1 Overview of Independent Component Analysis

Every problem becomes very childish when once it is explained to you.
— Sherlock Holmes (The Dancing Men, A.C. Doyle, 1905)

1.1 Introduction

It is often said that we suffer from "information overload," whereas we actually suffer from "data overload." The problem is that we have access to large amounts of data containing relatively small amounts of useful information. This is true both in our daily lives, and within many scientific disciplines. Independent component analysis (ICA) is essentially a method for extracting useful information from data.

ICA is of interest to a wide variety of scientists and engineers because it promises to reveal the driving forces which underlie a set of observed phenomena. These phenomena include the firing of a set of neurons, mobile phone signals, brain images (e.g., functional magnetic resonance imaging, fMRI), stock prices, and voices (see figure 1.1). In each case, a large set of signals are measured, and it is known that each measured signal depends on several distinct underlying factors, which provide the driving forces behind the changes in the measured signals. In other words, each measured signal is essentially a *mixture* of these underlying factors.

For example, the 100 stock prices in the London FTSE index represent a set of 100 time-varying measurements, each of which depends on a relatively small number of distinct time-varying causal factors (e.g., the latest retail sales figures, unemployment rates, and weather conditions). Thus each stock price can be viewed as a different mixture of these factors. If these factors could be extracted from the 100 measured signals then they could (in principle) be used to predict future movements of those 100 stock prices.

Similarly, if the time-varying outputs of 100 neurons in the visual cortex of the brain were measured then ICA could be used to test the extent to which all 100 neurons depend on a small set of causal factors, corresponding (for example) to luminance and the orientation of contrast edges. Having identified these factors, it would then be possible to estimate the extent to which each individual neuron depended on each factor, so that neurons could be classified as coding for luminance or edge orientation.

In every case, it is these factors or *source signals* that are of primary interest, but they are buried within a large set of measured signals, or *signal mixtures*. ICA can be used to extract the source signals underlying a set of measured signal mixtures.

1.2 Independent Component Analysis: What Is It?

ICA belongs to a class of *blind source separation* (BSS) methods for separating data into underlying informational components, where such data can take the form of images, sounds, telecommunication channels or stock market prices. The term "blind" is intended

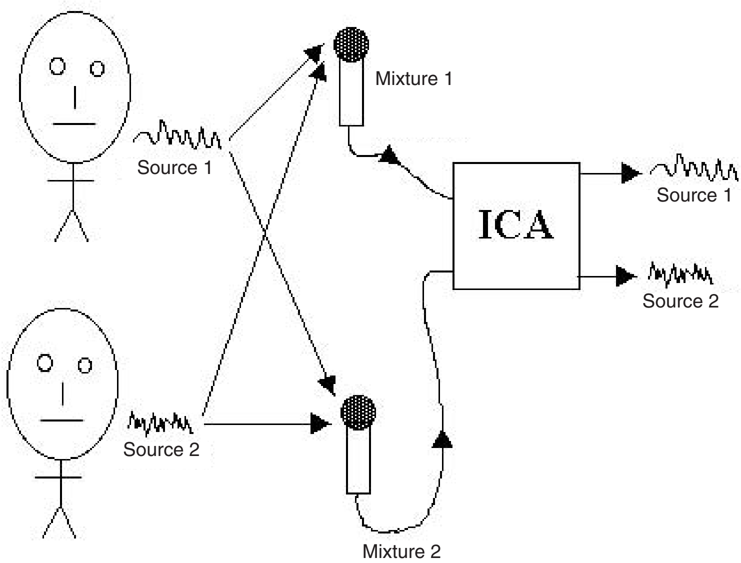

Figure 1.1
ICA in a nutshell. If two people speak at the same time in a room containing two microphones then the output of each microphone is a *mixture* of two voice signals. Given these two *signal mixtures*, ICA can recover the two original voices or *source signals*. This example uses speech, but ICA can extract source signals from any set of two or more measured signal mixtures, where each signal mixture is assumed to consist of a mixture of source signals (see text).

to imply that such methods can separate data into source signals even if very little is known about the nature of those source signals.

As an example, imagine there are two people speaking at the same time in a room containing two microphones, as depicted in figure 1.1. If each voice signal is examined at a fine time scale then it becomes apparent that the amplitude of one voice at any given point in time is unrelated to the amplitude of the other voice at that time (see figure 1.2). The reason that the amplitudes of the two voices are unrelated is that they are generated by two unrelated physical processes (i.e., by two different people). If we know that the voices are unrelated then one key strategy for separating voice mixtures into their constituent voice components is to look for unrelated time-varying signals within these mixtures. Using this strategy, the extracted signals are unrelated, just as the voices are unrelated, and it follows that the extracted signals are the voices. So, simply knowing that each voice is unrelated to the others suggests a strategy for separating individual voices from mixtures of voices. This apparently mundane observation is a necessary prerequisite for understanding how

