments with recorded signals any noise was assumed to be inherent to the two source signals rather than a third, additive signal. Although ICA does not rely on signal-to-noise ratios, and is largely unaffected by amplitude differences between signals, electrical noise may be found to be a limiting factor in maternal-fetal ECG separation.

Additional experiments could be performed to address further questions, such as the impact of source signal frequency on the performance of ICA. Different pairs of non-invasive abdominal maternal-fetal ECG mixtures, that were recorded simultaneously, could also be tried with ICA.

The method of recording the ECG signals may also have affected the signal separation results. Sampling rates and signal resolutions that are acceptable for adult ECG recordings

may not be able to accurately capture the smaller and faster fetal ECG. Further work could assess the impact of recording fidelity on the successful sampling of the fetal R-wave peaks, and in turn its impact on the extracted fetal heart rate.

Related to the ability to reliably and accurately separate mixed signals, is the ability to verify that the separation has been successful and quantify how closely the extracted and source signals match. This thesis proposed five measures to qualitatively and quantitatively assess the success of ICA on maternal-fetal ECG separation, though there is still more to be done. Further testing and analysis will help to verify that these are meaningful measures, and work could be done to establish a ranking of the measures within the context of maternal-fetal ECG separation. Other measures could also be identified and investigated, such as the coefficient of variation.

Finally, the ICA algorithm required that pre-recorded data be used as input. It would be particularly advantageous for the maternal-fetal ECG separation problem, among other blind source separation problems, to be able to separate the signals in real-time. While this was outside the scope of this project, doing so would increase the effectiveness of the signal separation technique as a continuous fetal heart monitoring tool.

References

[Albe93] L. L. Albers and C. J. Krulewitch, "Electronic fetal monitoring in the United States in the 1980s," *Obstetrics and Gynecology*, vol. 82(1), pp. 8–10, 1993.

[Anan13] C. V. Ananth, S.P. Chauhan, H.-Y. Chen, M. E. D'Alton, and A. M. Vintzileos, "Electronic fetal monitoring in the United States: temporal trends and adverse perinatal outcomes," *Obstetrics and Gynecology*, vol. 121(5), pp. 927–933, 2013.

[Anan03] K. V. K. Ananthanag and J. S. Sahambi, "Investigation of blind source separation methods for extraction of fetal ECG," *in Proceedings of the Canadian Conference on Electrical and Computer Engineering (CCECE 2003)*, pp. 2012–2024, IEEE, Montreal, QC, 2003.

[Bell95] A. J. Bell and T. J. Sejnowski, "An information-maximization approach to blind separation and blind deconvolution," *Neural Computation*, vol. 7(6), pp. 1129–1159, 1995.

[Besi14] N. Besic, G. Vasile, J. Chanussot, and S. Stankovic, "Polarimetric incoherent target decomposition by means of independent component analysis," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53(3), pp. 1236–1247, 2014.

[Bren89] R. L. Brent, "The effect of embryonic and fetal exposure to x-ray, microwaves, and ultrasound: counseling the pregnant and nonpregnant patient about these risks," *Seminars in Oncology*, vol. 16(5), pp. 347–368, 1989.

[Buxt63] T. M. Buxton, I. Hsu, and R. H. Barter, "Fetal electrocardiography: a study of 254 patients using a preamplifer and a standard electrocardiograph," *Journal of American of the American Medical Association*, vol. 185(6), pp. 441–444, 1963.

[Calh03] V. D. Calhoun, T. Adali, L. K. Hansen, J. Larsen, and J. Pekar, "ICA of functional MRI data: An overview," *in Proceedings of the 4th International Symposium on Independent Component Analysis and Blind Signal Separation*, pp. 281–288, Nara, Japan, 2003.

[Cama11] J. L. Camargo-Olivares, R. Martín-Clemente, S. Hornillo-Mellado, M. M. Elena, and I. Román, "The maternal abdominal ECG as input to MICA in the fetal

ECG extraction problem," *IEEE Signal Processing Letters*, vol. 18(3), pp. 161–164, 2011.

