Published in final edited form as:

J Neurosci Methods. 2018 October 01; 308: 88–105. doi:10.1016/j.jneumeth.2018.06.019.

Geometric classification of brain network dynamics via conic derivative discriminants

Matthew F. Singh^{1,2,3,*}, Todd S. Braver², and ShiNung Ching^{3,4}

¹Department of Neuroscience, Washington University in St. Louis, St. Louis, MO, USA

²Department of Psychology, Washington University in St. Louis, St. Louis, MO, USA

³Department of Electrical and Systems Engineering, Washington University in St. Louis, St. Louis, MO, USA

⁴Department of Biomedical Engineering, Washington University in St. Louis, St. Louis, MO, USA

Abstract

Background: Over the past decade, pattern decoding techniques have granted neuroscientists improved anatomical specificity in mapping neural representations associated with function and cognition. Dynamical patterns are of particular interest, as evidenced by the proliferation and success of frequency domain methods that reveal structured spatiotemporal rhythmic brain activity. One drawback of such approaches, however, is the need to estimate spectral power, which limits the temporal resolution of classification.

New Method: We propose an alternative method that enables classification of dynamical patterns with high temporal fidelity. The key feature of the method is a conversion of time-series data into temporal derivatives. By doing so, dynamically-coded information may be revealed in terms of geometric patterns in the phase space of the derivative signal.

Results: We derive a geometric classifier for this problem which simplifies into a straightforward calculation in terms of covariances. We demonstrate the relative advantages and disadvantages of the technique with simulated data and benchmark its performance with an EEG dataset of covert spatial attention. We reveal the timecourse of covert spatial attention and, by mapping the classifier weights anatomically, its retinotopic organization.

Comparison with Existing Method: We especially highlight the ability of the method to provide strong group-level classification performance compared to existing benchmarks, while providing information that is complementary with classical spectral-based techniques. The robustness and sensitivity of the method to noise is also examined relative to spectral-based techniques.

*Corresponding Author. f.singh@wustl.edu 600 S. Euclid Ave., 63110, St. Louis, MO, USA.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Conclusion: The proposed classification technique enables decoding of dynamic patterns with high temporal resolution, performs favorably to benchmark methods, and facilitates anatomical inference.

Keywords

Pattern Classification; Dynamical Systems; Neural Dynamics; EEG; Machine Learning; Directional Statistics

1. Introduction

Understanding how different brain regions interact in task-dependent ways is a key goal of cognitive neuroscience. In this regard, a frequent aim of neural data analysis is to characterize the spatiotemporal patterns present within brain activity and, subsequently, to enable the association of such patterns with specific cognitive states. Perhaps the most classic of such approaches, certainly within the domain of electrophysiological brain recordings, is time-frequency analysis which involves the projection of data into the Fourier domain in order to examine synchronization, spectral power and cross-channel (i.e., network) relationships.

Recently, increased attention has been directed toward characterizing neural recordings in terms of their underlying dynamics via more holistic descriptions [1, 2, 3, 4]. Such approaches include methods for characterizing chaos and instability reflected in neural recordings via estimation of Lyapunov exponents [5, 6] and manifold reconstruction methods such as Takens embedding [7, 8], which has been used to infer directed information flow between brain regions without overt analysis of rhythmicity [9]. As dynamical systems-based characterizations, these approaches are unified in that interest is placed on directly characterizing the underlying vector field that ultimately governs the time-evolution of the observed recordings. The vector field describes the instantaneous evolution of dynamical systems, i.e.-the time-derivatives of its state variables. In this spirit, here we propose a technique for elucidating properties of the vector field that involves analysis not of the observed time-series itself, but rather on the derivatives: the ‘velocity’ trajectories of neural activity. We use the term ‘velocity’ in the dynamical systems sense in which system states, such as a particular combination of cell voltages, are referred to as ‘positions’ within the state space and an observed sequence of ‘positions’ (states) forms a trajectory. The ‘velocity’ is thus considering how the system states change in time. It turns out that this rather straightforward transformation from considering system position to system velocity brings about several conceptual and practical advantages for classifying neural dynamics. Our approach is motivated in part by recent results in monotone dynamical systems theory that provide explicit characterizations of how the attractor landscape of a dynamical system relates to the geometry of its derivative phase flow ([10, 11, 12, 13]). In other words, by analyzing the observed velocity trajectories we can make certain deductions regarding the properties of the underlying neural circuit dynamics. The spatial resolution of these deductions match the spatial resolution of the recordings themselves. Exploiting this notion, we describe a novel approach to decode different cognitive states by geometrically deducing and characterizing their corresponding velocity trajectories. As will be seen, this approach

enjoys many advantageous statistical properties. In the particular case focused on in our paper, it reduces to an intuitive comparison of covariance geometry of the derivative time-series, which is amenable to classification using conic boundaries. We proceed to demonstrate the efficacy and ease of the proposed paradigm in a variety of examples, then apply it to an electroencephalogram (EEG) dataset [14] wherein we show its interpretability and performance in decoding covert spatial attention.

2. Classification for Hypothesis Testing in Neuroscience

Classification techniques play two distinct roles within neuroscience: for developing basic scientific inference as with multivariate pattern analysis (MVPA); and for developing brain-based interfaces such as brain-computer interfaces (BCI). The technique that we now propose has been developed with the former objective (scientific inference) in mind, although BCI applications are also possible. MVPA is a pattern-based approach to studying brain function, in which investigators attempt to decode whether activation patterns contain significant information regarding task events or cognitive states. A common goal of MVPA studies is to determine the anatomical regions in the brain that contain such information ([15]). Other studies have used MVPA-type analyses to test temporal rather than spatial hypotheses regarding brain function, as with studies of mnemonic periodic-replay ([16, 17]).

For all of these research questions, linear classification methods are by far the most common, and include logistic regression, linear discriminants, correlation, or support-vector machines ([18]). However, a continual goal of methodological research on MVPA applications for neuroscience is to search for enhanced classification approaches. One potential direction is to leverage the recent advances with deep learning approaches (i.e., convolutional neural networks, CNNs) that have shown significant promise within machine learning applications. A few studies have successfully adopted such approaches to classify neural data, for example by using pretrained image-classification networks ([19],[20]). Another potential direction for MVPA methodological development that has been more widely explored is to improve upon the neural features that are utilized by classification algorithms, such as encoding models ([21, 18, 22]) or feature selection ([17]).

In the current study, we also examine changes to the neural features utilized by classifiers, but adopt a qualitatively different approach, which instead considers a new form of information, namely the temporal evolution of brain activity. In particular, previous work using MVPA for EEG/MEG decoding has employed classification analysis using band-limited power, so that the classifier input is a periodic representation of brain activity. In contrast, the proposed technique uses a temporal derivative, which is a powerful transformation that can convert any spatiotemporal pattern (not just periodic ones) into another spatial feature. However, as we later discuss, this form of information is almost never separable with a linear boundary. Consequently, we develop an analogous classifier using a (quadratic) conic boundary. In practice, a positive classification using our method generates the following information: 1) a characterization of how the neural activity for the chosen EEG channels evolves differently in time based upon the task classes (the information is dynamically coded) and 2) a detailed temporal signature of the differential patterns of temporal evolution in the two classes.

In contrast to using complex formalisms such as deep learning, our classifier is explicit and employs only one form of information: the direction (covariance) of the temporal derivative. As a result, there is an explicit mapping between the classifier and the geometry of the system's temporal evolution. This feature aids the classifier's interpretability. In contrast, limited interpretability is seen as one of the greatest drawbacks of approaches such as convolutional neural networks and, to a lesser extent, linear decoders as well ([23]). The primary contribution of the proposed technique for neuroscience investigations is that it provides information regarding how spatiotemporal patterns of neural activity relate to task states. This information is characterized in terms of the interaction of neural populations across time. Such a characterization promotes mechanistic inference regarding the nature of task-based interactions between neural generators. As opposed to simply linking task states with mathematical descriptions of brainwaves (i.e. spectral power) inferences are made in terms of the temporal interactions between neural populations that characterize a task state. Any experimental inferences drawn from applying the classifier are made at the spatial and temporal scales in which the data is recorded: e.g. interactions between large cortical populations as indexed by the EEG timeseries.

