

INFORMED SOURCE SEPARATION: A BAYESIAN TUTORIAL

Kevin H. Knuth

Intelligent Systems Division
NASA Ames Research Center
Moffett Field CA 94035 USA
phone: +1-650-604-4279, fax: +1-650-604-4036,
email: kevin.h.knuth@nasa.gov
web: <http://ic.arc.nasa.gov/people/knuth/>

ABSTRACT

Source separation problems are ubiquitous in the physical sciences; any situation where signals are superimposed calls for source separation to estimate the original signals. In this tutorial I will discuss the Bayesian approach to the source separation problem. This approach has a specific advantage in that it requires the designer to explicitly describe the signal model in addition to any other information or assumptions that go into the problem description. This leads naturally to the idea of *informed source separation*, where the algorithm design incorporates relevant information about the specific problem. This approach promises to enable researchers to design their own high-quality algorithms that are specifically tailored to the problem at hand.

1. UNDERSTANDING THE PROBLEM

To gather information about the physical world, we deploy sensors to make measurements and detect signals. Our sensors, if properly designed, will collect information about the signals of interest. However, very often the signals of interest are comprised of a set of discrete signals, which have been superimposed during propagation, often with signals that are not of interest. Thus our sensors almost invariably detect a mixture of signals—some interesting and some non-interesting. In more straightforward applications, careful design of the sensors and application of filters can limit the recordings to the signal of interest. However, when this is not possible, more extreme steps need to be taken. This leads to a class of problems called source separation problems.

There is no limit to the complications that may arise. Superposition may be linear, or one of the infinite varieties of nonlinear superposition. If a set of sensors are used, there may be time delays due to propagation of the signals from each source to each detector, or there could be convolutions due to reflections or differences in propagation speed through intervening media coupled with diffraction. To presume to be able to construct a single algorithm that can deal with all of these imaginable cases is unrealistic. Instead, we must focus our efforts on developing a methodology for designing robust algorithms that are specific to the application at hand. Only then can we take advantage of the specific prior knowledge we possess about each problem to increase our chances of reaching an accurate and optimal solution. I call this approach *informed source separation*, which should be

contrasted with blind source separation, where very little is assumed to be known about the problem.

The source separation problem can be viewed as an inference problem, where one models a set of detected signals as a mixture of a set of source signals. It is important to remember that inference is not deduction—it doesn't always work. In difficult problems, prior information goes a long way to help assure that we reach an accurate solution. This prior information can take many forms, and can come into the problem at several different points. I will show that this prior information can significantly transform the source separation problem, and subsequently, the algorithmic solution.

2. BAYESIAN PROBABILITY THEORY

In this section, I give a brief description of the Bayesian methodology. I will focus on the use of probability theory to describe our knowledge, and leave the details of the problem of searching the hypothesis space for the optimal solution to other authors. The crux of the methodology is Bayes' Theorem

$$p(\text{model}|\text{data},I) = p(\text{model}|I) \frac{p(\text{data}|\text{model},I)}{p(\text{data}|I)} \quad (1)$$

where I represents our prior information. The probability on the left $p(\text{model}|\text{data},I)$ is called the *posterior probability*. It is the probability that a specific *model* accurately describes the problem given the *data* and our prior information I . The first term on the right $p(\text{model}|I)$ is called the *prior probability*, or *prior* for short. The prior describes the degree to which we believe a specific model is the correct description before we see any data, and thus encodes our knowledge about the possible values of the model parameters. The term in the numerator $p(\text{data}|\text{model},I)$ is called the *likelihood*, which describes the degree to which we believe that the model could have produced the observed data. This term encodes both the process of making predictions with our hypothesized model and the process of comparing these predictions to our data, which is an important part of the scientific method. The term in the denominator $p(\text{data}|I)$ is called the *evidence*. In many parameter estimation problems, where we have a static model and are merely estimating the values of its parameters, this term simply acts a normalization factor. In problems where we are testing one of a set of several models, this term becomes extremely relevant as it can indicate the degree to which a model is favored by the data.

The space of all considered models is called the hypothesis space. Bayes' Theorem turns the source separation prob-

lem into a search problem, where we search the hypothesis space for the most probable model. We can also look at Bayes' Theorem as a learning rule since it tells us how to update our prior knowledge when we receive new data. What we have learned from this data combined with what we knew prior is described by the posterior probability. In the next section I will show how the Bayesian method applied to a well-defined set of prior information leads directly to Infomax ICA.

3. FIRST EXAMPLE: INFOMAX ICA

In this section I will demonstrate how the Bayesian methodology allows one to derive Infomax ICA [1]. While this particular derivation has been published previously [2], I present it here again in detail both to assist with understanding the later derivations in this paper, and also to clear up some common misconceptions surrounding ICA and source separation. ICA is commonly considered to be a blind source separation algorithm because we make a minimal number of assumptions. However, it is important to note that no algorithm is truly blind, and that the assumptions we make—even if minimal in some sense—will have an effect on the performance of an algorithm when applied to a given problem.

We begin by assuming that there are N sources whose signals propagate instantaneously to N distinct detectors. The signals are assumed to superimpose linearly so that each detector records a linear mixture of the source signals. Furthermore, the entire process is assumed to be noise-free. This leads to a simple mathematical model that describes the recorded signals in terms of the unknown sources¹

$$x_{it} = \sum_{j=1}^N A_{ij}s_{jt} \quad (2)$$

where x_{it} is the signal recorded at i^{th} detector at time t , s_{jt} is the source signal emitted by the j^{th} source at time t , and A_{ij} is the "mixing matrix". The elements of the mixing matrix serve to couple the sources to the detectors. Physically each A_{ij} describes how the signal propagates from the source to the detector. While the physical interpretation is kept vague in the blind algorithm, we will see that the physical interpretation is quite useful when deriving informed source separation algorithms.