[Card93] J.-F. Cardoso and A. Souloumiac, "Blind beamforming for non gaussian signals," *IEE Proceedings-F*, vol. 140(6), pp. 362–370, 1993.

[Clar07] S. L. Clark, M. A. Belfort, G. D. Hankins, J. A. Meyers, and F. M. Houser, "Variation in the rates of operative delivery in the United States," *American Journal of Obstetrics and Gynecology*, vol. 196(6), pp. 526.e1–526.e5, 2007.

[Como94] P. Comon, "Independent component analysis, A new concept?" *Signal Processing*, vol. 36(3287-314), 1994.

[Cove06] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, John Wiley & Sons, Inc., Hoboken, New Jersey, 2nd edn., 2006.

[Davi93] B. L. Davies, P. A. Niday, C. A Nimrod, E. R. Drake, A. E. Sprague, and M. Trépanier, "Electronic fetal monitoring: a Canadian survey," *Canadian Medical Association Journal*, vol. 148(10), pp. 1737–1742, 1993.

[DeLa00] L. De Lathauwer, B. De Moor, and J. Vandewalle, "Fetal electrocardiogram extraction by blind source subspace separation," *IEEE Transactions on Biomedical Engineering*, vol. 47(5), pp. 567–572, 2000.

[Decl06] E. Declercq, F. Menacker, and M. MacDorman, "Maternal risk profiles and the primary cesarean rate in the United States 1991-2002," *American Journal of Public Health*, vol. 96(5), pp. 867–872, 2006.

[Donk93] D. K. Donker, H. P. van Geijn, and A. Hasman, "Interobserver variation in the assessment of fetal heart recordings," *European Journal of Obstetrics & Gynecology and Reproductive Biology*, vol. 52(1), pp. 21–28, 1993.

[Doub95] P. M. Doubilet and C. B. Benson, "Embryonic heart rate in the first trimester: what rate is normal?" *Journal of Ultrasound in Medicine*, vol. 14(6), pp. 431–434, 1995.

[Ferr01] T. C. Ferree, P. Luu, G. S. Russell, and D. M. Tucker, "Scalp electrode impedance, infection risk, and EEG data quality," *Clinical Neurophysiology*, vol. 112(3), pp. 536–544, 2001.

[Gadh06] N. Gadhok, *A study of outliers for robust independent component analysis*, Master's thesis, University of Manitoba, 2006.

[Gadh04] N. Gadhok and W. Kinsner, "A study of outliers for robust independent component analysis," *in* 3, editor, *Canadian Conference on Electrical and Computer Engineering*, pp. 1421–1425, IEEE, 2004.

[Gao13] J. Gao, B. M. Gurbaxani, J. Hu, K. J. Heilman, V. A. Emanuele II, G. F. Lewis, M. Davila, E. R. Unger, and J.-M. S. Lin, "Multiscale analysis of heart rate

variability in non-stationary environments," *Frontiers in Physiology*, vol. 4(119), pp. 1–8, 2013.

[Gao03] P. Gao, E.-C. Chang, and L. Wyse, "Blind separation of fetal ECG from single mixture using SVD and ICA," *in Proceedings of the 2003 Joint Conference of the Fourth International Conference on Information, Communications and Signal Processing*, vol. 3, pp. 1418–1422, IEEE, 2003.

[Gian98] X. Giannakopoulos, J. Karhunen, and E. Oja, "An experimental comparison of neural ICA algorithms," *in Proceedings of the International Conference on Artificial Neural Networks*, Skövde, Sweden, 1998.

[Gian99] X. Giannakopoulos, J. Karhunen, and E. Oja, "A comparison of neural ICA algorithms using real-world data," *in International Joint Conference on Neural Networks*, vol. 2, pp. 888–893, IEEE, Washington, DC, 1999.

[Gold00] A. L. Goldberger, L. Amaral, L. Glass, J. M. Hausdorff, P. Ch. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C-K Peng, and H. E. Stanley, "PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals," *Circulation Electronic Pages*, vol. 101(23), pp. e215–e220, 2000.