In the current study, we first provide simulated examples in which the system-states and corresponding inferences are made at the cellular scale. Next we provide analysis of the cortical populations that underly task-states in humans using an experimental dataset from a benchmark EEG experiment ([14]). This benchmark dataset has subsequently been treated using conventional MVPA approaches involving extraction of rhythmic activity, in efforts to reveal circuit dynamics that underlie covert spatial attention ([24]). Critically, we demonstrate the power of the ability of our approach to generate new neuroscientific insights regarding both the anatomical interactions and the precise temporal dynamics that give rise to covert spatial attention (see 3.4).

3. Results

3.1. Conic Analysis Reveals Underlying System Dynamics

We consider time-series (here, neural recordings) denoted as

$$x_j^{(i(t))}(t) = [x_{j,1}^{(i(t))}(t), x_{j,2}^{(i(t))}(t), \dots, x_{j,n}^{(i(t))}(t)]^T \quad (1)$$

where n denotes the number of observed variables, $j \in \mathbb{Z}^+$ is a trial index and $i(t) = 1, 2, \dots, K$ denotes a class label (e.g., a cognitive state). Our overall goal is to deduce the class label (at each time) by analyzing the derivatives of $x_j(t)$.

The proposed method for neural data analysis and classification in this context consists of two steps. First, for each class (i), we obtain an estimate of the derivative time-series, $\dot{x}^{(i)}(t)$. Second, we use generalized cones as a geometric classifier for individual points of the derivative time series. A generalized cone can be understood as an object in Euclidean space that is invariant to scaling: if a vector $\mathbf{v}$ is an element of a cone, then so is $\alpha \mathbf{v}$ for any non-negative $\alpha \in \mathbb{R}^+$. Thus, the cone is a fundamentally directional object. Our approach will

involve using such cones to partition Euclidean space into sets of derivatives attainable only by specific classes. This procedure can reveal how a system transitions between states (but not necessarily the speed at which these transitions occur).

Our subsequent results will demonstrate that this approach complements spectral analysis (e.g., assessment of oscillatory power), which, while an assay of dynamics, does not explicitly include directional information. More specifically, unlike many existing methods for spatiotemporal classification, we do not operate upon summary parameters such as power within a specific frequency band ([16], [17]). Instead, it is the individual data points of the derivative time series that are used as the basis for classification since these contain inherent spatiotemporal information. In this way, our classification approach can capture the instantaneous rules, (i.e., the dynamics), governing how the (class-specific) time-series evolve. This approach is particularly favorable as it reduces a spatio-temporal classification problem in the original time series to a spatial-only problem in the derivative time-series. As we demonstrate later, this reduction significantly aids analysis and interpretation.

The motivation for such an approach stems from recent results in the theory of smooth monotone dynamical systems. For our purposes, a smooth dynamical system consists of a state space (here, we consider this to be $\mathbb{R}^n$) coupled with a phase flow, or vector field that describes, for each state, how the system evolves forwards (and, backwards) in time. Recent research has linked geometric knowledge about the vector field (i.e., the derivatives along system trajectories or paths) to the asymptotic behavior of the system state variables. Such characterizations have been made possible by formally studying the extent to which the derivatives along a path may be confined within generalized cones ([10],[11],[12],[13]). Although a detailed discussion of monotone dynamics is beyond our current scope (see e.g., [25]), these recent advances motivate our use of cones for classification in derivative space.

To the extent that spatial patterns are useful for classification, the derivatives can provide us crucial information regarding the underlying system dynamics (spatiotemporal patterns). Perhaps an intuitive analogy is that of vertical ascent for two sorts of aircraft: (fixed-wing) airplanes and helicopters. While helicopters can make a direct vertical translation, airplanes require strong lateral velocity before any vertical motion is obtainable. Thus, for an airplane, reaching a point just above the current one requires a complex (spiral-like) velocity trajectory. Observing this velocity trajectory is thus highly informative in deducing the vehicle in question. Aerodynamic constraints also greatly constrain the turning radius of an airplane. As a result, the set of obtainable velocities for an airplane relative to its body-frame differs greatly from that of a helicopter, despite both being able to reach any point in airspace. Similarly, neural circuits are seemingly able to produce a wide variety of spatial activity patterns, although the set of attainable derivatives may be more limited. For example, in Figure 1 we consider two different toy networks, each composed of three continuous, recurrent Hopfield-model neurons ([26], see Methods; Fig 1 A,B). The two networks (systems) feature a common connection scheme, but differ in their internal dynamics. In the second system ("System B") every cell exhibits greater decay (in neuroscience parlance, leak conductance) (1B). The systems produce qualitatively similar trajectories wherein all initial conditions lead to a damped oscillatory response for both cases (1 C,E). Despite this similarity, the modeled systems differ in the oscillatory envelope

of their response (1 F). In the ideal case of noise-free uninterrupted observations, the trajectories of these systems are easily discriminable in the original phase space (1E). However, trial-based designs that are the primary workhorse of experimental neuroscience recordings (e.g., single- or multi-unit electrophysiology, EEG, fMRI etc.) generate large numbers of short duration observations, in which the variability of initial states present at the trial start can result in significant entanglement of activation trajectories. When we replicate these conditions for the two simulated networks, their recordings become inextricable in the phase-space of cell activity (1C). However, the systems remain easily separable in the phase space of their derivative signals (1D). As the derivative space is indicative of instantaneous changes in activity, dynamical features of the system are captured in individual observations, so the systems may be differentiated regardless of observation length or initial conditions. Of course, this classification can be performed to a certain degree via direct methods (e.g., linearization) if the full analytical model is already known. Our goal, in contrast, is to make the same dynamical inferences (i.e., about the structure and equations governing the system) based only upon observation of the activity. Using our proposed approach, we observe that derivatives are invariant within a particular directional region (1D). Based upon this invariance, we can directly surmise that System A (1A) cannot directly move across the Cell 1 axis as its derivatives can never point in that direction (1). Instead, its derivative space suggests System A could only traverse the space by moving indirectly in that direction, just as physical constraints prevent vertical translation of a fixed-wing aircraft. One possible solution, is for System A to generate a series of oscillations leading to traversal of the space along the Cell 2 axis, which is of course consistent with the actual dynamics (1E,F). In contrast, derivatives of System B are rarely orthogonal to the Cell 2 axis (1D) and the corresponding prediction holds that System B traverses the Cell 2 axis more directly (with fewer oscillations) than System A (1E,F).

To operationalize the above analysis, we propose a classifier to discriminate systems based upon their directional derivatives. Most classifiers used to analyze neural data, such as linear support vector machines ([27]), assume that the classes differ in some combination of their univariate means and thus form linear boundaries ([28]). For the case of derivatives, this assumption is not valid. In fact, a well-known result in dynamical systems theory, which we rely upon, is that the derivatives of a smooth, bounded dynamical system must be balanced (e.g. [11]): over long scales the mean derivative is always zero. Otherwise, the system would grow infinitely along the direction of the mean derivative. In the case of a neural system, the mean change in the measured activity variable (e.g., voltage) must be zero over a sufficiently long recording, lest the activity grow infinitely. As the mean derivative is thus irrelevant, the motivating results in dynamical systems theory concern invariance along cones. Due to scale-invariance, conic boundaries discriminate between classes based upon the directional components of the data, rather than magnitude features (such as speed of evolution), and are always centered at zero. With very mild assumptions upon the way derivatives are distributed for each class, we can derive a Bayes-optimal classifier for directional data (see Materials and Methods) which reduces to a comparison of covariance matrices. The boundaries of this classifier form a quadratic cone.