We can now apply Bayes' Theorem (1) to express the probability of our model, \mathbf{A} and \mathbf{s} , given our data, \mathbf{x}

$$p(\mathbf{A}, \mathbf{s} | \mathbf{x}, I) = p(\mathbf{A}, \mathbf{s} | I) \frac{p(\mathbf{x} | \mathbf{A}, \mathbf{s}, I)}{p(\mathbf{x} | I)}, \quad (3)$$

where \mathbf{A} represents the entire matrix, \mathbf{s} represents all of the source signals emitted by the sources, and \mathbf{x} represents all of our recorded data. Given the fact that the source signals are independent of the propagation, we can factor the prior probability $p(\mathbf{A}, \mathbf{s} | I)$ into two terms $p(\mathbf{A} | I)$ and $p(\mathbf{s} | I)$.

$$p(\mathbf{A}, \mathbf{s} | \mathbf{x}, I) = p(\mathbf{A} | I) p(\mathbf{s} | I) \frac{p(\mathbf{x} | \mathbf{A}, \mathbf{s}, I)}{p(\mathbf{x} | I)}. \quad (4)$$

Once we assign these probabilities, the problem is in some sense solved. All we will need to do is to search all possible values of the matrix \mathbf{A} and the source waveshapes \mathbf{s} to

¹I have purposefully kept this in component form so that it may be more easily compared with other algorithms presented later in this tutorial.

find the case that is most probable.² When we perform this search, the evidence in the denominator will never contribute to the calculation since it doesn't depend on the model parameters. So we can simplify the problem further by writing (4) as a proportionality

$$p(\mathbf{A}, \mathbf{s} | \mathbf{x}, I) \propto p(\mathbf{A} | I) p(\mathbf{s} | I) p(\mathbf{x} | \mathbf{A}, \mathbf{s}, I). \quad (5)$$

Now with the basic model in hand, we can construct a likelihood function. In this case, we are assuming that the recording process is noise-free; thus we will assign a delta-function likelihood function for each datum point

$$p(x_{it} | \mathbf{A}, \mathbf{s}_t, I) = \delta(x_{it} - \sum_{j=1}^N A_{ij}s_{jt}), \quad (6)$$

where I have used \mathbf{s}_t to represent all of the N source amplitudes emitted at time t . This delta function likelihood states, very strongly, that we believe that our model of the source separation problem (2) is correct. The recordings are independent of one other, so the likelihood function for our entire data set is merely the product of likelihoods (6) for each detector and each time step

$$p(\mathbf{x} | \mathbf{A}, \mathbf{s}, I) = \prod_{i=1}^N \prod_{t=1}^T \delta(x_{it} - \sum_{j=1}^N A_{ij}s_{jt}), \quad (7)$$

where T is the number of time steps.

Next we assume that the probability density of the amplitude of the individual source signals has a positive kurtosis (also known as leptokurtotic or super-Gaussian). With this assumption, we can assign a prior probability for the amplitudes of the signals emitted by the sources. Without such a prior, this problem has an infinite number of perfectly good solutions. This prior information will serve to make the problem soluble in cases where it is correct, while risking the possibility of incorrect solutions in cases where this assumption does not hold. We will write

$$p(s_{jt} | I) = q_j(s_{jt}), \quad (8)$$

where $q_j(s_{jt})$ is the probability that the j^{th} source could have a given amplitude at any time t . We could easily follow Bell & Sejnowski [1] and assign the derivative of a sigmoid function, which is a leptokurtotic density function. However, for the purposes of generalization, we will just write it as q_j . Assuming that the sources are independent of one another, we have

$$p(\mathbf{s} | I) = \prod_{j=1}^N \prod_{t=1}^T q_j(s_{jt}). \quad (9)$$

We now assume that we know nothing about the mixing matrix. We will encode this knowledge by assigning a uniform prior³ for the value of any given matrix element A_{ij} as long as it is within a "reasonable" range

$$p(A_{ij} | I) = \begin{cases} c & \text{if } A_{\min} \geq A_{ij} \geq A_{\max} \\ 0 & \text{if } A_{ij} < A_{\min}, A_{\max} > A_{ij} \end{cases} \quad (10)$$

²In practice, conducting this search is often the most difficult part of the problem.

³The astute reader will recognize that each matrix element acts as a scaling parameter in the problem. For this reason, a more accurate noninformative prior would be the appropriate Jeffrey's prior for the matrix \mathbf{A} .

where $c = (A_{max} - A_{min})^{-1}$. For the entire mixing matrix, we can assign a uniform joint prior

$$p(\mathbf{A}|I) = \prod_{i=1}^N \prod_{j=1}^N p(A_{ij}|I) \quad (11)$$

$$= \begin{cases} C & \text{if } \forall A_{ij}, A_{min} \geq A_{ij} \geq A_{max} \\ 0 & \text{if } \exists A_{ij}, \text{ s.t. } A_{ij} < A_{min}, A_{max} > A_{ij} \end{cases}$$

where $C = cN^2$. With our likelihood and priors all defined, we are now ready to re-write the posterior probability (5) and begin searching for the most probable parameter values.

However, this search will not be easy. Much of the effort in Bayesian inference is to limit the number of parameters to search over, or to come up with a clever heuristic to perform the search. Furthermore, it is often easier to work with the logarithm of the posterior (4). This neatly separates the posterior into a sum of the log priors plus the log likelihood.

