version 20230423

What and *where* in electromagnetic brain imaging

Alain de Cheveigné, Centre National de la Recherche Scientifique and Ecole normale supérieure.

Abstract

To understand the brain, we need to observe the nature and dynamics of the activity (the "what"), and the location or distribution of its sources (the "where"). This paper proposes a new approach in which these two elements are derived in parallel from separate columns of a data-driven analysis matrix. A subset of columns enhances the activity of interest, and its complement defines spatial filters that suppress that activity (null filters). Each null filter is combined with an anatomy-dependent source model to estimate a set of locations within the sources space in which the source might reside (the "zero set" of that filter) and the final estimate is derived from the intersection of zero sets. There is no theoretical limit to the accuracy with which the location of a source can be estimated in this way, but practical limits may arise from noise in the data, imperfect calibration, or an incomplete or inaccurate source model. The method is illustrated with simulated and real data from EEG, MEG and SEEG.

Introduction

In Plato's allegory of the cave, prisoners inferred the world from shadows shifting on a wall. According to Helmholtz's doctrine of unconscious inference, perception proceeds by inference from incomplete sensory evidence (Helmholtz 1967; Alhacen 1030; Hatfield 2002; Imbert 2020; Kawato 1993; Friston 2002), and a similar idea applies also to the analysis of electromagnetic (EM) brain data, because the sparsely-sampled measurements from techniques such as electroencephalography (EEG), magnetoencephalography (MEG), or invasive electrode arrays, do not suffice to characterize the full activity of the brain (Baillet 2001, 2010). At best, we can infer that activity by combining observations with priors, and the question is how best to do so. In Plato's cave, shadows offered only weak clues as to *what* outside objects were, in the absence of color or texture information, but strong clues as to *where* they were. Here, I transpose the what/where distinction, which curiously is also observed between pathways in the brain (Schneider 1969), to the problem of locating brain activity.

In electromagnetic (EM) brain imaging, observations (MEG, EEG, etc.) are combined with source and forward models to infer the spatial distribution of brain sources. Three main approaches are*: dipole fitting, beamforming*, and *minimum current norm estimate* (MNE), and here I suggest a fourth, *zero-set imaging* (ZSI). All approaches require a source model that constrains the nature and spatial characteristics of brain sources, and a forward model that relates them to observations at the sensor or electrode array. Source and forward models are informed by anatomical measurements (e.g. magnetic resonance imaging, MRI), priors (e.g. regions of interest, ROI), knowledge about tissue conductance, and sensor/electrode geometry. Sophisticated software tools are available to create such models and apply them to source localization and imaging (Oostenveld 2011; Tadel 2011; Gramfort et al 2014).

The principle

The approach involves two steps. First, a data-driven linear analysis method is applied to the electromagnetic recordings to produce an analysis matrix. Second, a subset of columns of the analysis matrix is combined with the source and forward model to derive the location of the source. The rationale behind both steps is explained in this section, and illustrated in the next using an easy-to-visualize "toy model". Additional aspects are explored in the Supplementary Information (SI).

EM observations are related to brain activity via a mixing process that is assumed to be instantaneous and linear:

$$\mathbf{X} = \mathbf{SM} \qquad (1)$$

where matrices $\mathbf{X}$ and $\mathbf{S}$ represent the time series of observations and source activity respectively ($\mathbf{S}$ includes both brain and artifact or noise sources). There are I sources and J observation channels, with $I \gg J$. The mixing matrix $\mathbf{M}$ and the brain source time series $s_i(t)$ are unknown. A useful viewpoint is that those time series span a *vector space* of dimension I, containing all their linear combinations, and that the observations $x_j(t)$ span a *subspace* of that space, of dimension J, upon which brain sources are projected. Since $I > J$, multiple sources may map to the same observation, which accounts for the indeterminacy of the inverse problem.

Linear analysis involves applying a matrix to the observations:

$$\mathbf{Y} = \mathbf{XU} \qquad (2)$$

where $\mathbf{Y}$ is a matrix of time series ("components") that hopefully reflect the brain activity in some useful fashion. For example, one column $\mathbf{y}_k$ might approximate the time course of a process or source of interest. In that case, the corresponding column $\mathbf{u}_k$ of the analysis matrix can be seen as a *spatial filter* (set of J coefficients) that enhances that source. In some

situations, a group of N sources are active together and a subset of N columns of $\mathbf{U}$ enhances them collectively.