For an unlabeled time-series $x^{(i)}(t)$ we define our quadratic-cone classifier as:

$$i(t) = \arg \min_k \left(\frac{\dot{x}(t)^T \Sigma_{(\dot{x}_k)}^{-1} \dot{x}(t)}{\det |\Sigma_{(\dot{x}_k)}^{-1}|^{1/n}} \right) \quad (2)$$

Here, $\Sigma_{(\dot{x}_k)} \in \mathbb{R}^{n \times n}$ is the covariance matrix associated with the derivative time-series of class k . At an implementation level, this covariance can be estimated numerically by collecting labeled training data from each class. Thus, the classifier operates on a point-wise basis: assigning a class label to each individual data-point of the derivative time-series. In the two class case with class labels ± 1 we describe the cone in terms of a single matrix P .

$$i(t) = \text{sgn}(\dot{x}(t)^T P \dot{x}(t)) \quad (3)$$

$$P := \frac{\Sigma_{(\dot{x}_{[-1]})}^{-1}}{\det |\Sigma_{(\dot{x}_{[-1]})}^{-1}|^{1/n}} - \frac{\Sigma_{(\dot{x}_{[+1]})}^{-1}}{\det |\Sigma_{(\dot{x}_{[+1]})}^{-1}|^{1/n}} \quad (4)$$

We denote the induced cone as $C(P) := \{y \in \mathbb{R}^n \mid y^T P y \leq 0\}$. The decoding rule may be restated in terms of the cone $C(P)$ as $\dot{x}(t) \in \text{Int}(C(P)) \Leftrightarrow i(t) = -1$ with Int indicating the interior (i.e. when the inequality is strict). We present further information concerning practical and efficient calculations of these terms and the case of singular covariances (i.e. fewer samples than channels) in Materials and Methods. Formally, we use the term ‘derivative’ in the sense of backward approximation by Newton’s difference quotient, however the scale-invariance property of cones makes such formalism unnecessary and the backward derivative may be replaced by the difference time series: $\dot{x}(t) \simeq x(t) - x(t - \Delta t)$. In general, we have had success implementing the derivative in this form without further preprocessing, even though an extensive literature exists concerning more sophisticated methods of numerical differentiation (i.e. [29]).

3.2. Conic Geometry Blindly Identifies Upstream Network States

To test the degree to which conic-invariance describes changes in functional structure ([30]), we simulated the case in which an experimenter tries to study a complex neural network but may only access a small number of outputs (Fig. 2A). We formalized this scenario by simulating a chaotic ([31]) system’s effects on a downstream recurrent network of four bursting Hindmarsh-Rose model neurons ([32]) with recordings only available for the first three cells (Fig. 2B,C). Cells are either excitatory (green) or inhibitory (red). The fourth, unobserved, inhibitory neuron forms the main recipient of upstream input and is normally quiescent save when activated by the upstream component. When the upstream component is in the “off” state, the network acts as a delayed negative feedback loop. When the upstream component is in the “on” state, a new indirect inhibitory path is opened via Cell 4.

Thus, by altering the activity of a mediating cell, the upstream component's activation dictates possible interactions in the downstream network.

In order to blindly determine the upstream component's activity state we performed unsupervised clustering of downstream neural voltage. Using the three observed neurons we performed 2-means clustering ([33]) for either the spectral power of the original voltage trace ($V(t)$) or, for conic analysis, the covariance of the derivative voltage traces ($\Sigma_{(\dot{V})}$).

Spectral power and covariances were calculated over equal-length non-overlapping time windows. Using the k-means algorithm ([33]) produces both a set of centroids (cluster means) and labels assigning each data point to the cluster to which it is closest. For the spectral classification we directly used the labels assigned to each time bin by k-means as the training class. For conic classification we discarded the original labels for each time bin and instead generated new labels through our proposed conic method. The two covariance centroids were treated as class covariances and we then classified each time point according to (2). Because the conic classifier provides instantaneous class labels, we assigned whole-bin classes post hoc in a winner-take-all paradigm based on the sum:

$$\arg \min_k \left(\sum_{t \in T_l} \left(\frac{\dot{x}(t)^T \Sigma_{(\dot{x}_k)}^{-1} \dot{x}(t)}{\det |\Sigma_{(\dot{x}_k)}^{-1}|^{1/n}} \right) \right), \quad (5)$$

where T_l denotes the l^{th} bin.

We repeated this simulation/analysis sequence while varying simulation parameters for a total of 136 simulations per case (see Materials and Methods). The parameters of interest were the variance of intrinsic system noise (a Brownian motion), measurement noise, coupling strength with the upstream component, and the window length for each bin. Results demonstrate a general advantage of conic classification over spectral clustering (Fig. 2D). This relationship was most pronounced for strong coupling strength or low measurement noise. Unlike conic decoding, spectral decoding was largely unaffected by the white-noise added to the recorded voltages. Spectral performance exceeded conic decoding in extreme cases of measurement noise, which is well known to degrade estimation of derivative time series. However, both methods were resistant to intrinsic noise within the system. These two noise varieties (intrinsic vs. measurement) had equal variance ranges, but unlike measurement noise, intrinsic noise interacts with system states. For instance, Brownian motion in the voltage variable has less impact during the fall of an action potential than near the firing threshold, so this differential sensitivity may preserve those portions of dynamics to which the conic method is most sensitive. We conclude that for most parameter choices within this simulation, the proposed conic decoding method exceeded spectral classification of upstream system states but is susceptible to idealized measurement noise. However, we later demonstrate that the conic method is actually resistant to many (realistic) sources of experimental noise and artifact (see Discussion).

3.3. The Conic Method Produces Instantaneous Decoding of Covert Spatial Attention

In this section, we transition from simulated to real experimental data. Specifically, we consider the method's application to EEG data recorded during a spatial covert-orienting task, taking the data from a publicly available database ([34],[14]). During the task a briefly presented central cue indicated where a subsequent task-relevant target would appear with 80% accuracy. The central cues had near identical visuo-spatial features and subjects were instructed to focus upon the cue (see Materials and Methods). This design feature implies that the putative neural dynamics associated with particular orientations correspond to covert spatial attention rather than simply the direction of gaze. The original results using this dataset have demonstrated that the spatial distribution of alpha power (across channels) may be used to determine to which attended location subjects are covertly orienting. These previous analyses ([34]) were performed pair-wise using L1-regularized logistic regression for alpha-band power in parietooccipital electrodes and the authors presented the best pair accuracy for each subject as well as the number of significantly classified pairs per subject out of a total of 15 (see Table 1). In a first set of analyses, we use all trials with delays of at least 1300 ms and only consider data after the first 250 ms post-cue. We consider the same metrics to those previously reported for comparison ([14]) and all reported accuracies for the current analyses are based upon leave-one-out cross-validation. Accuracies are averaged between conditions to prevent spurious decoding from unequal class sizes.

3.3.1. Group-Level Conic Performance Exceeds the Limits of alpha-band Classification—In the case of pairwise classification, current results appear favorably [Fig. 3 A,B] to those previously published ([14]). In their analysis, Schmidt and colleagues ([14]) provide two metrics: the number of significantly discriminated pairs and the maximal pairwise accuracy. We emphasize comparisons in terms of number of significant pairs as these take into account more subtle spatial comparisons whereas the previously reported maximal pairwise accuracy corresponded to spatially opposite locations for seven out of eight participants in [14]. Results demonstrate that the conic classifier is, on average, superior to α -band based decoding in terms of number of significant pairwise classifications (Table 1). However, the conic advantage in maximal pairwise accuracy was insignificant (Table 1).