We will begin by reducing the number of parameters of interest, and will conclude by taking the logarithm. First we reason that if we knew the mixing matrix \mathbf{A} , or better yet, its inverse \mathbf{A}^{-1} , we could apply it to the data to recover an estimate of the source signals. Surely this is not ideal, but it is easier than searching the entire multidimensional parameter space, which would be $N^2 + NT$ parameters. To do this we use the fact that probabilities sum to one, and *marginalize* over all possible values of the source signal amplitudes, written symbolically as

$$p(\mathbf{A}|\mathbf{x}, I) = \int ds p(\mathbf{A}, \mathbf{s}|\mathbf{x}, I) \quad (12)$$

$$\propto \int ds p(\mathbf{A}|I) p(\mathbf{s}|I) p(\mathbf{x}|\mathbf{A}, \mathbf{s}, I)$$

$$\propto p(\mathbf{A}|I) \int ds p(\mathbf{s}|I) p(\mathbf{x}|\mathbf{A}, \mathbf{s}, I)$$

where the integral sign represents all NT integrals over each of the s_{jt} . Substituting our likelihood (7), and source density prior (9), we have NT integrals to solve

$$\int ds_{11} \cdots \int ds_{NT} \prod_{j=1}^N \prod_{t=1}^T q_j(s_{jt}) \delta(x_{it} - \sum_{j=1}^N A_{ij} s_{jt}). \quad (13)$$

The delta functions easily allow us to solve each of the integrals simply by introducing a change of variables where $w_{it} = x_{it} - \sum_{j=1}^N A_{ij} s_{jt}$. We then have that $\det \mathbf{A} ds = d\mathbf{w}$, and that $s_{jt} = \sum_{i=1}^N A_{ij}^{-1} (x_{it} - w_{it})$, so that the integral becomes

$$\frac{1}{\det A} \int \cdots \int dw_{NT} \prod_{j=1}^N \prod_{t=1}^T q_j \left(\sum_{i=1}^N A_{ij}^{-1} (x_{it} - w_{it}) \right) \delta(w_{it}). \quad (14)$$

The delta functions now all select $w_{it} = 0$, and we have as a result

$$\frac{1}{\det A} \prod_{j=1}^N \prod_{t=1}^T q_j \left(\sum_{i=1}^N A_{ij}^{-1} x_{it} \right). \quad (15)$$

Substituting this result into (12) we get

$$p(\mathbf{A}|\mathbf{x}, I) \propto p(\mathbf{A}|I) \frac{1}{\det A} \prod_{j=1}^N \prod_{t=1}^T q_j \left(\sum_{i=1}^N A_{ij}^{-1} x_{it} \right), \quad (16)$$

which when taking the logarithm gives

$$\log p(\mathbf{A}|\mathbf{x}, I) = K + \log p(\mathbf{A}|I) - \log \det A + \sum_{j=1}^N \sum_{t=1}^T \log q_j \left(\sum_{i=1}^N A_{ij}^{-1} x_{it} \right), \quad (17)$$

where K is the logarithm of the constant implicit in the proportionality. By varying \mathbf{A} to maximize the log posterior above, we can solve for the optimal mixing matrix. The way this is done in ICA is to take the derivative with respect to the inverse of \mathbf{A} and to use this in a gradient ascent learning rule. Specifically, if we assign the mixing matrix prior according to (11), and write $W_{ij} = A_{ij}^{-1}$, we get the familiar Infomax ICA gradient ascent learning rule [1, 2]

$$\Delta W_{ij} = \frac{\partial}{\partial W_{ij}} \left[-\log \det A + \sum_{j=1}^N \sum_{t=1}^T \log q_j \left(\sum_{i=1}^N A_{ij}^{-1} x_{it} \right) \right]$$

$$= A_{ji} + \sum_{j=1}^N \sum_{t=1}^T x_{jt} \left(\frac{q'_j(u_{it})}{q_j(u_{it})} \right)_j, \quad (18)$$

where $u_{it} = \sum_{i=1}^N A_{ij}^{-1} x_{it}$.

However, strictly speaking, this rule doesn't lead to the optimal separation matrix \mathbf{W} , since the maximum value of the posterior with respect to variations of \mathbf{A} will not be identical to the maximum value of the posterior with respect to \mathbf{A}^{-1} . This is due to the fact that probability densities transform differently than functions. Since

$$p(\mathbf{A}^{-1}|\mathbf{x}, I) = p(\mathbf{A}|\mathbf{x}, I) \left| \frac{\partial \mathbf{A}^{-1}}{\partial \mathbf{A}} \right|^{-1}, \quad (19)$$

if we define

$$\hat{\mathbf{A}} = \arg \max_{\mathbf{A}} p(\mathbf{A}|\mathbf{x}, I) \quad (20)$$

$$\check{\mathbf{W}} = \arg \max_{\mathbf{A}^{-1}} p(\mathbf{A}|\mathbf{x}, I) \quad (21)$$

$$\hat{\mathbf{W}} = \arg \max_{\mathbf{A}^{-1}} p(\mathbf{A}^{-1}|\mathbf{x}, I) \quad (22)$$

in general, we have that

$$\check{\mathbf{W}} \neq \hat{\mathbf{W}} \text{ and } \hat{\mathbf{W}} \neq \hat{\mathbf{A}}^{-1}, \quad (23)$$

where $\hat{\mathbf{W}}$ is the optimal separation matrix, $\hat{\mathbf{A}}$ is the optimal mixing matrix, and $\check{\mathbf{W}}$ is the Infomax ICA solution. Thus the inverse of the optimal estimate of the mixing matrix does not equal the optimal estimate of the separation matrix.⁴ However, the ICA solution is actually neither of these. If we were really interested in finding the most probable inverse of the mixing matrix, we should have used (19) to write the posterior for \mathbf{A}^{-1} and solved for its most probable value (22). As a result the standard technique (18) and (21) leads to a biased separation.

That being said, there is much one can learn from this derivation of the Infomax ICA algorithm. Certainly one obtains the same answer when deriving it from the information-theoretic viewpoint; this being due to the duality between

⁴The classic example of this is the fact that the frequency at which a given blackbody spectrum has maximum energy density is different than the wavelength at which it has maximum energy density.

probability theory and information theory [3]. However, the Bayesian derivation has several distinct advantages. First, all of the assumptions that go into the algorithm are made explicit. We see that the sigmoidal nonlinearity in the original derivation is merely related to the derivative of the prior probability for the source amplitude density. This answers one of the common questions that arises: Why does ICA have problems separating pure sinusoids? The answer is clear; sinusoids have bimodal amplitude histograms, which is a severe deviation from our prior expectation of the super-Gaussian prior probability that we have assigned explicitly, and Bell & Sejnowski assigned implicitly [4]. Modifications to this prior to allow for sub-Gaussian densities typically do not improve the situation mainly because they are essentially smoothed uniform densities, which are non-informative. If you want to separate sinusoids, you need to include this relevant information in the design of the algorithm.