Several data-driven methods may be used to find such a matrix $\mathbf{U}$. Among them are *principal component analysis* (PCA), *independent component analysis* (ICA), *canonical correlation analysis* (CCA), *correlated components analysis* (CorrCA) (Parra et al 2019), *joint decorrelation* (JD) (de Cheveigné and Parra 2014), or others. Usually, the aim is to observe the dynamics of a target source of interest, using one column of the analysis matrix as a spatial filter to isolate that source. Here, we take a different perspective and focus instead on the *other* columns of the analysis matrix, that produce components in which the source is not present. Those columns consist of spatial filters that cancel the target source (called *null filters* for that source), from which its location is derived in the second step, as explained in the next paragraph.

Suppose that we have a source model that fits the target source for some value $\tilde{\imath}$ of its parameter space, and a forward model that defines the source-to-channel vector $\mathbf{g}_{\tilde{\imath}}$ for the source position $\tilde{\imath}$ .Multiplying this gain vector by a null filter $\mathbf{u}_k$ for that source yields zero by definition: $\mathbf{g}_{\tilde{\imath}} \cdot \mathbf{u}_k = 0$. However, there may be other locations for which that product is zero, collectively these form the *zero set* of this filter. If several null filters are available, each with a zero set, the source must be located within their intersection. This is the core idea of the zero-set method. To find the intersection of the zero sets, it is convenient to add the gain×filter products over k to obtain a cost function:

$$o(\tilde{\imath}) = \sum_k (\mathbf{g}_{\tilde{\imath}} \cdot \mathbf{u}_k)^2 / b_{\tilde{\imath}} \qquad (3)$$

that is zero when the index $\tilde{\imath}$ matches the real source i. The normalization factor $b_{\tilde{\imath}}$, which doesn't affect the zero, is convenient to counter systematic variations of gain over the source space, as explained in SI. If the search space is small enough, the location of the zero may be found by grid search over the source model parameter space. A larger space may be searched with methods such as the Nelder-Mead simplex algorithm.

The location estimate thus is derived from a combination of data observations encapsulated within the null filters, and prior knowledge encapsulated within the source model. Related to the source model, a rendering model may be used to translate the index within source parameter space into a realistic image. Factors that determine realism (e.g. smoothness of the distribution of activity over cortical tissue) are encapsulated within the rendering model. This pipeline is summarized in Fig. 1.

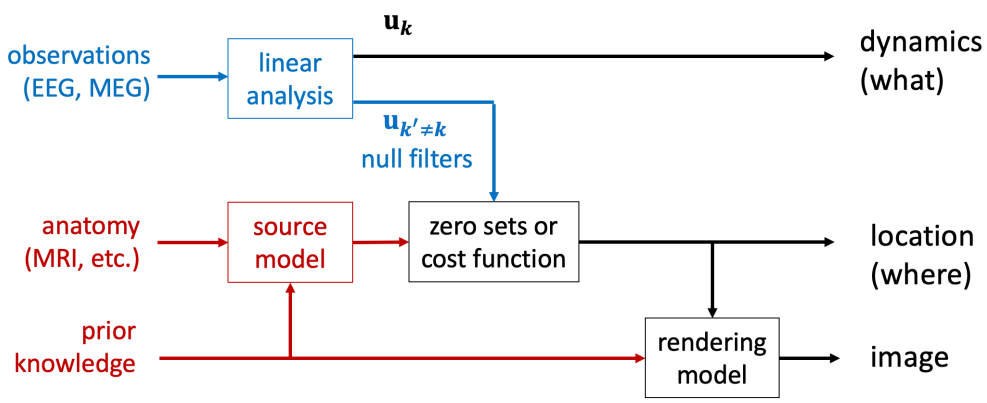

Fig. 1 Data analysis pipeline. Electromagnetic data (EEG, MEG, etc.), in blue, are processed to obtain an analysis matrix $\mathbf{U}$ that is split into one column that isolates the dynamics of the source of interest, and remaining columns that serve as null filters for that source. Anatomical and sensor position data together prior knowledge, in red, inform a source/forward model that, combined with the null filters, produces zero sets (or a cost function) that indicate the location of the source. Optionally, a rendering model produces a realistic image of brain activity.