3.3.2. Conic and Spectral Decoding Form Complementary Techniques—As with traditional approaches applied to the dataset ([14]), conic decoding exhibits variable performance across subjects. Interestingly, however, conic and alpha-band decoding performance across individuals tended to be negatively correlated in terms of the number of significant pairs per subject (Fig. 3C). Due to the small sample size and suspected outlier (Subject 8), we performed a rank-order test with the outlier and parametric test both with and without the outlier. Using the full data with rank-order correlation demonstrated a negative relationship between conic and spectral performance (Table 2). The parametric relationship with the full data did not reach significance due to the low sample size, but this relationship greatly increases when removing subject 8 for whom no pairs were significantly classified by the conic method (Table 2). However, the low sample size makes it difficult to determine whether the suspected outlier does in fact represent anomalous data or just the extreme poor end of conic classification ability. Thus, while the small sample size interferes

with statistical inferences, results suggest that not only are these decoding methods complementary, but also that many subjects exhibiting poor spectral-based decoding may actually possess strong conic classification. In summary, the proposed conic method and spectral decoding may form complementary approaches, which is to say that they may be decoding different forms of information. Consequently, subjects insensitive to one technique may benefit from the other. Although conic performance was significantly better on average, spectral decoding was superior for two subjects (#1 and 8) so the conic classifier may still be missing some information that is present in the spectral domain and using the two methods in conjunction might provide superior results.

3.3.3. Conic Weights Generate Retinotopic Maps of Covert Attention—The conic method also appears to replicate previously published results showing that spatial aspects covert-orienting of covert orienting are most prominent of posterior electrodes. In fact, spatial location is highly prominent in the conic weights (see Materials and Methods) and, in this case, provides anatomical localization of dynamical information (Fig. 4) without requiring search-light type analyses ([35]). To determine whether conic weights provide anatomically meaningful information we performed a spatial mapping of the weights for each cone-generating matrix. We defined the map for each condition in terms of the contrast ‘location x’ vs. the average hemifield (3 locations) opposite ‘x’. We averaged the covariances across subjects for viewing purposes (group level), but in most cases similar results held at the individual subject level as well. As a matrix, the conic classifier is most naturally visualized as a weighted graph. However, to produce a spatial map we assigned each channel a weight based upon its contribution to each of the matrix’s eigenvectors (see Materials and Methods). We find that the regions representing the covertly attended location experience decreases in the amplitude of derivatives (negative weights) while regions representing the opposite spatial location experience relative increases in the amplitude of derivatives (positive weights) (Fig. 4). Thus, covert attention may be associated with neural activity whose variability (as distinct from variance) is more restricted in regions representing the covertly attended location.

3.3.4. Instantaneous Decoding Tracks the Evolution of Neural States—The conic-geometry of these instantaneous changes is easily visible for the 2-class case (i.e. lower left vs. upper locations) when projecting the derivative time series into the eigen-coordinates (see Materials and Methods) with the largest eigenvalues (Fig. 5A,B). These may be thought of as ‘principal components’, but as they relate the differential activation between classes they also possess sign information. Further, the absolute magnitude of each eigenvalue indicates how strongly selective it is for the class associated with its sign (Fig. 5B). The time-series of these components indicate that the information differentiating between classes is associated with increased dynamics of certain spatial patterns (the eigenvectors). Critically however, these dynamics are not necessarily associated with consistent increases in particular frequency bands (Fig. 5B). Point-wise classifier performance may be visualized by projecting data into the plane of positive and negative eigenlengths (see Materials and Methods) which indicates the combined ‘magnitude’ of dynamics associated with each class (Fig. 5C). Data is displayed for the median-performing

subject in terms of number of significantly classified pairs (number four). Analogous plots for the best and worst-performing subjects are included in Materials and Methods.

3.4. Conic Decoding Reveals Distributed Interactions at High Temporal Resolution

In a recent paper, Samaha and colleagues ([24]) used the same dataset to test the hypothesis that the objects of spatial covert attention are differentially encoded in phase and power of the alpha band using both decoding classifiers and encoding models. The authors found an immediate phasecoupling response between frontal and posterior regions during the early post-cue interval (0-400 ms) that was not spatially tuned, but conversely identified a sustained spatially-tuned response in alpha-power that began at roughly 450 ms post-cue. The authors interpreted these results to indicate that previously reported observations of increased anterior-posterior phase coupling during spatial covert attention reflects top-down attentional modulation of posterior sites. In contrast, they argue that the results rule out phase coupling as a mechanism for encoding the contents of spatial covert attention. In contrast to phase coupling, the conic classifier enables investigation of more general interactions between channels. Therefore, we tested whether the same conclusions—namely a lack of spatial information in anterior-posterior interactions and an onset of spatial tuning at 450 ms (in our case significant decoding), would hold for the conic classification approach.

To distinguish between local and distributed features, we compared the full conic classifier to reduced versions that only utilized univariate (variance) or multivariate (correlation) information (see Materials and Methods). The variance-only version of the conic classifier results in a weighted sum of univariate classification rules similar to signal power while having no sensitivity to information such as signal phase. In contrast, the correlation-only version ensures that the classifier can only consider the relationship between channels but not channel-specific information such as power. We then compared 6-way classification accuracy (all 6 locations at once) for the three conic classifiers (full/covariance, variance, and correlation) with the results recently reported by Samaha and colleagues ([24]).

Compared to the previously reported results by Samaha and colleagues with this dataset ([24]), we found that the full and reduced conic classifiers performed favorably when applied to the most posterior channels. The full conic classifier achieved a mean accuracy of $32.1 \pm 7\%$ which was not only significantly better than chance levels, but also significantly better than the classification performance achieved by Samaha and colleagues $23.1 \pm 4\%$ based upon alpha-power (Table 1). The reduced conic classifiers corresponding to just correlation or just variance also demonstrated significant classification ability, relative to chance accuracy (Table 3). Their performance was worse than that of the full conic classifier and the two versions were equivalent to each other in classification accuracy as well as to the alpha-power classifier used by Samaha and colleagues (Table 3). Therefore, significant information is present in both single-channel measures (variance/power) as well as in cross-channel measures (correlation) for posterior channels. Moreover the combination of these variables (covariance) in the full cone provides far more information than either the correlation or variance alone. We found no additional information in the anterior channels above what was present in posterior channels alone (see Materials and Methods) replicating

the findings of Samaha and colleagues ([24]) that anterior-posterior phase coupling is not spatially tuned.

Next, we tested whether the conic classifier produces a similar temporal profile of covert spatial attention as the inverted encoding models utilized by Samaha and colleagues ([24]). Their reported results indicated that significant channel tuning for alpha-power to the attended location began at roughly 450 ms post-cue. In contrast, the conic classifier identifies significant and sustained information regarding the attended location at a substantially earlier post-cue time period. Significance over the time course was assessed by permutation testing $n = 350,000$ (see Materials and Methods). For the desired levels of significance thresholds did not vary significantly over time or classification procedure (all $\sigma^2 < .005\%$) so we plot them as static. Mean classification accuracy for the full conic method exceeded chance ($p < .05$) starting at roughly 310 ms post-cue (Fig. 6 A), thus preceding the identified onset of alpha-band tuning ([24]) by nearly 150 ms. We note that this initial period was not included in the classifier training and significant decoding is attained from that point onwards possibly reflecting similarly sustained information. Thus, the decoding pattern meets the temporal persistency required of sustained attentional markers. Still, one possibility for the earlier onset of significant decoding is that the discrepant interval reflects cue-driven responses. However, in the very early trial interval (before 250 ms) classification accuracy was consistently at or below chance. We therefore consider it unlikely that residual activity from visual responses drives the earlier onset of conic decoding (see Materials and Methods). Results from central and parietal electrodes suggest that the earlier decoding ability stems from posteriorly generated neural signals (see Materials and Methods). Lastly, we considered whether the earlier onset of significant decoding with the conic classifier could be attributed to any preprocessing steps, yet this seems unlikely. Our only preprocessing step consisted of a conservative wavelet smoothing (see Materials and Methods). In particular, the kernel we employed is far shorter than the temporal discrepancy between methods (cones vs. alpha-power tuning) and wavelet convolution moves information forward in time so any biases would actually diminish the effect.