Second, why does ICA assume the same number of sources as detectors? In this derivation one can see that the integral is not analytically solvable if we do not make such an assumption. In addition, if we would have assumed that the recorded signals were noisy and assigned a Gaussian likelihood, we would again not be able to perform the integration analytically past the first integral. The noise-free square mixing matrix allows for an analytic marginalization over the source waveshapes resulting in a straightforward and elegant solution. However, this elegant solution will break when pushed too far.

Last, another common question that arises is: Are Gaussian-distributed signals separable? Often the answer is ‘yes’—you just have to rely on additional or different prior information. This is why understanding the information that goes into the design of an algorithm is so important. It allows you to better understand the range of applicability of an algorithm and how to fix it when it doesn’t work. This is why I prefer the Bayesian approach to source separation. It requires you to make all of this explicit.

4. INCORPORATING PRIOR KNOWLEDGE

In this section I will demonstrate another advantage to the Bayesian approach. We will modify the algorithm to account for a simple piece of prior information. Let’s say that we know that the speeds of propagation of the signals remain instantaneous, but that the signals follow an inverse-square propagation law. Such knowledge implies that the coefficients of the mixing matrix are dependent on the relative positions of the sources and detectors. How can we use this information if we have no knowledge of the source-detector distances?

First, if we did know the distance r_{ij} from source j to detector i , the mixing matrix element A_{ij} would follow the inverse-square propagation law

$$A_{ij} = \frac{1}{4\pi r_{ij}^2}. \quad (24)$$

However, we may know that the source must be within some maximum distance R from the detector. If it could be anywhere in the three-dimensional space surrounding the detector, we can assign a uniform probability for its position (r, θ, ϕ) within any volume element of that space

$$p(r, \theta, \phi|I) = \frac{1}{V} = \frac{3}{4\pi R^3}, \quad (25)$$

where V is the spherical volume of radius R surrounding the detector. This is the prior probability that the source is at any position with respect to the detector. However, we only need the probability that the source is some distance r from the detector. We obtain this by marginalizing over all possible values of the angular coordinates

$$p(r|I) = \int_0^{2\pi} d\phi \int_0^\pi \sin\theta d\theta r^2 p(r, \theta, \phi|I) \quad (26)$$

$$= \int_0^{2\pi} d\phi \int_0^\pi \sin\theta d\theta r^2 \frac{3}{4\pi R^3} \quad (27)$$

$$= \frac{3r^2}{R^3}.$$

The prior on the source-detector distance is very reassuring since it is naturally invariant with respect to coordinate rescaling (change of variables). Specifically if we introduce a new coordinate system so that $\rho = ar$ and $P = aR$ with $a > 0$, we find equating the probabilities around $\rho + d\rho$ and $r + dr$ that

$$|p(\rho|I)d\rho| = |p(r|I)dr| \quad (28)$$

$$p(\rho|I) |d\rho| = p(r|I) |dr|$$

$$p(\rho|I) = p(r|I) \left| \frac{d\rho}{dr} \right|^{-1}$$

$$p(\rho|I) = \frac{3r^2}{R^3} \left| \frac{d\rho}{dr} \right|^{-1}$$

$$p(\rho|I) = \frac{3a\rho^2}{P^3} |a|^{-1}$$

$$p(\rho|I) = \frac{3\rho^2}{P^3}.$$

Since this prior is invariant with respect to coordinate rescaling, we can measure distances using any units we wish.

We can now use this to derive a prior for the mixing matrix element. First write the joint probability using the product rule

$$p(A_{ij}, r|I) = p(r|I)p(A_{ij}|r, I). \quad (29)$$

The first term on the right is the source-detector distance prior, and the second term is a delta function described by the hard constraint of the inverse-square law.⁵ These assignments give

$$p(A_{ij}, r|I) = \frac{3r^2}{R^3} \delta(A_{ij} - (4\pi r^2)^{-1}). \quad (30)$$

We now marginalize over all possible values of r

$$p(A_{ij}|I) = \int_0^R dr \frac{3r^2}{R^3} \delta(A_{ij} - (4\pi r^2)^{-1}). \quad (31)$$

⁵Some readers may wonder why I go through the difficulty of using delta functions rather than computing the Jacobians and just performing a change of variables with the probability densities as we did before when demonstrating invariance with respect to rescaling. The reason is that in more complex problems where the parameter of interest depends on multiple other parameters, the change of variables technique becomes extremely difficult. Care must be taken when using delta functions, however, since the argument needs to be written so that it is solved for the parameter of interest, in this case A_{ij} rather than another parameter such as r .

To do this we will need to make a change of variables again by defining $u = (A_{ij} - (4\pi r^2)^{-1})$, so that

$$r^2 = [(4\pi)(A_{ij} - u)]^{-1} \quad (32)$$

$$r^3 = [(4\pi)(A_{ij} - u)]^{-3/2} \quad (33)$$

and $du = (2\pi r^3)^{-1} dr$, which can be rewritten as

$$dr = 2^{-3/2}(2\pi)^{-1/2}(A_{ij} - u)^{-3/2} du \quad (34)$$

giving us

$$p(A_{ij}|I) = 2^{-4}\pi^{-3/2} \frac{3}{R^3} \int_{-\infty}^{u_{max}} du (A_{ij} - u)^{-5/2} \delta(u), \quad (35)$$

where $u_{max} = A_{ij} - (4\pi R^2)^{-1}$. The delta function will select $u = 0$ as long as it is true that $u_{max} > 0$ or equivalently that $r < R$. If this is not true, the integral will be zero. If this hard constraint of the sources being within a distance R of the detectors causes problems, your choice for R was wrong. The result is

$$p(A_{ij}|I) = \frac{3}{16\pi^{3/2}R^3} A_{ij}^{-5/2}. \quad (36)$$

so that the prior for the mixing matrix elements is proportional to $A_{ij}^{-5/2}$. Readers more familiar with statistics may note (and perhaps worry) that this prior is improper since it blows up as A_{ij} goes to infinity. This is not a practical concern as long as the sources are not allowed to get off to infinity. Other readers may note that this prior depends explicitly on the value of the maximum source-detector distance R . More accurate knowledge about the value of R will lead naturally to a more appropriate prior probability, resulting in more accurate source separation results. Once the value of R has been chosen, this prior can be inserted into (17) to generate a source separation algorithm that accounts for this prior knowledge.