Concretely, suppose that we wish to localize the source of a stimulus-evoked brain response. For that, we gather EM recordings with an appropriate paradigm (Fig. 1, blue), together with anatomical and co-registration data to inform source and forward models (Fig. 1 red). Based on the recordings, a data-driven linear analysis method produces an analysis matrix $\mathbf{U}$ with one column $\mathbf{u}_k$ that isolates the time course of stimulus-locked activity, and the other columns $\mathbf{u}_{k'\neq k}$ that serve as null filters. Combining the latter with a source and forward model results in zero sets (or a cost function) from which a location in source parameter space is

derived. Optionally, as a final step, a rendering model may be applied to produce a satisfying image of the distribution of neural activity (Fig. 1).

At this point, I must temper the reader's enthusiasm. A basic assumption is that the source model is *complete*, so that any real source can be matched by the model. If not, there are no zero sets, and Eq. 3 is not satisfied for any parameter value. Unfortunately, that seems to be the case empirically for real brain data, at least for standard source/forward models (e.g. a single dipole). We are thus forced to treat the method as an *approximation*, and the rest of the paper investigates to what degree it can bear fruit. Before that, the next section attempts to improve our insights into its properties.

A toy world

The reader might feel confused by the various "spaces" invoked: a *physical* space within the brain, a *signal* space of neural dynamics, and a *parameter* space of the source model, not all of which can be visualized. Here, a simplified "toy world" captures the essence of the full problem while allowing for intuitive visualization. Sources and sensors are both in a plane (the former within a disc, the latter on its periphery), and source-to-sensor gain varies as the inverse of the square of their distance, $1/d^2$ (Fig. 2 A). The planar source space is homeomorphic to the cortical surface, while the inverse quadratic dependency mimics Maxwell's equations. Real-world complexities such as three-dimensional location and orientation, anatomy, and inhomogeneous or anisotropic conduction, are conveniently ignored.

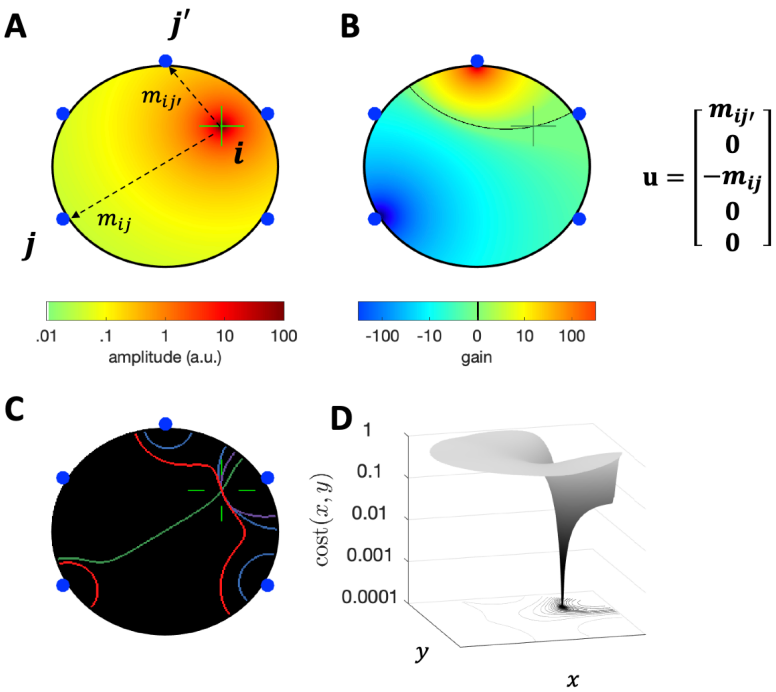

Fig. 2. "Toy world" consisting of an array of sensors (blue dots) placed at the periphery of a disc. A: A source within the disc (green cross) radiates within the plane with an amplitude that decreases with distance as $1/d^2$. B: Gain map for a linear combination of two sensors with weights chosen to cancel the source at the green cross (*null filter*). The weight vector (spatial filter) is shown on the right with coefficients ordered clockwise from the lower right. The black line is the *zero set* of this filter. C: Zero sets of four null filters for this same source. Their intersection signals the source position. D: Cost function of position within the source space (Eq. 3). The global minimum (zero) indicates the source position.