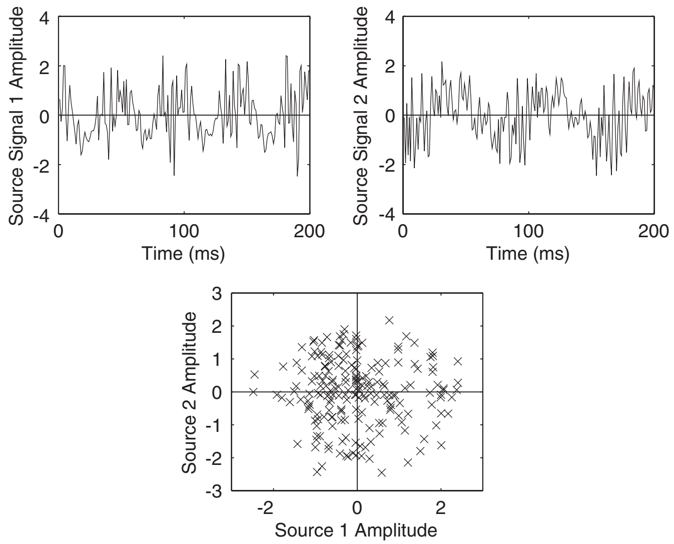

Figure 1.2
ICA exploits the fact that two signals, such as voices, from different physical sources are independent. This implies that, if the two different source voice signals shown in the top panels are examined at a fine time scale then the amplitude of one voice (top left) at any given time provides no information regarding the amplitude of the other voice (top right) at that time. This can be confirmed graphically by plotting the amplitude of one voice at each time point against the corresponding amplitude of the other voice (bottom panel). The resultant distribution of points does not indicate any obvious pattern, suggesting that the two voice signals are independent.

ICA works. The property of being unrelated is of fundamental importance, because it can be used to separate not only mixtures of sounds, but mixtures of almost any type (e.g., images as in figure 1.4, radio).[1]

While it is true that two voice signals are unrelated, this informal notion can be captured in terms of *statistical independence*.[2] If two or more signals are statistically independent of each other then the value of one signal provides no information regarding the value of the other signals.

1. In fact, sounds present a harder separation problem than electromagnetic signals, such as radio. This is because sound travels sufficiently slowly that it arrives at different sensors (microphones) at different times. This differential delay can be overcome in practice (e.g., Lee *et al*., 1997). For simplicity, we will assume there is no such delay in the speech examples considered here.

2. For brevity, we will usually use the term *independence*.

Before considering how ICA works, we need to introduce some terminology. As its name suggests, independent component analysis separates a set of *signal mixtures* into a corresponding set of statistically independent component signals or *source signals*. The mixtures can be sounds, electrical signals, e.g., electroencephalographic (EEG) signals, or images (e.g., faces, fMRI data). The defining feature of the extracted signals is that each extracted signal is *statistically independent* of all the other extracted signals.

1.3 How Independent Component Analysis Works

ICA is based on the simple, generic and physically realistic assumption that if different signals are from different physical processes (e.g., different people speaking) then those signals are statistically independent. ICA takes advantage of the fact that the implication of this assumption can be reversed, leading to a new assumption which is logically unwarranted but which works in practice, namely: if statistically independent signals can be extracted from signal mixtures then these extracted signals must be from different physical processes (e.g., different people speaking). Accordingly, ICA separates signal mixtures into statistically independent signals. If the assumption of statistical independence is valid then each of the signals extracted by independent component analysis will have been generated by a different physical process, and will therefore be a desired signal.

The preceding description represents the high-level strategy implicit in ICA. The mathematical nuts and bolts of precisely how ICA works are described in subsequent chapters. While the nuts and bolts are necessary, grasping the essential physically motivated underpinnings of independent component analysis is the key to understanding these nuts and bolts.

1.4 Independent Component Analysis and Perception

The problem of blind source separation solved by independent component analysis is analogous to the problem encountered by every newborn animal: how to decompose perceptual inputs into their underlying physical causes. For example, if an animal looks at an object then each retinal receptor has an output which is a function of several physical causes, including the luminance, reflectance, slant, tilt, and motion of the object's surface. ICA methods have the potential to provide a rigorous model of how the decomposition of perceptual inputs can be learned by making use of generic and physically plausible constraints, such as statistical independence and spatiotemporal continuity (see chapter 11 for an example using stereo disparity). This is not intended to suggest that the brain implements ICA, but simply that ICA and neuronal computation are based on a common set of underlying principles (see Barlow, 1981) for a classic and lucid account of this type of approach with respect to uncorrelatedness rather than independence).