As one can imagine, these algorithms can be made arbitrarily more detailed depending on the prior information available. If one has information about the absolute positions of the detectors, such as in a sensor web, and probable locations of the sources, one can derive more accurate prior probabilities for the source-detector distances [5] ⁶. This leads naturally to more accurate prior probabilities for the mixing matrix elements, which in turn lead to better results.

In an ICA-style gradient ascent learning rule (17, 18) these mixing matrix priors act as an additive term biasing the update rule toward the solutions suggested by the prior knowledge. From this perspective, the prior can be viewed as a regularizer. However, one shouldn't get too carried away with this viewpoint since rather than being devised in an *ad hoc* manner, these priors can be carefully designed based on the specific prior information possessed by the algorithm designer. It would have been very difficult to guess the prior we have just derived above.

⁶The author would like to thank Vivek Nigam for pointing out errors in the SPIE98 paper, which will be corrected in a future version available at <http://www.arxiv.org/abs/physics/0205069>

5. SEPARATION AND LOCALIZATION

Now that we have introduced the idea of the relative positions of the sources and detectors, we can take this problem even further and attempt not only to separate the sources, but also to localize them. Here I present results from an earlier paper where we considered the relationship between source separation and source localization [6].

We consider the problem of neural source estimation in electroencephalography (EEG) where we have multiple neural sources in the brain and multiple recording electrodes. Each source j emitting a signal s_j will have some three-dimensional position in the brain \mathbf{p}_j . In addition, these sources are such that they often emit dipolar current fields, so we must also be concerned about their orientation \mathbf{q}_j . The mixing matrix \mathbf{A} again describes the coupling between the sources and the detectors. In the case of electrophysiology, often much is known about this coupling since the electrodynamics of current flow through tissue is well-understood and can be modelled in detail using magnetic resonance imaging (MRI) derived head models. In this problem, the electric currents propagate nearly instantaneously throughout the head and superimpose linearly resulting in signals \mathbf{x} recorded by the detectors. Using Bayes' Theorem, we can write the posterior symbolically as

$$p(\mathbf{p}, \mathbf{q}, \mathbf{A}, \mathbf{s} | \mathbf{x}, I) \propto p(\mathbf{p}, \mathbf{q}, \mathbf{A}, \mathbf{s} | I) p(\mathbf{x} | \mathbf{p}, \mathbf{q}, \mathbf{A}, \mathbf{s}, I), \quad (37)$$

where \mathbf{p} represents the positions of all the sources in the model, and similarly for the other non-subscripted parameters. The model we have chosen is redundant in the sense that the mixing matrix depends on the relative positions and orientations of the sources to the detectors. This allows us to simplify some of the terms above, and factor the prior using the product rule

$$p(\mathbf{p}, \mathbf{q}, \mathbf{A}, \mathbf{s} | \mathbf{x}, I) \propto p(\mathbf{x} | \mathbf{A}, \mathbf{s}, I) \times p(\mathbf{A} | \mathbf{p}, \mathbf{q}, I) p(\mathbf{p} | I) p(\mathbf{q} | \mathbf{p}, I) p(\mathbf{s} | I). \quad (38)$$

These priors show that some model parameters are dependent on others, such as the prior for the mixing matrix. The orientations of the sources depend on their position in the brain since the orientations are determined by the histology of that particular neural source.

Now if we assume that we know nothing about the source positions, nor how they affect the mixing matrix, the prior $p(\mathbf{A} | \mathbf{p}, \mathbf{q}, I)$ reduces to $p(\mathbf{A} | I)$. Marginalizing over all source positions and orientations using uniform priors, we recover the basic source separation problem (4)

$$p(\mathbf{A}, \mathbf{s} | \mathbf{x}, I) \propto p(\mathbf{x} | \mathbf{A}, \mathbf{s}, I) p(\mathbf{A} | I) p(\mathbf{s} | I). \quad (39)$$

With the appropriate probability assignments, we could recover the Infomax ICA algorithm, or perhaps another source separation algorithm that is better suited for the job.

However, let's see what happens if we change our focus and concentrate on the source positions rather than the mixing matrix. We will describe the propagation of the signals from the sources to the detectors with a *forward model* that takes into account the electrodynamics of the physical situation. For simplicity, we will write this symbolically as a function

$$A_{ij} = F(\mathbf{d}_i, \mathbf{p}_j, \mathbf{q}_j), \quad (40)$$

where \mathbf{d}_i is the position of the i^{th} detector, \mathbf{p}_j and \mathbf{q}_j are the position and orientation of the j^{th} source. This function will play a role in our prior probability for A_{ij}

$$p(\mathbf{A}|\mathbf{p}, \mathbf{q}, I) = \prod_{i=1}^M \prod_{j=1}^N \delta(A_{ij} - F(\mathbf{d}_i, \mathbf{p}_j, \mathbf{q}_j)), \quad (41)$$

where we are assuming that there are N sources and M detectors.

Next we will assign a Gaussian likelihood to encode that we have uncertainties about how well the model describes the experimental data. The idea is to assume we know the expected squared deviation σ^2 between the predicted and observed results. Given that we know this expected squared deviation, the principle of maximum entropy [7] says that the Gaussian distribution is the most honest quantification of this knowledge.

$$p(\mathbf{x}|\mathbf{A}, \mathbf{s}, I) = \prod_{i=1}^M \prod_{t=1}^T \exp\left[-\frac{(x_{it} - \sum_{l=1}^N A_{il}s_{lt})^2}{2\sigma_i^2}\right]. \quad (42)$$

It is important to note that this does not imply that we believe the noise is Gaussian distributed, it merely implies that we know something about the expected squared deviation.⁷ If in fact we do not know the actual value of σ , we can always marginalize over it to obtain a more conservative probability density related to Student's t distribution.