A source (green cross) radiates within the disc as shown as a color scale (Fig. 2 A), with a different amplitude at each sensor. Adding two sensors with appropriate weights produces a spatial filter with zero gain at the source position (a *null filter* for that source) (Fig. 2 B). As explained in SI, null filters for a given source form a subspace of dimension $J - 1$ (here 4) of the space of spatial filters over J sensors (SI), and thus there may be multiple null filters for the same source, each with a different gain pattern. The gain of a null filter is also zero at other positions in the disc, that form the *zero set* of that filter (black line in Fig. 2 B). The zero sets of all null filters intersect at the location of the source (Fig. 2 C), and if a cost function is defined

as in Eq. 3, its value is zero at the position of the source (Fig. 2 D). This illustrates the principle of zero set imaging explained earlier.

Is the zero unique? The question is difficult to answer in full generality, but some insights can be gained within the toy world. With J sensors and K null filters, the zero set for filter k consists of positions (x, y) such that:

$$\sum_j \frac{u_{jk}}{(x-x_j)^2+(y-y_j)^2} = 0, \qquad (4)$$

based on the toy-world version of the cost function of Eq. 3. Multiplying each term by the denominators of the other terms (to give all terms the same denominator), we have:

$$\sum_j u_{jk} \prod_{j' \neq j} \left[(x - x_{j'})^2 + (y - y_{j'})^2\right] = 0, \qquad (5)$$

which can be expanded into a multivariate polynomial equation in x, y of order $2(J-1)$. For two unknowns (x, y), we need at least two such equations (and thus at least two null filters). Bezout's theorem tells us that there may be up to $n = [2(J-1)]^2$ zeros, and thus there might indeed be multiple zeros of the cost function, albeit finite in number. Empirically, the zero seems often to be unique but it is worth keeping in mind the possibility that localization might sometimes be ambiguous. This and other issues are discussed in SI.

Simulated MEG

The toy world offers insight, but the effectiveness of the method needs checking within more realistic configurations. A source model with 15002 dipoles, positioned on the cortical surface with orientations normal to that surface, was calculated using Brainstorm (Tadel et al. 2011), based on a data set from the Brainstorm tutorial (SI). A dipolar source position was chosen (Fig. 3 A), and a single-column matrix of Gaussian noise, representing simulated source activity, was multiplied by the corresponding gain vector to obtain 274-channel simulated MEG data. PCA was applied to these data, and columns 2-274 of the transform matrix were retained as null filters for the source. The cost function (Eq. 3), plotted in Fig. 3 B, shows a sharply-defined

minimum at a position that matches that of the source. The precision of the location estimate depends only on the resolution of the source model.

Other configurations and source models are explored in SI, as are the sensitivity of the method to noise in the data, jitter of the forward model coefficients, and mismatch between real and model sources. In a nutshell: the method is quite tolerant to noise and jitter, but sensitive to a mismatch between real and model source, an issue that will become apparent with real data.

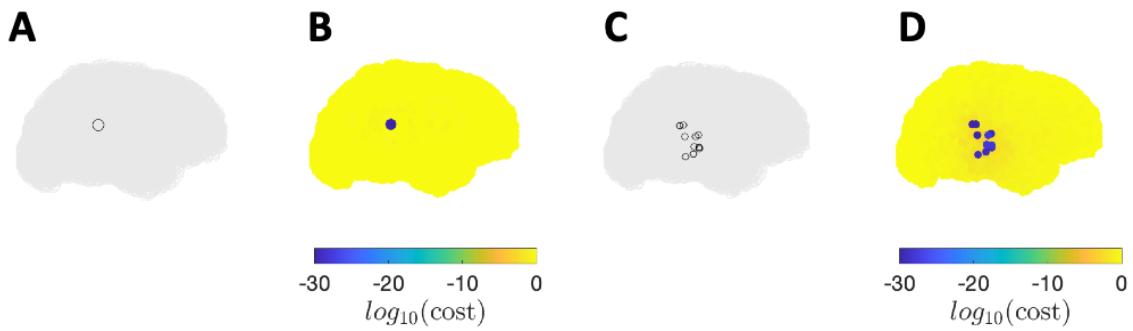

Figure 3. Localization of simulated MEG sources. A: Location of simulated source. B: Cost function, assuming a constrained orientation cortical surface model. C: Locations of 10 simulated sources with independent activities. D: Cost function.