[Gupt07] A. Gupta, M. C. Srivastava, V. Khandelwal, and A. Gupta, "A novel approach to fetal ECG extraction and enhancement using blind source separation (BSS-ICA) and adaptive fetal ECG enhancer (AFE)," *in Proceedings of the 6th International Conference on Information, Communications & Signal Processing*, pp. 1–4, IEEE, Singapore, 2007.

[Hare60] W. B. Harer Jr. and R. E. Davis, "The study of fetal problems with electronic methods," *American Journal of Medical Science*, vol. 239, pp. 766–772, 1960.

[Harr11] F. Harriott, "ECG waveform generator for Matlab/Octave," URL `http://www.physionet.org/physiotools/matlab/ECGwaveGen/`.

[Hayk05] S. Haykin and Z. Chen, "The Cocktail Party Problem," *Neural Computation*, vol. 17(9), pp. 1875–1902, 2005.

[Hayk01] S. Haykin and B. Kosko, editors, *Intelligent Signal Processing*, Wiley-IEEE Press, Piscataway, NJ, USA, 2001.

[Hild01] K. E. Hild II, D. Erdogmus, and J. Príncipe, "Blind source separation using Reny's mutual information," *IEEE Signal Processing Letters*, vol. 8(6), pp. 174–176, 2001.

[Horn01] P. Horn, "Autonomic computing: IBM's perspective on the state of information technology," Tech. rep., IBM, Armonk, NY, 2001, URL `http://people.scs.carleton.ca/~soma/biosec/readings/autonomic_computing.pdf`.

[Hyva13] A. Hyvärinen, "The FastICA package for MATLAB," URL `http://research.ics.aalto.fi/ica/fastica/`.

[Hyva99] A. Hyvärinen, "Fast and Robust Fixed-Point Algorithms for Independent Component Analysis," *IEEE Transactions on Neural Networks*, vol. 10(3), pp. 626–634, 1999.

[Hyva01] A. Hyvärinen, J. Karhunen, and E. Oja, *Independent component analysis*, Wiley, New York, NY, 2001, 481 pp.

[Hyva00] A. Hyvarinen and E. Oja, "Independent Component Analysis: Algorithms and Applications," *Neural Networks*, vol. 10(3), pp. 626–634, 2000.

[Jafa05] M. G. Jafari and J. A. Chambers, "Fetal electrocardiogram extraction by sequential source separation in the wavelet domain," *IEEE Transactions on Biomedical Engineering*, vol. 52(3), pp. 390–400, 2005.

[Jung00] T. P. Jung, S. Makeig, T. W. Lee, M. J. McKeown, G. Brown, A. J. Bell, and T. J. Sejinowski, "Independent component analysis of biomedical signals," *in Proceedings of the 2nd International Workshop on Independent Component Analysis and Signal Separation*, pp. 633–644, Swartz Center for Computational Intelligence, Helsinki, Finland, 2000.

[Jutt91] C. Jutten and J. Hérault, "Blind separation of sources, part I: An adaptive algorithm based on neuromimetic architecture," *Signal Processing*, vol. 24(1), pp. 1–10, 1991.

[Karh97] J. Karhunen, A. Hyvärinen, R. Vigário, J. Hurri, and E. Oja, "Applications of neural blind separation to signal and image processing," *in IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 131–134, IEEE, Munich, 1997.

[Kawa97] I. Kawachi, "Heart Rate Variability," URL `http://www.macses.ucsf.edu/research/allostatic/heartrate.php`.

[Khal13] J. Khalaf, B. Moslem, O. Bazzi, and M. O. Diab, "Fetal ECG extraction from abdominal composite recordings - a preliminary study," *in Proceedings of the 2nd International Conference on Advances in Biomedical Engineering*, pp. 133–136, IEEE, Tripoli, Lebanon, 2013.

[Kins05] W. Kinsner, "A unified approach to fractal dimensions," *in Proceedings of the 4th IEEE Conference on Cognitive Informatics*, pp. 58–72, IEEE, 2005.

[Kins10] W. Kinsner, "System complexity and its measures: How complex is complex," *in* Y. Wang, D. Zhang, and W. Kinsner, editors, *Advances in cognitive informatics and cognitive computing*, vol. 323 of *Studies in cognitive intelligence*, pp. 265–295, Springer, 2010.