We now marginalize the posterior to get rid of the nuisance parameters A_{ij}

$$p(\mathbf{p}, \mathbf{q}, \mathbf{s}|\mathbf{x}, I) \propto p(\mathbf{p}|I)p(\mathbf{q}|\mathbf{p}, I)p(\mathbf{s}|I) \int d\mathbf{A} p(\mathbf{x}|\mathbf{A}, \mathbf{s}, I)p(\mathbf{A}|\mathbf{p}, \mathbf{q}, I). \quad (43)$$

With our probability assignments (42) and (41), and writing $F_{ij} = F(\mathbf{d}_i, \mathbf{p}_j, \mathbf{q}_j)$ the integrals become

$$\int d\mathbf{A}_{ij} \prod_{i=1}^M \prod_{t=1}^T \exp\left[-\frac{(x_{it} - \sum_{l=1}^N A_{il}s_{lt})^2}{2\sigma_i^2}\right] \times \prod_{j=1}^N \delta(A_{ij} - F_{ij}), \quad (44)$$

which gives

$$\prod_{i=1}^M \prod_{t=1}^T \exp\left[-\frac{(x_{it} - \sum_{l=1}^N F_{il}s_{lt})^2}{2\sigma_i^2}\right]. \quad (45)$$

Simplifying the notation further by writing $\hat{x}_{it} = \sum_{l=1}^N F_{il}s_{lt}$,⁸ the marginalized posterior is then

$$p(\mathbf{p}, \mathbf{q}, \mathbf{s}|\mathbf{x}, I) \propto p(\mathbf{p}|I)p(\mathbf{q}|\mathbf{p}, I)p(\mathbf{s}|I) \times \prod_{i=1}^M \prod_{t=1}^T \exp\left[-\frac{(x_{it} - \hat{x}_{it})^2}{2\sigma_i^2}\right]. \quad (46)$$

⁷Jaynes has an excellent chapter where he works through this common misconception [7] ch. 7, specifically 7.7.

⁸Note that $\hat{\mathbf{x}}$ are the predicted recordings based on the sources \mathbf{s} and the forward model.

The remaining priors provide the potential for the introduction of a significant amount of prior information. In this demonstration however, I will simply assign uniform priors. Taking the logarithm of the posterior results in

$$\log p(\mathbf{p}, \mathbf{q}, \mathbf{s}|\mathbf{x}, I) = -\sum_{i=1}^M \sum_{t=1}^T \frac{(x_{it} - \hat{x}_{it})^2}{2\sigma_i^2} + C, \quad (47)$$

where C is the logarithm of the implicit proportionality constant. Maximizing this log posterior results in minimizing the familiar chi-squared 'cost function'

$$\chi^2 = \sum_{i=1}^M \sum_{t=1}^T \frac{(x_{it} - \hat{x}_{it})^2}{2\sigma_i^2} + C, \quad (48)$$

which is a common procedure in electromagnetic source localization.

From this example we see that based on the parameters of interest and the prior information we include, a source separation problem can become a source localization problem. The lesson here is that the Bayesian formalism is the structure that underlies not only source separation and source localization problems, but rather signal processing in general. In fact, many familiar techniques—even the Fourier transform [8]—have their basis in the Bayesian methodology and can be significantly improved by understanding the underlying models and assumptions that go into the algorithm.

6. BEYOND SEPARATION

We can take these ideas further by developing signal models that include parameters that allow us to characterize or describe the signals in different ways. In this example, we again consider EEG signals. Typically an experimenter will design an experiment and record data from multiple experimental trials. The standard analysis technique consists of averaging the data across trials to reduce the effects of noise ('noise' meaning signals that are either not understood or not interesting). The implicit signal model that is employed is the signal plus noise (SPN) model where we assume that we have a single stereotypic source waveshape $s(t)$ that is produced every trial in addition to ongoing noise $\eta(t)$.⁹ The data that is recorded in an electrode can be modelled as

$$x_r(t) = s(t) + \eta_r(t), \quad (49)$$

where r indexes one of the R trials and t indexes the measurements at T discrete time points.¹⁰ Using Bayes' Theorem we have,

$$p(\mathbf{s}|\mathbf{x}, I) \propto p(\mathbf{x}|\mathbf{s}, I)p(\mathbf{s}|I). \quad (50)$$

Relying on arguments laid out in the previous section, I will assign a Gaussian likelihood, and a uniform prior for s . The log posterior is then

$$\log p(\mathbf{s}|\mathbf{x}, I) = -\sum_{r=1}^R \sum_{t=1}^T \frac{(x_r(t) - s(t))^2}{2\sigma^2} + C. \quad (51)$$

⁹I have changed notation here slightly where I am now writing these signals as functions of time. This notation is more clear later when we are required to describe latency shifts of the neural response.

¹⁰Note that the Bayesian methodology does not require that these measurements be equally spaced in time. This is a distinct advantage when dealing with 'missing data' problems.

We can find the maximum of the log posterior by taking the derivative with respect to each $s(t)$ and setting it equal to zero. For a particular time t' we have

$$\begin{aligned} \frac{\partial}{\partial s(t')} \log p(\mathbf{s}|\mathbf{x}, T) &= \frac{\partial}{\partial s(t')} \left(- \sum_{r=1}^R \sum_{t=1}^T \frac{(x_r(t) - s(t))^2}{2\sigma^2} + C \right) \\ &= - \frac{1}{2\sigma^2} \sum_{r=1}^R \frac{\partial}{\partial s(t')} \left(\sum_{t=1}^T (x_r(t) - s(t))^2 \right) \\ &= -\sigma^{-2} \sum_{r=1}^R (x_r(t') - s(t')). \end{aligned} \quad (52)$$

Setting this equal to zero and solving for $s(t')$ we get

$$s(t') = \frac{1}{R} \sum_{r=1}^R x_r(t'), \quad (53)$$

which shows that if you believe that there is only one stereotypic signal in the data then averaging the data over trials will yield the optimal estimator of the source signal.