An interesting case is *multiple sources* with independent activity. The time courses of $N = 10$ such sources (Fig. 3 C) were synthesized as a 10-column matrix of Gaussian-distributed values, which was multiplied by the corresponding columns of the model gain matrix to obtain 274-channel simulated MEG. Columns 1-10 of the PCA matrix capture all the variance, and thus columns 11-274 constitute null filters for all 10 sources. The cost function, plotted in Fig. 3 D, shows sharply-defined minima at positions that match those of the sources. This result holds even if the source waveforms are correlated, but breaks down if they are are linearly dependent, in particular identical (SI). The multiple-source case is interesting because it often

happens that linear analysis isolates a group of components from background activity, but cannot resolve them individually.

Real MEG and EEG

The method was applied to 274-channel MEG data recorded in response to repeated auditory stimulation in an odd-ball detection task (publicly-available data used by the Brainstorm tutorial, SI). The JD algorithm was applied repeatedly: first to isolate stimulus-locked activity, second to isolate button-press-locked activity, and third to isolate narrowband activity near 10 Hz (alpha). In each case, null filters were chosen from columns of the transform matrix excluding the first N ($N = 1$ for auditory and button press, $N = 10$ for alpha). A cost function was calculated over locations of a volume source model (dipoles on a 2 mm grid with arbitrary orientation). For the auditory response, sensors and source locations were restricted to left or right hemispheres and analyzed separately.

Figure 4 A and B show the cost function for the auditory response, for the left and right hemispheres respectively. Each pixel represents the minimum cost over a stack of voxels perpendicular to the viewing plane, thus deep sources are visible. Small values of the cost function are clustered in temporal regions of both hemispheres, as might be expected for auditory activity. However, compared to the simulation (Fig. 3), cost values remain relatively high (note the range on the colorbar), implying that the real source does *not* fit the model of a single dipole.

Visually, these patterns are roughly consistent with those found with techniques such as dipole fitting, beamforming, and minimum norm estimate (see SI for examples). An important difference is that those methods estimate the *distribution of activity* over source space, whereas the plots in Fig. 4 reflect the *distribution of fit* between the real source and model as a function of the model source position in its parameter space. Note that the model parameter space also includes orientation, not visible in these plots.

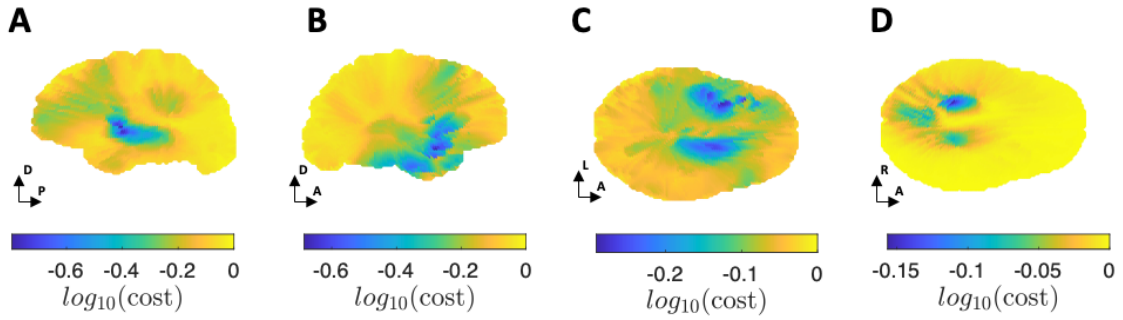

Fig. 4. Localization of sources in MEG. A: Cost function associated with the auditory response, left hemisphere. B: Same, right hemisphere. C: Cost function associated with the button press. D: Cost function associated with 10 Hz narrowband activity. Each pixel represents the minimum cost over a stack of voxels perpendicular to the viewing plane (minimum-take-all transparent brain).

Similarly, Fig. 4 C shows the cost function for the response preceding the button press, and Fig. 4 D shows the cost function associated with the first 10 components of the alpha response. Again, the relatively large cost values suggest that the cortical source is not a single dipole.