[Kins13] W. Kinsner, "Fractal and chaos engineering course notes," Tech. rep., University of Manitoba, 2013.

[Klei87] R. E. Kleiger, J.P. Miller, J. T. Bigger Jr., and A. J. Moss, "Decreased heart rate variability and its association with increased mortality after acute myocardial infarction," *American Journal of Cardiology*, vol. 59(4), pp. 256–262, 1987.

[Lark58] S. D. Larks and K. Dasgupta, "Fetal electrocardiography, with special reference to early pregnancy," *American Heart Journal*, vol. 56(5), pp. 701–714, 1958.

[Lear03] E. G. Learned-Miller and J. W. Fisher III, "ICA Using Spacings Estimates of Entropy," *Journal of Machine Learning Research*, vol. 4, pp. 1271–1295, 2003.

[Lee04] J. Lee, S. P. Cho, and K. K. J. Lee, "A study of fetal ECG separation from small channel abdominal ECGs using ICA," *in Proceedings of the 30th Annual Conference of the IEEE Industrial Electronics Society*, vol. 3, pp. 3216–3218, IEEE, Busan, Korea, 2004.

[Lins89] R. Linsker, *An application of the principle of maximum information preservation to linear systems*, vol. 1, pp. 186–194, Morgan Kaufmann Publishers Inc., San Francisco, CA, 1989.

[Liu11] S.-J. Liu, D.-L. Liu, J.-Q. Zhang, and Y.-J. Zeng, "Extraction of fetal electrocardiogram using recursive leasst squares and normalized least mean squares algorithms," *in Proceedings of the 3rd International Conference on Advanced Computer Control*, pp. 333–336, IEEE, Harbin, 2011.

[Llam13] M. Llamedo, A. Martín-Yebra, P. Laguna, and J. P. Martínez, "Noninvasive fetal ECG estimation based on linear trasformations," *in Computing in Cardiology Conference*, pp. 285–288, IEEE, Zaragoza, 2013.

[Matt03] T. G. Matthews, P. Crowley, A. Chong, P. McKenna, C. McGarvey, and M. O'Regan, "Rising caesarean section rates: a cause for concern," *BJOG: An International Journal of Obsstetrics and Gynaecology*, vol. 110(4), pp. 346–349, 2003.

[McSh12] P. McSharry and G. Clifford, "ECGSYN: a realistic ECG waveform generator," URL `http://www.physionet.org/physiotools/ecgsyn/`.

[McSh03] P. McSharry, G. Clifford, L. Tarassenko, and L. Smith, "A dynamical model for generating synthetic electrocardiogram signals," *IEEE Transactions on Biomedical Engineering*, vol. 50, pp. 289–294, 2003.

[Meni01] C. A. Menihan, *Electronic fetal monitoring: concepts and applications*, Lippincott, Philadelphia, USA, 2001.

[Pott02] N. G. Michael Potter and W. Kinsner, "Separation performance of ICA on simulated EEG and ECG signals contaminated by noise," *in Proceedings of the*

Canadian Conference on Electrical and Computer Engineering (CCECE 2002), vol. 2, pp. 1099–1104, IEEE, Winnipeg, MB, 2002.

[Miet02] J. Mietus, C-K Peng, I. Hentry, R. L. Goldsmith, and A. L. Goldberger, "The pNNx files: re-examining a widely used heart rate variability measure," *Heart*, vol. 88, pp. 378–380, 2002.

[Miho02] M. Mihoko and S. Eguchi, "Robust blind source separatino by beta divergence," *Neural Computation*, vol. 14(8), pp. 1859–1886, 2002.

[Mous13] Z. Moussavi, "Biomedical instrumentation and signal processing course notes," Tech. rep., University of Manitoba, 2013.

[Naja06] F. S. Najafabadi, E. Zahedi, and M. A. M. Ali, "Fetal heart rate monitoring based on independent component analysis," *Computers in Biology and Medicine*, vol. 36, pp. 241–252, 2006.