However, researchers are well aware that there are multiple simultaneous signals, and that these signals vary from trial-to-trial. We have shown that both amplitude and latency variability play a role in the variations of the signals emitted by neural sources [9]. This has led us to a new model of the recorded signal from a set of neural sources

$$x_r(t) = \sum_{n=1}^N \alpha_{nr} s_n(t - \tau_{nr}) + \eta_r(t), \quad (54)$$

where α_{nr} describes the amplitude scale of the n^{th} component during the r^{th} trial, and τ_{nr} similarly describes its latency shift forward or backward in time. This allows us to account for and to characterize amplitude changes and response delays in the neural responses during the course of an experiment or under different experimental conditions. This model assumes that each of the N sources has a distinct stereotypic wave-shape. In our work we have found that by simply describing these additional characteristics of the neural responses, we can separate source signals that vary differentially from trial to trial. The algorithm that results from this model, and our subsequent probability assignments, is called *differentially Variable Component Analysis* (dVCA) [10, 11, 12]. To accommodate multiple detectors, we simply modify the signal model accordingly

$$x_{mr}(t) = \sum_{n=1}^N C_{mn} \alpha_{nr} s_n(t - \tau_{nr}) + \eta_{mr}(t), \quad (55)$$

where \mathbf{C} is the mixing matrix, or coupling matrix as we call it since it describes the coupling between the sources and the detectors. With this new signal model in hand, we are already making interesting new discoveries in our old data sets.

7. CONCLUSION

In this tutorial I have introduced the idea of *informed source separation*. My motivations here are those of a physical scientist, where I have specific problems in need of accurate solutions. In these cases, it is much more advantageous to begin with the appropriate model, introduce the known prior

information, and derive an algorithm specifically engineered for the task.

Historically, while source separation had its beginnings in neural networks and information theory [13, 14, 15, 1], it was recognized early on that these results were related to the maximum likelihood formalism [16, 17, 18]. From this point, one is easily led to the Bayesian methodology [4, 19, 2, 20, 21]. A distinct advantage of the Bayesian approach is that it breaks the problem into three pieces: the signal model, the cost function, and the search algorithm. The researcher begins by choosing an appropriate signal model for the physical problem. Once this model has been chosen, the researcher uses probability theory to derive the posterior probability, which is the cost function to be optimized. With a cost function in hand, a search algorithm is employed to identify the optimal model parameter values. Each of these three pieces can be modified leading to different algorithms that vary in applicability, accuracy and efficiency.

Missing from this short tutorial is a discussion of the numerous techniques and algorithms that can be used to search the parameter space to identify solutions with high probabilities. I will attempt to refer the reader to a variety of useful and important techniques that have been presented in the literature. These methods include: gradient ascent search [1, 22], iterative fixed point algorithms [23, 10], Markov chain Monte Carlo (MCMC) [24, 25], sequential MCMC (also known as particle filters) [26, 27], mean field and ensemble methods [28, 29, 30, 31], variational Bayes [32, 33], as well as Bayesian techniques which utilize sparsity [25]. Last, as an aid to better understanding Bayesian methods, I would recommend the following introductory references [34, 35, 36].

I would also recommend that the reader seek out the other papers presented in this special session to get a taste for the wide array of methods and applications. My hopes are that this tutorial will inspire and enable readers to engineer algorithms for their specific problems.

REFERENCES

- [1] A. J. Bell and T. J. Sejnowski, "An information-maximization approach to blind separation and blind deconvolution," *Neural Comp.*, 7, 1129–1159, 1995.
- [2] K. H. Knuth, "A Bayesian approach to source separation," in J.-F. Cardoso, C. Jutten and P. Loubaton (eds.), *Proc First Int Workshop on ICA and Signal Separation: ICA'99*, Aussois, France, Jan. 1999, pp. 283–288, 1999.
- [3] K. H. Knuth, "Lattice duality: The origin of probability and entropy," In press: *Neurocomputing*, 2005.
- [4] K. H. Knuth "Difficulties applying recent blind source separation techniques to EEG and MEG," in G. J. Erickson, J. T. Rychert and C. R. Smith (eds.), *Maximum Entropy and Bayesian Methods, Boise 1997*, Dordrecht:Kluwer, pp. 209–222, 1998.
- [5] K. H. Knuth "Bayesian source separation and localization," in A. Mohammad-Djafari (ed.), *SPIE'98 Proceedings: Bayesian Inference for Inverse Problems*, San Diego, July 1998, pp. 147–158, 1998. Note that errors are corrected in the online version <http://arxiv.org/abs/physics/0205069>
- [6] K. H. Knuth and H. G. Vaughan, Jr., "The Bayesian origin of blind source separation and electromagnetic source estimation," in W. von der Linden, V. Dose, R.