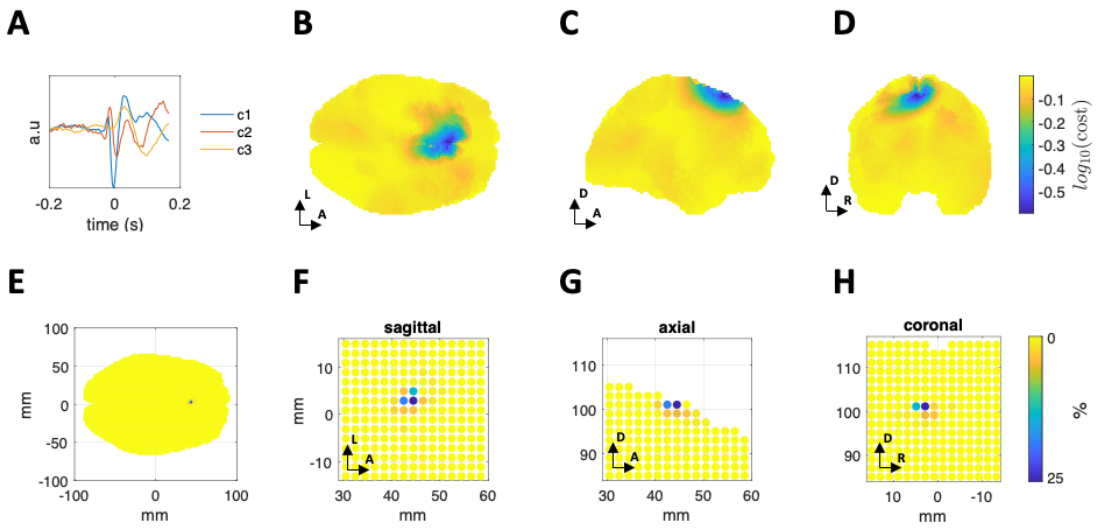

Fig. 5. Localization of interictal spikes in EEG. A: Waveforms of three components of a JD analysis time locked to interictal spikes. B-D: Cost function associated with the first (earliest) component,

based on a volume source model on a 2mm grid, projected on sagittal, axial, and coronal planes (minimum-take-all). E-H: empirical probability distribution of the position of the minimum cost, based on bootstrap resampling (n=1000).

The method was also applied to 29-channel EEG data recorded from an epileptic patient between seizures (Brainstorm tutorial). These data were analyzed with the JD method to isolate activity time locked to markers aligned with interictal spikes. Figure 5 A shows the time-averaged waveforms of the first three JD components, and Fig. 5 B-D shows the cost function associated with the first of these (earliest), viewed in sagittal, axial and coronal planes. As in Fig. 4, each pixel represents the minimum over a stack of voxels (transparent brain). The cost function is smallest at a superficial position on the left margin of the longitudinal fissure, roughly consistent with results from dipole fitting, beamforming, and minimum norm estimation (SI). Figure 5 E-H shows the empirical distribution of the position of the minimum, based on a bootstrap resampling (n=1000). The distribution is tight: 95% of trials fall within 4 mm of its centroid (however in the absence of ground truth one cannot say if that centroid coincides with the true source). Accurate localization of ictal and interictal activity is clinically important. Another example based on SEEG is described in SI.

Discussion

Zero-set imaging (ZSI) assumes distinct processing pathways for source activity waveforms (*what*) and their locations (*where*), relying on different columns of the data-driven analysis matrix. Null filters play the role of *observation* and the source model that of *prior* in the inference process from which the location is derived. Null filters can be understood as "sufficient statistics" of the EM recordings for this purpose.

Caveats. The theory assumes a perfect fit between the real source and a model source. For real data, however, a perfect fit was not found, and so we are left to interpret a pattern of imperfect fit to an inexact model (Fig. 4, 5 B-D). That pattern seems informative and consistent with other methods, but the theoretical basis to interpret it is uncertain. Hopefully, progress

in understanding underlying brain processes will lead to a more complete model that allow a perfect fit, obviating the need to interpret the pattern of imperfect fit. The question is further explored in the SI. In the meanwhile, caution is required.

Relation with data-driven linear analysis. Typical processing pipelines focus either on the *time course* of neural activity, location playing a secondary role, or on its *location*, the time course playing a secondary role. Linear analysis methods (PCA, ICA, etc.) are prominent in the first, but often relegated to an artifact-rejection role in the second. For ZSI, linear analysis is centre stage, but in contrast to standard pipelines in which non-target columns of the analysis matrix are discarded, here those columns are retained to feed the localization pathway.

Relation with dipole fitting. ZSI resembles dipole fitting in that a source model is scanned for a fit to the data. In dipole fitting, the fit is made directly to values at the sensors, for ZSI the fit is via the null filters which gather information over a larger temporal span and (potentially) factor out interference. The dipole fit draws on J values, zero-set imaging on up to $J-1$ values. Dipole fitting fails if non-target activity leaks into the target component (i.e. the pattern across sensors is noisy), ZSI fails only if the target leaks into non-target components (i.e. non-target columns of $\mathbf{U}$ are not perfect null filters).