In summary, the conic decoding method is able to identify the contents of covert spatial attention from approximately 310 ms post-cue from posterior sites, roughly 150 ms prior to inverted encoding models ([24]; Fig. 6A). Follow-up analyses suggest that this discrepancy is unlikely to be caused by visual responses to the cue, transient cue-related responses in neighboring regions, or preprocessing steps (see Materials and Methods). Moreover, the time-course of the deconstructed cone (Fig. 6 B) indicated that this initial period of significant decoding is more attributed to local changes (the variance component) than distributed changes (the correlation component; see Materials and Methods). Moreover, in this period and the surrounding ± 100 ms were marked by a general anti-correlation between decoding accuracy based upon variance vs. correlation (see Materials and Methods; Fig. 6B). To be clear, we do not claim that this result alone is sufficient to support an alternative account of the timecourse of spatial attention, but we use it to illustrate how the proposed method may contribute to neuroscience investigations through its improved temporal sensitivity and flexibility (i.e. the ability to construct reduced models that enable selective probing of univariate versus multivariate effects).

3.5. Invariant Properties Promote Robust Classification

The proposed approach yields a number of fruitful properties that minimize the contaminating effects of variance due to both ‘noise’/artifact and within-subject variance. We conducted a number of simulations using the EEG dataset to demonstrate that the conic classifier is robust to three sorts of noise: (slow) scalar multiplication (Fig. 7 A), temporal warping (Fig. 7 B), and multivariate box noise (Fig. 7C) (see Materials and Methods). The first two relate to the scale-invariant properties of cones. The first case (Fig. 7 A) results from simply differentiating the product of the original time series $x_k(t)$ and a scalar function $f(t)$. When f is slow, we have $d/dt[f(t)x_k(t)] = f(t)\dot{x}_k(t) + \dot{f}(t)x_k(t) \approx f(t)\dot{x}_k(t)$. Due to the scale invariance of cones, $\dot{x}_k(t) \in \text{Int } C(P) \Leftrightarrow f(t)\dot{x}_k(t) \in \text{Int } C(P)$ when $\dot{f}(t) = 0$. This case is especially relevant for multiplicative noise in which a multivariate signal is multiplied by an unknown scalar function, as could be the case for amplifier noise or changes in reference electrode conductance for EEG recordings. However, while the scaling will not affect how datapoints are labeled for a fixed classifier, it is possible that bias in the scaling function’s magnitude can affect the calculated covariance, hence the corresponding classifier. Normalizing individual data points of $\dot{x}(t)$ can remove this effect, although this procedure also can change the calculated covariance.

When sampling rates are high so that the time-series are relatively smooth, the conic method is also invariant to temporal scaling (Fig. 7B). More formally, we define a class invariant interval as a continuous interval $\tau := [T_1, T_2] \subset \mathbb{R}^+$ with $t_1, t_2 \in \tau \Rightarrow \dot{x}(t_1) = \dot{x}(t_2)$. Let $y(t \in \tau) = x(f(t))$ for any strictly positive monotone $f: \tau \rightarrow \tau$ in which τ is a nonnegative class-invariant section of the time series. Then $\dot{x}(f(t)) \in \text{Int } C(P) \Leftrightarrow \dot{y}(t) \in \text{Int } C(P)$. This case follows immediately by applying the chain rule and noticing that strict monotonicity implies $\dot{f}(t) \neq 0$. As the conic method is invariant to scalar multiplication, so too is it invariant to temporal “stretching” which corresponds to a conserved order of events despite irregularity in the precise timing. This problem of temporal “stretching” is referred to as dynamic time warping (DTW), and has been discussed by many previous authors, including in the use of DTW on derivative time series for classification ([36],[37],[38]). In general, dynamic time warping algorithms give a similarity measure between time series evolving at different rates by comparing the order of events between time-series rather than on a point-by-point basis. The recently proposed derivative-DTW variants operate similarly, but on the time series of first and second derivatives. Thus, both DTW and the currently proposed method are able to perform classification on warped time series. Unlike DTW, however, our approach does not explicitly consider the order of events, but rather operates on the set of realizable changes (the derivative geometry) and is thus applicable to systems in which different initial conditions also produce different orders of events. For example, in the Lorenz attractor ([31]), different initial conditions not only affect how long a trajectory stays within one of the attractor’s two “loops”, but also how many cycles it completes before exiting. The resultant time series from very similar initial conditions within the same system thus produce time series whose order of events differ greatly, despite generating identical state-space geometries after sufficient time. As with the previous case of multiplication by a scalar function, temporal warping will not affect how data points are classified for a fixed cone, but

it can affect the calculated covariance, hence the classifier. As before, this effect can be removed by normalizing individual data points of the derivative time series.

The method is also resistant to additive multivariate box noise (Fig. 7C). This follows simply from the fact that the box function has derivative zero everywhere save at a finite number of points (corresponding to steps) and thus has minimal influence on the derivative time series. When the sampling rate is sufficiently high, many slow signals such as those induced by motion artifact may be approximated via a series of box functions, making conic decoding robust to artifactual noise. Thus the filtering of slow signal by derivatives combined with the scale-invariant properties of cones make the classifier robust to various sources of artifact and within-subject variability. Lastly, we will show that in addition to robustness to noise, our conic method is especially tuned towards task-related signals as opposed to contaminant spontaneous activity.

3.6. Sensitivity to fast-slow interactions

In general, the current literature suggests that very slow changes in the EEG amplitude are particularly prone to drift and motion artifact (i.e. [39]). In contrast, increasing evidence suggests that fast neural activity is associated with task performance ([40]). However, this does not necessarily indicate that slow components are always artifactual, but simply that fast components, when present, may be more informative. There are also many cases in which neural activity manifests meaningful interactions between fast and slow signals ([41], [42], [43]). Thus, a method is needed that 1) can analyze slow signals when fast signals are weak/absent, 2) penalizes slow signals in the presence of faster ones, and 3) prioritizes slow components that exhibit slow-fast coupling over those that don't. The derivative covariance matrix exhibits all of these properties, especially when the components have similar variance (related to spectral power). Formally, consider two multivariate processes:

$X(t), Y(t) : \mathbb{R}^+ \rightarrow \mathbb{R}^n$ and class labels $\dot{x}(t) \in \{k\}$. As before, we will denote the class-dependent derivative covariances as $\Sigma_{\dot{X}_k}$ and $\Sigma_{\dot{Y}_k}$. Now consider a second, slower process $B(t)$

and a constant $0 < c < 1$ defined such that for any class-invariant interval τ (as defined above) and non-negative scalars $t, \epsilon : [t, t + \epsilon] \subseteq \tau$ we have $B(t + c\epsilon) - B(t) = Y(t + \epsilon) - Y(t)$. In other words, $B(t)$ evolves c times the speed of $Y(t)$. The covariance matrix for the combined signal $S(t) := X(t) + B(t)$ is:

$$\Sigma_{\dot{S}_k} = \Sigma_{\dot{X}_k} + c^2 \Sigma_{\dot{Y}_k} + c(\text{cov}(\dot{X}_k, \dot{Y}_k) + \text{cov}(\dot{Y}_k, \dot{X}_k)) \quad (6)$$

The notation $\text{cov}(A, B)$ indicates the covariance between two multivariate processes, whereas we have previously indicated single-process covariance Σ_A as shorthand for $\text{cov}(A, A)$.