[Nasi11] M. Nasiri, K. Faez, and A. M. Nasrabadiw, "A new method for extraction of fetal elecrocardiogram signal based on adaptive nero-fuzzy inference system," *in Proceedings of the 2011 IEEE International Conference on Signal and Image Processing Applications*, pp. 456–461, IEEE, Kuala Lumpur, 2011.

[Novo59] C. A. Novotny, W. K. Hass, and D. A. Callagan, "Early diagnosis of multiple pregnancy: use of electroencephalograph in prenatal examination," *Journal of American of the American Medical Association*, vol. 171(7), pp. 116–120, 1959.

[Oja00] E. Oja, K. Kiviluoto, and S. Malaroiu, "Independent component analysis for financial time series," *in The IEEE Adaptive Sysmtes for Signal Processing, Communications, and Control Symposium*, pp. 111–116, IEEE, Lake Louise, Alberta, 2000.

[ORei73] F. J. O'Reilly and C. P. Quesenberry, "The conditional probability integral transformation and applications to obtain composite chi-square goodness-of-fit tests," *The Annals of Statistics*, vol. 1(1), pp. 74–83, 1973.

[Oudi04] M. A. Oudijk, G. H. A. Visser, and E. J. Meijboom, "Fetal tachyarrhythmia - part 1: diagnosis," *Indian Pacing Journal and Electrophysiology*, vol. 4(3), pp. 104–113, 2004.

[Ozte09] D. Oztekin, O. Oztekin, F. I. Aydal, S. Tinar, and Z. H. Adibelli, "Embryonic heart rate as a prognostic factor for chromosomal abnormalities," *Journal of Ultrasound in Medicine*, vol. 28(5), pp. 609–614, 2009.

[Pott08] M. Potter, *Convergence of dynamical features under ICA with application to fetal ECG*, Phd thesis, University of Manitoba, 2008.

[Pott01] M. Potter and W. Kinsner, "Competing ICA techniques in biomedical signal analysis," *in IEEE Canadian Conference on Electrical and Computer Engineering (CCECE 2001)*, vol. 2, pp. 987–992, IEEE, Toronto, ON, 2001.

[Pott09] M. Potter and W. Kinsner, "Improved event interval reconstruction in synthetic electrocardiograms," *IEEE Transactions on Biomedical Engineering*, vol. 56(3), pp. 900–903, 2009.

[Pour06] M. Poursoltani, A. Boroomand, and A. Ayatollahi, "Comparing higher order statistics of three ICA methods in wavelets-based single-channel fetal ECG extraction," *in International Conference on Biomedical and Pharmaceutical Engineering*, pp. 162–167, IEEE, Singapore, 2006.

[Rist99] T. Ristaniemi and J. Joutsensalo, "On the performance of blind source separation in CDMA downlink," *in Proceedings of the International Workshop on Independent Component Analysis and Blind Source Separation*, pp. 437–441, Aussois, France, 1999.

[Robe01] S. Roberts and R. Everson, *Independent component analysis: principles and practice*, Cambridge University Press, Cambridge, United Kingdom, 2001.

[Rode08] C. H. Rodeck and M. J. Whittle, *Feal medicine: basic science and clinical practice*, Churchill Livingstone, London, UK, 2nd edn., 2008.

[Rudi07] M. Rudinac, M. Uscumlic, S. Rudinac, B. Milovanovic, I. Reljin, and B. Reljin, "Fractal and multifractal analysis of heart rate variability," *in Proceedings of the 8th International Conference on Telecommunications in Modern Satellite, Cable and Broadcasting Services*, pp. 325–328, IEEE, Serbia, 2007.

[Sabr01] M. Sabry-Rizk, W. Zgallai, A. McLean, E. R. Carson, and K. T. V. Grattan, "Virtues and vices of source separation using linear independent component analysis for blind source separation of non-linearly coupled and synchronised fetal and mother ECGs," *in Proceedings of the 23rd Annual EMBS International Conference*, pp. 1985–1989, IEEE, Istanbul, Turkey, 2001.