Relation with beamforming. Beamforming and ZSI both rely on spatial filtering to cancel sources. Beamforming involves a single spatial filter that cancels up to $J-1$ *interfering* sources, whereas ZSI involves up to $J-1$ spatial filters that all cancel the same *target* source. Beamforming is used in two distinct ways, either to design a spatial filter to isolate a target *based on knowledge of its location*, or to locate a source by scanning the source space in search of positions for which the spatial filter recovers the target's activity, *based on knowledge of the target's time course*. ZSI assumes neither. Toy-world simulations suggest that beamforming may lead to ambiguous localization (SI). While beam-forming is known to be sensitive to correlation between sources, ZSI is relatively robust.

Relation with MNE. MNE, in contrast to ZSI and other methods that estimate location parameters, directly estimates an "inverse" matrix $\mathbf{U}$ to transform the observation time series matrix $\mathbf{X}$ into an estimated source time series matrix $\hat{\mathbf{S}}$, using a constraint of minimum current norm to overcome indeterminacy. "What" and "where" are both packed into $\hat{\mathbf{S}}$. The minimum norm constraint enforces a prior, that is a priori rational and leads to plausibly smooth distributions of estimated activity, however that prior is hardwired into the calculation of the inverse matrix. With ZSI a similar prior could be introduced explicitly in the design of the source model (and/or rendering model).

Figure S11 compares estimates from several methods applied to the same data. Methods cannot be ranked on this basis, given the lack of ground truth and the sensitivity to tuning and visualization parameters, The safest conclusion is that they are roughly consistent, but diverge in detail. A crucial difference is that the plot for ZSI represents *fit* to a dipole model space, whereas plots for other methods are better interpreted in terms of *activity* over space (and time). The accuracy of classic methods is limited, as quantified in terms of peak localization error or point spread function (Barnes et al 2004; Hauk et al 2022; Jaiswal 2020; Lütkenhöner 2003), but there is no such limit in ZSI. Whereas beamforming and minimum norm require a noise covariance matrix, ZSI does not, although similar information may intervene in the data-driven analysis stage (de Cheveigné and Parra 2014).

Strengths. The conceptual basis of ZSI is unusual, suggesting it might be *complementary* to established approaches. An appeal is the clear distinction between "what" and "where" pathways and, within the latter, between observations (encapsulated in the null filters) and priors (encapsulated in the source model). In neuroscience it is common to speak of distinct what/where pathways *within* the brain (Schneider 1969), which is possibly coincidental but might also reflect similar computational constraints (Chikkerur 2010; Cisek 2019). Unlimited spatial resolution is a strong point for ZSI, however this must be tempered by its sensitivity to a mismatch between real and model source, as reflected by high values of the cost function in Figs. 4 and 5 B-D. However, that weakness might be turned into a strength, as it offers a

basis to develop better models and quantify their benefit, within a wider search loop that is not just algorithmic but also scientific. The method is applied here to brain data, but it might be more widely useful, e.g. in acoustics or geophysics. Based on the reciprocity theorem, sensor positions can in principle be localized from a calibrated array of sources, thus contributing to instrumentation. These and other aspects remain to be explored.

Returning to Plato's allegory, if the shape of the shadow were scaled, turned into a mask, and placed *between* the light source and the object, it would perfectly obliterate that object. This is an almost perfect metaphor for a null filter. Just as multiple light sources cast multiple shadows of the same object, together informing us of its volume in space, multiple null filters constrain the location of a source within the brain. Looking at an illuminated object tells us *what* that object is, and looking at the shadows that it casts tells us accurately *where* it is, an idea that works also for the brain.

Acknowledgments

Previous drafts benefitted from comments of Israel Nelken, Lucas Parra, and Malcolm Slaney.

Code availability

https://github.com/adc2/WhatVSWhere

Supplementary Information

Available from the BioRxiv site. That document contains important details not included in this manuscript, and should be kept at hand when reading this document.

References

al Haytham I. (1030) Book of optics (Hatfield 2002, Imbert 2020).

Baillet, S., Mosher, J.C., and Leahy, R.M. (2001) Electromagnetic brain mapping. IEEE Signal Processing Magazine, 18(6):14–30.

Baillet, S. (2010) The Dowser in the Fields: Searching for MEG Sources, in MEG, an Introduction to Methods, edited by Hansen, P.C., Kringelbach, M.L., Salmelin, R., Oxford University Press, 83-123.