Clearly $c^2 < c$ and if $\dot{X}_k, \dot{Y}_k$ are independent, $\text{cov}(\dot{X}_k, \dot{Y}_k) = 0$. When $\dot{X}_k$ is sufficiently small

(with the frequency profile held constant), the classifier's inner term $\Sigma_{\dot{S}_k}^{-1} / \det |\Sigma_{\dot{S}_k}^{-1}|^{1/n}$

approaches $\Sigma_{\dot{Y}_k}^{-1} / \det |\Sigma_{\dot{Y}_k}^{-1}|^{1/n}$ which satisfies the first condition: sensitivity to slow

components in the absence of fast ones. Likewise as c becomes small (for $\Sigma_{\dot{X}_k} \neq 0$), the classifier's inner term is biased towards the higher frequency X term even if the components have similar amplitude thus satisfying our second condition. Lastly, when X, Y have amplitude, the classifier is more sensitive to the interaction between fast and slow components via $cov(\dot{X}_k, \dot{Y}_k)$ than those that are independent of X (hence restricted to $\Sigma_{\dot{Y}_k}$ as $c^2 < c$). Thus our classifier satisfies all three conditions to balance the influence of fast/slow signal components in a manner that emphasizes task-related components. If all three scenarios are mixed throughout a single class, direct calculation of the covariance based upon $\dot{S}_k(t)$ will be biased in favor of periods featuring fast components due to their greater derivative magnitude. Again, this bias can be removed by normalizing individual data-points of the derivative time series (using any norm). This step will change the calculated covariance for each class, but will not affect the data's classifiability as $\dot{x}(t) \in \text{Int}(C(P)) \Leftrightarrow \alpha \dot{x}(t) \in \text{Int}(C(P))$ for any nonnegative α . The classifier's performance regarding the three criteria is exemplified by the continued sensitivity in the bursting network example (Fig. 2) which featured fast-slow coupling of downstream/upstream dynamics and within-cell fast-slow coupling during burst activity ([32]).

4. Discussion

We have introduced a novel method for performing classification of spatiotemporal signals. The key innovative feature of the approach is the use of the derivative rather than the original time-series as the basis for classification. By using the derivative time-series, our conic approach to classification is sensitive to temporal information, even though the classification method does not explicitly consider time as a variable. After computing the derivative time series, the method proceeds identically to static classification methods, with each time point of the derivative time series corresponding to one sample. From a dynamical-perspective the conic approach emphasizes decoding based upon the vector field of a system rather than the time course of trajectories.

We have given one example of an explicit form for conic classification which bears some similarity to a special case of quadratic discriminant analysis. However, unlike conventional quadratic discriminants, the balancing properties of derivatives and scale-invariant property of cones enable a natural explicit form for the n -class boundaries. These properties also lead to a number of invariance properties that cancel certain forms of noise such as scalar-multiplicative noise (i.e. amplifier noise), multivariate box-wave noise and low frequency noise (as elaborated on below). The form we provide relies exclusively upon the covariance matrices for each class's derivative time series and is essentially a comparison of the degree to which their derivative covariance matrices differ.

4.1. Continuous-Time Classification Reveals the Evolution of Neural States

As illustrated in Figure 5 the conic approach not only provides an accurate prediction for the trial, but also generates a prediction for each time point. Together, these predictions form a time series describing the evolution of the system between states with high temporal resolution. The key step of our conic classification technique is to capture the dynamic

properties of states by converting to the derivative time series. Previous dynamic approaches have generated state time-series using parameter estimates at fixed or sliding windows ([16], [17],[44],[45]). In contrast, the current approach generates a time series of individual predictions from a single classifier. Users have additional flexibility in choosing the desired resolution via smoothing etc., or application of secondary classification techniques to the cone-generated time series for higher-order analyses.

4.2. Extension to hemodynamic signals

Presently, we have only presented results of conic decoding for data with relatively high sampling rates, namely neuronal simulations and EEG. However, the potential exists for applications involving much slower modalities provided the events of interest are observable on a similarly long timescale. Task-related dynamics have long been observed in the BOLD signal (e.g. [46]) and more recent studies have also found dynamic relations between regions (e.g. [47],[48],[49]). Although the low sampling rate of fMRI is nonideal for derivative estimation, the aforementioned work suggests BOLD dynamics may be observable over longer intervals. The temporal derivative of interpolated data might then prove a suitable, smoother proxy. As such, the proposed method of conic classification has potential application to fMRI studies of temporally extended cognitive states.

4.3. The Conic Method Extends Classification to Dynamic Features

The special relevance of our proposed derivative-based classification technique to neuroscience is that it specifically decodes dynamic features (versus spatial patterns). The mechanism of the classifier is to decode how a system evolves across time, which we believe is particularly salient to neuroscience. Indeed, an influential view of brain function is to view the brain as a dynamical system. Consequently, the most productive approach would be to map cognitive constructs not only onto regions of the brain, but also the “patterns” of how those regions interact in time. For example, a primary interest of the cognitive neuroscience community is to determine the mapping from characteristic waveforms (event-related potentials or spectral components) to cognitive states. Classical characterizations have been used to inform dynamic models, which then predict anatomical regions of interest. Indeed, several modeling and empirical results now concern how specific waveform features, such as frequency, are formed. However, primary results derived from these characterizations still ultimately link cognitive states to specific mathematical properties of signals (e.g. patterns of frequency decomposition). For example, previous analyses of the covert spatial orienting dataset ([14], [24]) relate band-limited power (in the alpha band) to covert attention. In contrast, our method does not rely on deriving secondary measures from data, but rather examines dynamics in their most primitive form: the temporal derivatives of the signals themselves.

5. Materials and Methods

5.1. Derivation of the Conic Classification Criteria (2)

To derive our equation for conic classification we assume that the distributions for the two classes are elliptical in the sense that they may be described in terms of $f_i \sim g_i(x^T \Sigma_i^{-1} x)$.

Unlike the traditional definition of an elliptic function such as those described in standard quadratic discriminant analysis we do not assume that the function f is non-increasing. In fact, this must not be the case from the perspective of autonomous dynamical systems as the origin/'mean' in derivative space corresponds to a fixed point and should thus have low density in the case of neural signals which are highly transient. Thus we assume nothing regarding the distributions f_i save that they are non-degenerate. Instead of restricting the distribution, we restrict our set of classifiers to be scale-invariant (conic) and thus we may instead consider the distribution projections on the surface of the $(n-1)$ -dimensional sphere (S^{n-1}) without loss of generality. What follows involves deriving the probability density of an elliptic distribution onto $\partial\mathcal{S}^{n-1}$. These densities immediately lead to the derivation of a maximum likelihood classifier of the form (2).

We impose zero density at the origin as it has no direction and hence is never classifiable. From a practical standpoint this assumption should be insignificant for continuous valued data with a finite number of data-points. Our treatment of the general elliptic case is in essence the same as the Gaussian case with mean zero long-considered in directional statistics [50]. We then apply simple monotone transformations to the distribution to ease computation in high dimensional settings.