[Sach98] F. Sachse, C.D. Werner, K. Meyer-Waarden, and O. Dössel, "Applications of the visible man dataset in electrocardiology: calculation and visualization of body surface potential maps of a complete heart cycle," *in The Second Visible Human Project Conference Proceedings*, Maryland, USA, 1998.

[Schr96] T. Schreiber and D. T. Kaplan, "Signal separation by nonlinear projections: the fetal electrocardiogram," *Physical Review E*, vol. 53(5), pp. R4326–R4329, 1996.

[Silv12] N. Silver, *The Signal and the Noise: Why So Many Predictions Fail - But Some Don't*, Penguin, 2012.

[Skem58] J. T. Skemp and F. J. Millen, "Electroencephalographic tracings of the fetal heart in utero," *Obstetrics and Gynecology*, vol. 11(2), pp. 148–155, 1958.

[Stog04] H. Stögbauer, A. Kraskov, S. A. Astakhov, and P. Grassberger, "Least Dependent Component Analysis Based on Mutual Information," *Physical Review E*, vol. 70, pp. 066 123–1–066 123–17, 2004.

[Ston04] J. V. Stone, *Independent Component Analysis: A Tutorial Introduction*, MIT Press, Cambridge, Massachusetts, 2004.

[Subh14] S. Subhashini, D. J. Jagannath, and A. Immanuel Selvakumar, "Extricating non invasive fetal ECG by adaptive optimization technique," *in Proceedings of the 2014 International Conference on Electronics and Communication Systems*, pp. 1–5, IEEE, Coimbatore, India, 2014.

[Sugu14] D. Sugumar, P. T. Vanathi, and S. Mohan, "Joint blind source separation algorithms in the separation of non-invasive maternal and fetal ECG," *in Proceedings of the 2014 International Conference on Electronics and Communication Systems*, pp. 1–6, IEEE, Coimbatore, India, 2014.

[Tara11] D. Taralunga, M. Ungureanu, R. Strungaru, and W. Wolf, "Performace comparison of four ICA algorithms applied for fECG extraction from transabdominal recordings," *in Proceedings of the 10th International Symposium on Signals, Circuits and Systems*, pp. 1–4, IEEE, 2011.

[Vale14] G. Valenza, N. Vanello, M. Milanesi, and E. Pasquale, "Decoding underlying brain activities in time and frequency domains through complex indpendent component analysis of EEG signals," *in Proceedings of the 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 3192–3195, IEEE, Chicago, IL, 2014.

[Vand87] J. Vanderschoot, D. Callaerts, W. Sansen, J. Vandewalle, G. Vantrappen, and J. Janssens, "Two methods for optimal MECG elimination and FECG detection from skin electrode signals," *IEEE Transactions on Biomedical Engineering*, vol. 34(3), pp. 233–243, 1987.

[Venn14] M. Vennila and S. M. Yacin, "A modified approach on fetal ECG extraction using extended state Kalman filtering," *in Proceedings of the 2014 International Conference on Green Computing Communication and Electrical Engineering*, pp. 1–5, IEEE, Coimbatore, India, 2014.

[Viga00] R. Vigário, J. Särelä, V. Jousmäki, M. Hämäläinen, and E. Oja, "Independent component approach to the analysis of EEG and MEG recordings," *IEEE Transactions on Biomedical Engineering*, vol. 47(5), pp. 589–593, 2000.

[Weie95] K. Weierstrass, "On continuous functions of a real argument that do not have a well-defined differential quotient," *Mathematische Werke*, vol. 2, pp. 71–74, 1895.

[Widr75] B. Widrow, J. R. Glover Jr., J. M. McCool, J. Kaunitz, C. S. Williams, R. H. Hearn, J. R. Zeidler, E. Dong Jr., and R. C. Goodlin, "Adaptive noise cancelling: principles and applications," *Proceedings of the IEEE*, vol. 63(12), pp. 1692–1716, 1975.

[Wiki06] Wikimedia Commons, "Heart," URL `http://commons.wikimedia.org/wiki/File:Heart.svg`.