Barnes, G. R., Hillebrand, A., Fawcett, I. P., & Singh, K. D. (2004). Realistic spatial sampling for MEG beamformer images. *Human Brain Mapping*, *23*(2), 120–127. https://doi.org/10.1002/hbm.20047

Cisek, P. (2019). Resynthesizing behavior through phylogenetic refinement. Attention, Perception, & Psychophysics, 81(7), 2265–2287. https://doi.org/10.3758/s13414-019-01760-1

de Cheveigné, A., Parra, L. (2014) Joint decorrelation: a versatile tool for multichannel data analysis. *NeuroImage*, 98, 487-505, doi: 10.1016/j.neuroimage.2014.05.068.

de Cheveigné, A. (2020) ZapLine: A simple and effective method to remove power line artifacts, *Neuroimage*, 207, 116356, doi: 10.1016/j.neuroimage.2019.116356.

de Cheveigné, A (2021). Shared Component Analysis. *NeuroImage*, 226, 117614.

Frauscher, B., Von Ellenrieder, N., Zelmann, R., Doležalová, I., Minotti, L., Olivier, A., Hall, J., Hoffmann, D., Nguyen, D. K., Kahane, P., Dubeau, F., & Gotman, J. (2018). Atlas of the normal intracranial electroencephalogram: Neurophysiological awake activity in different cortical areas. *Brain*, *141*(4), 1130–1144. https://doi.org/10.1093/brain/awy035

Friston, K. (2002). Functional integration and inference in the brain. *Progress in Neurobiology*, *68*(2), 113–143. https://doi.org/10.1016/S0301-0082(02)00076-X.

Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., Parkkonen, L., & Hämäläinen, M. S. (2014). MNE software for processing MEG and EEG data. *NeuroImage*, *86*, 446–460. https://doi.org/10.1016/j.neuroimage.2013.10.027

Hatfield G. (2002) Perception as unconscious inference. In Heyer D., & Mausfeld R. (Eds) Perception and the physical world: Psychological and philosophical issues in perception (pp. 113–143). John Wiley and Sons.

Hauk, O., Stenroos, M., & Treder, M. S. (2022). Towards an objective evaluation of EEG/MEG source estimation methods – The linear approach. *NeuroImage*, *255*, 119177. https://doi.org/10.1016/j.neuroimage.2022.119177.

Helmholtz H. (1867). Handbuch der Physiologischen Optik (English tranl.: 1924 JPC Southall as Treatise on Physiological Optics) Voss.

Imbert M (2020). La fin du regard éclairant. Une révolution dans les sciences de la vision au XIe siècle. Ibn al-Haythan, Vrin, Paris.

Jaiswal, A., Nenonen, J., Stenroos, M., Gramfort, A., Dalal, S. S., Westner, B. U., Litvak, V., Mosher, J. C., Schoffelen, J.-M., Witton, C., Oostenveld, R., & Parkkonen, L. (2020). Comparison of beamformer implementations for MEG source localization. *NeuroImage*, *216*, 116797. https://doi.org/10.1016/j.neuroimage.2020.116797

Kawato, M., Hayakawa, H., & Inui, T. (1993). A forward-inverse optics model of reciprocal connections between visual cortical areas. *Network: Computation in Neural Systems*, *4*(4), 415–422. https://doi.org/10.1088/0954-898X_4_4_001

Lütkenhöner, B. (2003). Magnetoencephalography and its Achilles' heel. *Journal of Physiology-Paris*, *97*(4–6), 641–658. https://doi.org/10.1016/j.jphysparis.2004.01.020

Oostenveld, R., Fries, P., Maris, E., Schoffelen, J.-M. (2011) FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. Computational Intelligence and Neuroscience, 2011:156869.

Parra, L. C., Haufe, S., & Dmochowski, J. P. (2019). Correlated Components Analysis—Extracting Reliable Dimensions in Multivariate Data. *Neurons, Behavior, Data Analysis, and Theory*, *2*(1). https://doi.org/10.51628/001c.7125

Tadel, F., Baillet, S., Mosher, J. C., Pantazis, D., & Leahy, R. M. (2011). Brainstorm: A User-Friendly Application for MEG/EEG Analysis. Computational Intelligence and Neuroscience, 2011, 1–13. https://doi.org/10.1155/2011/879716

Schneider, G. E. (1969). Two Visual Systems. *Science*, *163*, 895–902.