Proposition: Consider a set of elliptical probability distributions $\{f_i\}_{i=1\dots m}$ over $\mathbb{R}^n$ with all means equal zero and $f_i(0)=0 \forall i \in \{1\dots m\}$. Denote the corresponding covariance matrices for each elliptical distribution as $\{\Sigma_i\}_{i=1\dots m}$ and a random variable X valued over $\mathbb{R}^n$. Define the projection operator $\Pi_S: \mathbb{R}^n \setminus \{0\} \rightarrow \partial\mathcal{S}^{n-1}$ and its derived distributions (h_i) as follows:

$$\Pi_S(x) = \frac{x}{\|x\|_2} \quad X \sim f_i \Rightarrow \Pi_S(X) \sim h_i \quad (7)$$

Then the following two assignment functions are equivalent:

$$\arg \max_i h_i(\Pi_S(X)) = \arg \min_i \frac{X^T \Sigma_i^{-1} X}{\det |\Sigma_i^{-1}|^{1/n}} \quad (8)$$

Proof. Consider the elliptic distribution f_i and covariance matrix Σ_i . Thus, by the definition of an elliptic function, there is a function $g_i: \mathbb{R} \rightarrow \mathbb{R}$ and a functional $\kappa: L^1 \rightarrow \mathbb{R}$ such that:

$$f_i(X) := \kappa(g_i) \det |\Sigma_i|^{1/2} |g_i(X^T \Sigma_i^{-1} X)| \quad (9)$$

Here K_{g_i} denotes the scaling function which depends upon g_i but not Σ_i .

The distribution for $\Pi_S(x)$ is the projection of f_i onto $\partial\mathcal{S}^{n-1}$. In order to perform this radial projection we convert to spherical coordinates $r(x) \in \mathbb{R}^+$ and $\phi(x) \in \mathbb{R}^{n-1}$ with radial variable $r(x) := \|x\|_2$ and angular vector $\phi(x)$ (for any full rank angular basis). We define the function Θ_i to represent the portion of the quadratic term, which depends upon the orientation of x (i.e., $\phi(x)$), but not its magnitude:

$$\Theta_i(\phi(x)) := (\Pi_S(x))^T \Sigma_i^{-1} (\Pi_S(x)) \quad (10)$$

The distribution of $\Pi_S(x) \in \partial\mathcal{S}^{n-1}$ is found by integrating Θ_i along the radial component of x ; $r := \|x\|_2$:

$$f_i(\Pi_S(x)) = \kappa(g_i) \det |\Sigma_i^{1/2}| \int_{r=0}^{\infty} r^{n-1} g_i(r^2 \Theta_i(\phi(x))) dr \wedge d\Omega \quad (11)$$

In which $d\Omega$ is the volume element of the $(n-1)$ -spherical shell. Using a simple change of variables ($u = r^2 \Theta(\phi)$) we remove the angular component $\Theta(\phi)$:

$$f_i(\Pi_S(x)) = \kappa(g_i) \frac{\det |\Sigma_i^{1/2}|}{2\Theta(\phi)} \int_{u=0}^{\infty} g_i(u) \left[\frac{u}{\Theta(\phi)} \right]^{\frac{n-1}{2}} du \wedge d\Omega \quad (12)$$

$$= \kappa(g_i) \frac{\det |\Sigma_i^{1/2}|}{2\Theta(\phi)^{n/2}} \int_{u=0}^{\infty} g_i(u) u^{\frac{n-1}{2}} du \wedge d\Omega \quad (13)$$

We will refer to the remaining integral as $\xi(g_i)$. Now consider a second distribution f_0 characterized by a separate scalar function g_0 g_i but identical covariance: $\Sigma_0 = \Sigma_i$, hence $\Theta_0(\phi(x)) = \Theta_i(\phi(x))$. The distributions may then be written:

$$f_i(\Pi_S(x)) = \kappa(g_i) \frac{\det |\Sigma_i^{1/2}|}{2\Theta_i(\phi)^{n/2}} \xi(g_i) \quad (14)$$

$$f_0(\Pi_S(x)) = \kappa(g_0) \frac{\det |\Sigma_i^{1/2}|}{2\Theta_i(\phi)^{n/2}} \xi(g_0) \quad (15)$$

As both functions are probability distributions they must sum to 1. As the distributions are defined over $\partial\mathcal{S}^{n-1}$ they are integrated over $d\Omega$. Factoring out the scalar terms produces integrals dependent only upon the angular component $\Theta(\phi)$.

$$\xi(g_0)\kappa(g_0)\frac{\det|\Sigma_i^{1/2}|}{2}\int_{\partial\mathcal{S}^{n-1}}\Theta_i(\phi)^{-\frac{n}{2}}d\Omega \quad (16)$$

$$= \xi(g_i)\kappa(g_i)\frac{\det|\Sigma_i^{1/2}|}{2}\int_{\partial\mathcal{S}^{n-1}}\Theta_i(\phi)^{-\frac{n}{2}}d\Omega = 1 \quad (17)$$

Thus there must be a constant η_n related to the geometry of the $n-1$ -sphere, such that $\xi(g_j)\kappa(g_j) = \eta_n$ for any appropriate function g_j . Dividing f_i by η_n and raising the result to the $-2/n$ completes the proof:

$$\arg \max_i f_i(X) = \arg \min_i \left((f_i(X) / \eta_n)^{-\frac{2}{n}} \right) = \arg \min_i \frac{X^T \Sigma_i^{-1} X}{\det|\Sigma_i^{-1}|^{1/n}} \quad (18)$$

which amounts to our classifier (2). $\square$

To recover a full maximum likelihood classification which considers initial population proportions (p_i) we simply multiply the derived probability density proportion of the probability density $f_i(X)$ by p_i before raising solutions to $-2/n$ hence:

$$\arg \max_i f_i(X)p_i = \arg \min_i \left((f_i(X)p_i / \eta_n)^{-\frac{2}{n}} \right) = \arg \min_i \frac{X^T \Sigma_i^{-1} X}{(p_i^2 \det|\Sigma_i^{-1}|)^{1/n}} \quad (19)$$

5.2. Computational Considerations and Limitations

A key computational step is calculation of the determinant. To avoid numerical instability associated with matrix inversion, we calculate the determinants via:

$$\det|M|^{\frac{1}{n}} = \prod_i^n (\sigma_i(M))^{\frac{1}{n}} \quad (20)$$

where $\sigma_{1\dots n}$ indicate the singular values of M . Nevertheless, retrieving the eigenvalues may be computationally expensive ($O(n^{k>2})$) as compared with linear classification algorithms and this remains a limitation of the proposed method. However, we have used the default MATLAB eigenvalue algorithm for conic classification of 2,000 dimensional data on a midrange laptop (x64, dual Intel(R) Core(TM) i7-6500U @ 2.50GHz).

Another limitation of the current method concerns when the covariance matrices are ill-conditioned. In some potential applications (e.g. fMRI) the dimension of the data (number of

channels) may exceed the number of datapoints per class. The simplest approach to this issue is dimensionality reduction or to (artificially) increase the number of datapoints through interpolation. However, an alternative method is to retain the original number of dimensions and use sparse-estimation for the population inverse matrix. Previous authors (e.g.[51]) have developed computationally efficient methods to perform this calculation. Many of these approaches rely upon LASSO or regularization methods so their accuracy may rely upon at least some channels being independent.

In practice, some eigenvalues calculated for the covariance matrices may be near zero due to dependency. To maintain stable estimates of the classifier we recommend setting these eigenvalues equal to zero in the pseudoinversion and removing them from the calculation of determinants. However, we did not need to threshold eigenvalues for any of the simulated classifications. For the empirical data we used a relative eigenvalue threshold of 10,000.