[Wiki06a] Wikimedia Commons, "Sinus Rhythm Labels," URL `http://en.wikipedia.org/wiki/File:SinusRhythmLabels.png`.

[Wiki07] Wikimedia Commons and M. Otto, "Systole QRS complex," URL `http://commons.wikimedia.org/wiki/File:Systole_QRS_Complex.png`.

[Wils08] J. D. Wilson, R. B. Govindan, J. O. Hatton, C. L. Lowery, and H. Preissl, "Integrated approach for fetal QRS detection," *IEEE Transactions on Biomedical Engineering*, vol. 55(9), pp. 2190–2197, 2008.

[Wool90] M. S. Woolfson, W. Peasgood, S. Sahota, and J. A. Crowe, "Signal processing of the fetal electrocardiogram," *Journal of Electrocardiography*, pp. 51–57, 1990.

[Xu12] B. Xu, H. Jin, X. Tan, Y. Hu, and X. Luo, "Extracting fetal electrocardiogram based on a modified fast independent component analysis," *in Proceedings of the 9th International Conference on Fuzzy Systems and Knowledge Discovery*, pp. 1787–1791, IEEE, Sichuan, 2012.

[Ye07] Y. Ye, X. Yao, Z.-L. Zhang, and Q. Mo, "A non-invasive fetal electrocardiogram extraction algorithm based on ICA neural network," *in Proceedings of the 1st International Conference on Bioinformatics and Biomedical Engineering*, pp. 806–809, IEEE, Wuhan, 2007.

[Zgal13] W. A. Zgallai, *Second- and Third-Order Statistical Characterization of Non-Linearity and Non-Gaussianity of Adult and Fetal ECG Signals and Noise*, chap. 2, pp. 25–54, InTech, 2013.

There are 101 references in this thesis that range in publication date from 1895-2014.

Appendix A

Heart Rate Calculation

There were three main steps to calculate the heart rate of both the recorded and simulated ECG signals:

(i) First, the ECG signal of interest was created in Matlab. This signal could be obtained from PhysioNet using the PhysioToolKit `rdsamp` command to specify the length of the imported signal. An ECG could also be simulated using the ECGsyn model by setting its sampling frequency, length, and average R-wave frequency.

(ii) Next, the R-wave peaks of the ECG signal were identified. To do this, a minimum distance between peaks and a minimum peak height needed to be manually specified. The code used to identify the peaks is given in the disc attached to this thesis. An example of a recorded ECG with the peaks identified by circles is shown in Figure A.1.

(iii) The time between R-wave peaks was calculated as $R_i - R_{i-1}$ where R_i is the i^{th} R-wave peak in the signal. The length of time between R-waves defines the HRV, as shown in Figure A.2.

The average heart rate in beats per minute, bpm_{mean}, can then be calculated as:

$$bpm_{mean} = \frac{60 \times F_s}{RR_{mean}} \tag{A.1}$$

where F_s is the sampling frequency of the ECG, and RR_{mean} is the average length of the RR-intervals in the signal.