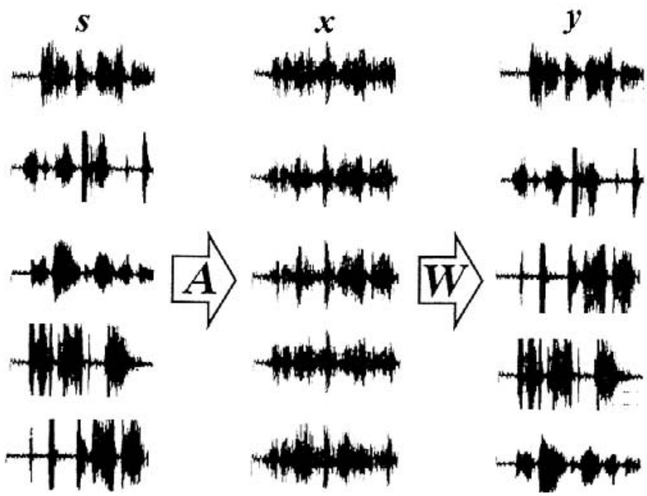

Figure 1.3
Speech Separation.
Left: Each of five people speaking simultaneously generates an independent voice *source signal*. The set of source signals is denoted **s**.
Middle: If there are five microphones present then the output of each microphone is a mixture of five independent source signals (i.e., voices). The (unknown) speaker-microphone distances are represented by the mixing process labeled **A**. The set of signal mixtures is denoted **x**.
Right: ICA extracts five independent components from the set of signal mixtures, where each extracted signal is an estimate of one of the original source signals (i.e. single voices). The unmixing process identified by ICA is denoted **W**, and the estimated source signals are denoted **y**. Note that ICA re-orders signals, so that an extracted signal y_i and its source signal s_i are not necessarily on the same row. From (Bell & Sejnowski, 1995).

1.5 Principal Component Analysis and Factor Analysis

ICA is related to conventional methods for analyzing large data sets, such as *principal component analysis* (PCA) and *factor analysis* (FA) (see chapter 10 and appendix F). Whereas ICA finds a set of independent source signals, PCA and FA find a set of signals with a much weaker property than independence. Specifically, PCA and FA find a set of signals which are *uncorrelated* with each other. This is a crucial distinction, to which we will return later. For example, PCA would extract a set of uncorrelated signals from a set of mixtures. If these mixtures were microphone outputs then the extracted signals would simply be a new set of voice mixtures. In contrast, ICA would extract a set of independent signals from this set of mixtures, so that the extracted signals would be a set of single voices.

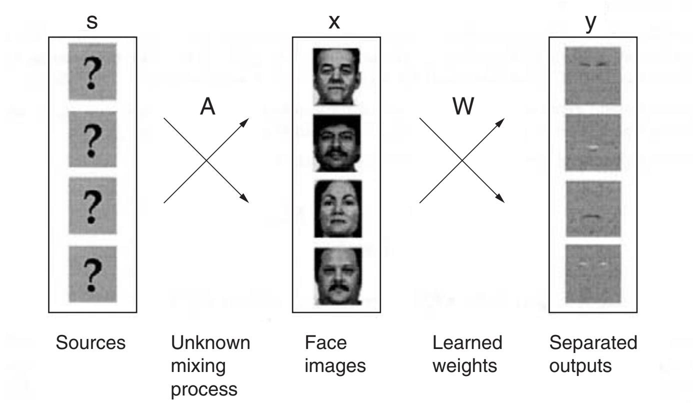

Figure 1.4
Face Recognition.
Left: Each of four prototypical unknown faces is a (spatial) independent source signal. The set of four source signals is labeled **s**.
Middle: The faces of four people are each assumed to be a different mixture of underlying prototypical faces (left), where the mixing process is labeled **A**. The set of four signal mixtures is labeled **x**.
Right: ICA extracts a set of signals, each of which is an estimate of one of the unknown spatial source signals. This unmixing process is labeled **W**, and the set of four estimated source signals is labeled **y**. Note how the estimated source signals contain spatially localized features corresponding to perceptually salient features, such as mouth and eyes. From (Bartlett, 2001).

The "forward" assumption that signals from different physical processes are uncorrelated still holds, but the "reverse" assumption that uncorrelated signals are from different physical processes does not. This is because lack of correlation is a weaker property than independence. In summary, independence implies a lack of correlation, but a lack of correlation does not imply independence.

1.6 Independent Component Analysis: What Is It Good For?

ICA has been applied to problems in fields as diverse as speech processing, brain imaging (e.g., fMRI and optical imaging), electrical brain signals (e.g., EEG signals), telecommunications, and stock market prediction. However, because independent component analysis is an evolving method which is being actively researched around the world, the limits of what ICA may be good for have yet to be fully explored.

Two contemporary applications of ICA are presented in figures 1.3– 1.4. Note that ICA can be used to find independent components which can take the form of speech, electrical signals or images.

In conclusion, ICA is based on a single physically realistic assumption: namely, that different physical processes generate outputs that are independent of each other. In the following chapters it will be shown how this assumption not only provides an intuitive insight into how ICA works but also how it provides insight into how the physical world works.