- Fischer and R. Preuss (eds.), *Maximum Entropy and Bayesian Methods, Munich 1998*, Dordrecht:Kluwer, pp. 217–226, 1999.
- [7] E. T. Jaynes *Probability theory: the logic of science*. Cambridge:Cambridge Univ. Press, 2003.
- [8] G. L. Bretthorst, *Bayesian Spectrum Analysis and Parameter Estimation, Lecture Notes in Statistics*, vol. 48, New York:Springer-Verlag, 1988.
- [9] W. A. Truccolo, M. Ding, K. H. Knuth, R. Nakamura, and S. L. Bressler, “Trial-to-trial variability of cortical evoked responses: implications for the analysis of functional connectivity,” *Clin. Neurophysiol.* 113(2), 206–226, 2002.
- [10] K. H. Knuth, W. A. Truccolo, S. L. Bressler and M. Ding “Separation of multiple evoked responses using differential amplitude and latency variability,” in T.-W. Lee, T.-P. Jung, S. Makeig, T.J. Sejnowski (eds.), *Proc Third Int Workshop on ICA and Blind Signal Separation: ICA 2001*, Dec 2001, San Diego, USA, pp. 463–468, 2001.
- [11] Truccolo, W.A., Knuth, K.H., Shah, A.S., Bressler, S.L., Schroeder, C.E., and Ding, M., “Estimation of single-trial multi-component ERPs: Differentially variable component analysis (dVCA),” *Biol Cybern.* 89, 426–438, 2003
- [12] K. H. Knuth, A. S. Shah, W. A. Truccolo, M. Ding, S. L. Bressler, C. E. Schroeder, “Differentially variable component analysis (dVCA): Identifying multiple evoked components using trial-to-trial variability”, submitted.
- [13] C. Jutten , J. Herault, “Blind separation of sources, Part 1: an adaptive algorithm based on neuromimetic architecture”, *Signal Process*, 24, 1–10, 1991.
- [14] D. T. Pham, P. Garrat, C. Jutten, “Separation of a mixture of independent sources through a maximum likelihood approach”, *Proc. EUSIPCO*, p.771–774, 1992.
- [15] A. Cichocki, R. Unbehauen, E. Rummert, “Robust learning algorithm for blind separation of signals”, *Electron Lett*, 30, 1386–1387, 1994.
- [16] J.-F. Cardoso, “Infomax and maximum likelihood for source separation”, *IEEE Signal Proc Let*, 4, 112–114, 1997.
- [17] D. J. C. MacKay, “Maximum likelihood and covariant algorithms for independent component analysis”, Draft Paper, 1996, available at: <http://wol.ra.phy.cam.ac.uk/mackay/>
- [18] B. A. Pearlmutter, L. C. Parra. Maximum likelihood blind source separation: A context-sensitive generalization of ICA. In *Advances in Neural Information Processing Systems* 9. MIT Press:Cambridge, 613–619, 1997.
- [19] S. J. Roberts, “Independent component analysis: Source assessment & separation, a Bayesian approach.” *IEE Proc: Vision, Image & Sig Proc*, 145(3), 149–154, 1998.
- [20] A. Mohammad-Djafari, “A Bayesian Approach to Source Separation” In: G.J. Erickson, J.T. Rychert and C.R. Smith (eds.), *Maximum Entropy and Bayesian Methods, Boise 1999*, AIP Conf Proc 567, Melville NY:AIP, 221–244, 2001.
- [21] D. B. Rowe, “A Bayesian Approach to Blind Source Separation”. *Journal of Interdisciplinary Mathematics*, 5(1), 49–76, 2002.
- [22] S.-I. Amari, “Natural gradient works efficiently in learning,” *Neural Comp*, 10, 251–276, 1998.
- [23] A. Hyvärinen , E. Oja, “A fast fixed-point algorithm for independent component analysis”, *Neural Computation*, 9, 1483–1492, 1997.
- [24] S. Sénécal and P.-O. Amblard, “MCMC methods for discrete source separation,” In: A. Mohammad-Djafari (ed.), *Bayesian Inference and Maximum Entropy Methods in Science and Engineering, Paris 2000*, AIP Conf Proc 568, Melville NY:AIP, pp. 350–360, 2001.
- [25] C. Févotte, S. J. Godsill. “A Bayesian approach for blind separation of sparse sources”, in press: *IEEE T Speech Audi P*.
- [26] R. M. Everson, S. J. Roberts, “Particle Filters for Non-stationary ICA,” in M. Girolami (ed.), *Advances in Independent Components Analysis*, Springer Verlag, pages 23-41, 2000.
- [27] J.-R. Larocque, J. P. Reilly, W. Ng, “Particle filters for tracking an unknown number of sources,” *IEEE T Signal Proces*, 50, 2926–2937, 2002.
- [28] H. Lappalainen, “Ensemble Learning for Independent Component Analysis” in J.-F. Cardoso, C. Jutten and P. Loubaton (eds.), *Proc First Int. Workshop on ICA and Signal Separation: ICA’99*, Aussois, France, Jan. 1999, pp. 7–12.
- [29] J. Miskin, D. J. C. MacKay, “Ensemble learning for blind image separation and deconvolution,” In: M. Girolami (ed.), *Advances in Independent Component Analysis*, Springer-Verlag, pp. 123–141, 2000.
- [30] P. A. d. F. R. Hojen-Sorensen, L. K. Hansen, O. Winther, “Mean Field Implementation of Bayesian ICA,” in T.-W. Lee, T.-P. Jung, S. Makeig, T.J. Sejnowski (eds.), *Proc Third Int Workshop on ICA and Blind Signal Separation: ICA 2001*, Dec 2001, San Diego, USA, pp. 439–444, 2001.
- [31] H. Valpola and J. Karhunen. An unsupervised ensemble learning method for nonlinear dynamic state-space models. *Neural Comp*, 14, 2647–2692, 2002.
- [32] N. D. Lawrence, C. M. Bishop, “Variational Bayesian independent component analysis” *Univ of Cambridge Tech Report*, 2000.
- [33] R. A. Choudrey, S. J. Roberts, “Flexible Bayesian Independent Component Analysis for Blind Source Separation” in T.-W. Lee, T.-P. Jung, S. Makeig, T.J. Sejnowski (eds.), *Proc Third Int Workshop on ICA and Blind Signal Separation: ICA 2001*, Dec 2001, San Diego, USA, pp. 90–95, 2001.
- [34] D. S. Sivia, *Data Analysis. A Bayesian Tutorial*, Oxford:Clarendon Press, 1996.
- [35] J. Skilling, “Probabilistic data analysis: an introductory guide,” *J. Microscopy* 190, 28–36, 1998.
- [36] V. Dose, “Bayes in five days,” Reprint no. 83, Centre for Interdisciplinary Plasma Science, 2002.