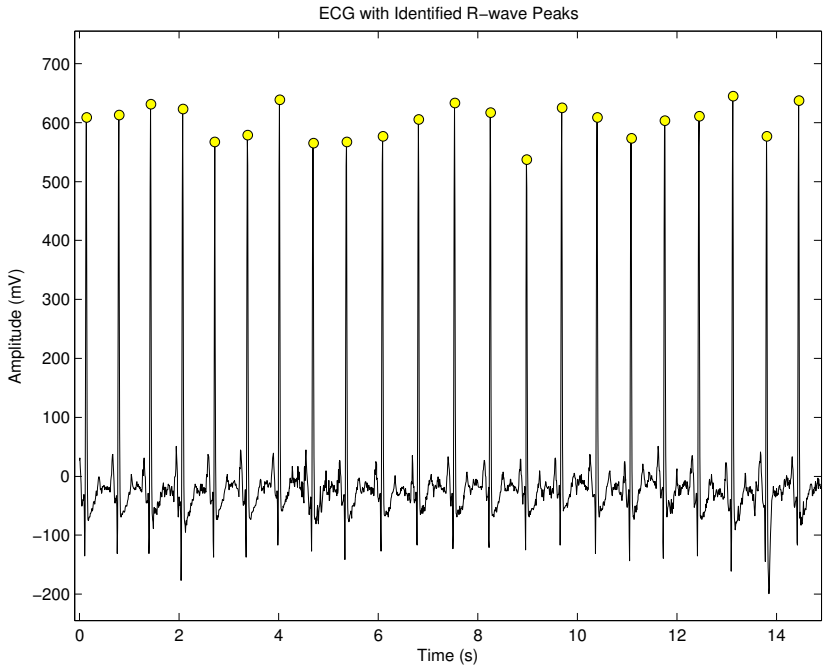

Fig. A.1: Recorded ECG with the R-wave peaks identified.

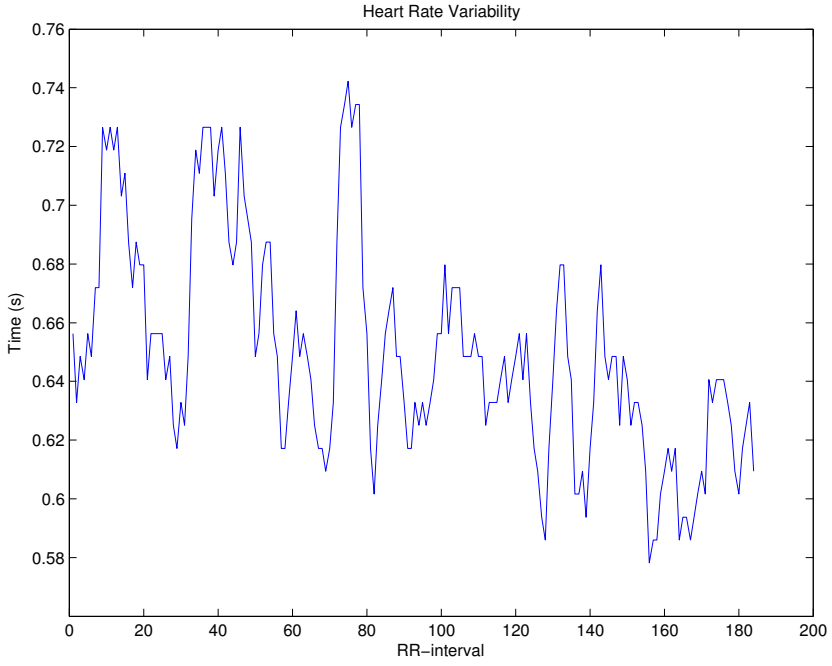

Fig. A.2: Heart rate variability of a recorded ECG.

Appendix B

Experiment Code

The Matlab code that was written for the experiments in Chapters 2 and 5 is provided on the attached disc. This includes the source code:

- to implement the Weierstrass function,
- to visually compare two signals on the same graph,
- to calculate the RMSE between two signals,
- to calculate signal entropy,
- to calculate the estimated probabilities of a signal,
- to calculate the spectral fractal dimension of a time series,
- to import PhysioNet recordings and implement the ECGsyn model, and
- to find the R-wave peaks in an ECG signal and calculate its heart rate and HRV.

The FastICA algorithm is available online as an open source Matlab file [Hyva13]. The ECGwaveGen and ECGsyn heart waveform simulators, as well as the nsrdb and nifecgdb databases of recorded ECG signals, are available on the open source PhysioNet website [Gold00].

Appendix C

Colophon

This thesis is typeset in LaTeXusing a modified version of the custom template created by Dario Schor in TeXShop version 3.11. BibDesk version 1.6.3 was used to manage the references using BibTeX. The body of the report is written in 11 point Times New Roman, while the figure captions are printed in 10 point Arial.

With the exception of Figures 3.1-3.3 and the microphones of the Visual Abstract, which are modified open source images, all figures were created in Matlab version 2010a and saved as Encapsulated PostScript (eps) files at 300dpi.

All of the work was performed using Mac OS X version 10.8.5 with a 2.3 GHz Intel Core i7 processor and 8 GB of 1600 MHz DDR3 memory.