5.3. Simulation 1

In Figure 1 we used a sigmoidal 3-neuron recurrent network to illustrate conic dynamics in a low dimensional space. The connection scheme is depicted schematically in Fig 1A,B and features two coupled negative feedback loops. A common hub (neuron 3) inhibits excitatory cells 1,2. Connections between neurons 2,3 were stronger than those with neuron 1. The model was as follows:

$$\frac{dx}{dt} = W(1 + \tanh(s \circ (x - b))) - D \circ x + c \quad (21)$$

We use the notation $\circ$ to indicate the Hadamard product (element-wise multiplication). We simulated this network twice using the same slope (s) and connection weights (W) for each case:

$$W = \begin{bmatrix} 0 & 0 & -3 \\ 0 & 0 & -25 \\ 3 & 25 & 0 \end{bmatrix} \quad s = \begin{bmatrix} 1.25 \\ 2.5 \\ 3 \end{bmatrix} \quad (22)$$

The two cases differed, however, in the values for decay (D), baseline (c), and threshold (b):

$$\text{FirstSim: } c = \begin{bmatrix} .7 \\ .584 \\ -2.407 \end{bmatrix}, \quad b = \begin{bmatrix} 7.909 \\ .597 \\ 3.374 \end{bmatrix}, \quad D = \begin{bmatrix} .11 \\ .14 \\ .13 \end{bmatrix} \quad (23)$$

$$\text{SecondSim: } c = \begin{bmatrix} 3.04 \\ 25.3 \\ -36.26 \end{bmatrix}, \quad b = \begin{bmatrix} 5.37 \\ -.194 \\ -.04 \end{bmatrix}, \quad D = \begin{bmatrix} 5.5 \\ 7 \\ 6.5 \end{bmatrix} \quad (24)$$

Simulations were performed with Euler integration $dt = .0005$ and a run time of $t = 3$ for each iteration and then down-sampled by a factor of 15. For each case 150 iterations were performed with starting positions selected pseudo-randomly. Segments of each time-series were also pseudorandomly selected using a random box-function ($B(t)$). The box function was formed by choosing 200 box-lengths distributed according to $(1 + r)^2$ with r distributed according to the standard normal distribution ($r \sim N(0,1)$). The resulting box-lengths were then rescaled so that they summed to the total window size ($t=3$). The amplitude for each box was generated from a standard normal distribution and the inclusion rule for displaying a point $x(t)$ was $B(t) > 1.5$. Thus, roughly 6.9% of the total points were displayed for each trajectory/initial condition.

All simulations were carried out in MATLAB2016b.

5.4. Simulation 2

For this simulation we considered a network of 4 bursting Hindmarsh-Rose neurons with single-exponential synapses. A chaotic Lorenz attractor provides input to the network as a representation of upstream network activity. The Lorenz attractor was given a standard parameterization with the addition of a Gaussian white-noise term (η_i) with standard deviation ϵ ("Internal Noise") :

$$\frac{dx}{dt} = 10(y - x) + \epsilon\eta_1(t) \quad (25)$$

$$\frac{dy}{dt} = x(28 - z) - y + \epsilon\eta_2(t) \quad (26)$$

$$\frac{dz}{dt} = xy - 8/3z + \epsilon\eta_3(t) \quad (27)$$

The Lorenz attractor was Euler integrated with $dt=.0025$ for duration $t=50$. Each time point was then upsampled (repeated) by a factor of 100 to ensure that the system did not fluctuate too quickly. The upstream system was considered 'on' when the first Lorenz variable (x) was greater than 0. Initial conditions for each iteration were normally distributed with mean $[-10 -10 27]$ and standard deviation .25.

The output current from the Lorenz attractor was a binary function of the first Lorenz variable's sign:

$$I_{Lorenz} = \kappa L(t) \begin{bmatrix} -.18 \\ .24 \\ .06 \\ 1.56 \end{bmatrix}, \quad L(t) = \begin{cases} 1 & x(t) > 0 \\ 0 & x(t) \leq 0 \end{cases} \quad (28)$$

The term κ is a scaling factor (“Stimulus Strength”) which we manipulated. Each neuron’s internal dynamics were governed by the standard Hindmarsh-Rose equations [32]:

$$\frac{du}{dt} = v - u^3 + 3u^2 + .5 + I(t) + \epsilon\gamma(t) \quad (29)$$

$$\frac{dv}{dt} = 1 - 5u^2 - v \quad (30)$$

$$\frac{dw}{dt} = .001(4(u + 1.6) - w) \quad (31)$$

Here $I(t)$ denotes the combined input to each cell and $\gamma(t)$ denotes a Gaussian white-noise process with standard deviation ϵ (“Internal Noise”). Input consisted of a constant baseline level of input I_{base} , the synaptic currents I_{Syn} and the currents from the Lorenz attractor I_{Lorenz} . The synaptic currents were the product of a synaptic weight and a single dynamic synapse variable per neuron.

$$I_{Syn}(t) = Wg(t) \quad W = \begin{bmatrix} 0 & 8 & -4 & -16.8 \\ 12 & 0 & 0 & -14 \\ 0 & 7 & 0 & -16 \\ 16 & 0 & -20 & 0 \end{bmatrix} \quad (32)$$

$$\frac{dg}{dt} = -5g + .5(1 + \tanh(4u - 2)) \quad (33)$$

$$I_{base} = \begin{bmatrix} 1.8 \\ 1.5 \\ 1.9 \\ -1.6 \end{bmatrix} \quad (34)$$

The Hindmarsh-Rose system was Euler integrated with step size .05 and end time $t=5,000$. Initial conditions for all neuron variables (u, v, w) were drawn from a standard normal Gaussian distribution. All synapses were given initial condition zero. The initialization period from $t=0$ to $t=1,000$ was removed to allow the system time to stabilize. The variables of interest consisted of the voltages (u) for the first three neurons only, which were down-sampled by a factor of 10. Gaussian noise was added with standard deviation ξ ("Measurement Noise"). We performed blind classification with clustering to assign system states (Lorenz input on/off) based upon the first three neurons' voltages. We performed classification using the covariance of the normalized derivative time series ($\dot{x}_t / \| \dot{x}_t \|_2$) for each time bin as input to the k-means algorithm. Clustering for covariances was performed using k-means for 2 clusters for the upper triangle of the covariance (six elements). Similar clustering for spectral analyses were based upon 2-cluster k-means for the vector composed of the separate real and imaginary components of the spectrum evaluated at 6 equi-spaced points spanning half the smallest window size. The conic decoding was then performed using the two centroids derived from the k-means clustering as the covariance matrix for each class. The prediction weights for all points within a window were averaged to assign a class to the full window. For spectral clustering, labels were assigned directly based upon the centroid cluster to which each bin belonged. As the analysis involved blind clustering, cluster labels were arbitrary. In assigning cluster numbers to system states (i.e. does cluster 1 or 2 indicate input 'on') we used the assignment that provided maximal accuracy for each simulated trial. As such, the minimum accuracy for each trial is .5 and the expected accuracy is greater than .5 using this assignment. Therefore, we derive the null distribution for this case to evaluate bias. As the accuracy for a random assignment (random mapping between cluster number and 'on'/'off' state) is uniform $U[0,1]$, their mean is Bates distributed. For larger numbers of samples (as was the case for each trial), the Bates distribution for mean .5 over $[0,1]$ (as there are two classes) limits to the Gaussian $\mathcal{N}(.5, 1/(12n))$. Thus the chance accuracy for random label assignment limits to $\mathcal{N}(.5, 1/(12n))$. For our biased case (using the assignment that provides greatest accuracy), the assignment function is $\max(\mu(A), 1 - \mu(A))$ with A the accuracy obtained by arbitrary assignment. For several hundred samples, as was our case for each trial, the expected values of this distribution is essentially .5 (limiting to $.5 + \frac{E[|N(0,1)|]}{\sqrt{12n}}$).

To analyze the effect of parameterization we manipulated each parameter (intrinsic/measurement noise, stimulus strength, bin length) independently while keeping all others at their median value. There were a total of 136 trials per condition (parameter level) and 9 equi-spaced conditions per variable. As demonstrated in Figure 2D, long run times ensured that despite the obvious noise/chaos within each trial (Fig. 2B) standard deviations for decoding accuracy were relatively small. Combined with the large sample size, these factors generated highly significant results for nearly every statistical comparison. As such we focused upon displaying performance. Note that the shading in Figure 2D indicates standard deviation as opposed to standard error (which would not be visible).

Experiment 3

The data and task were obtained from an open repository associated with ([34],[14]). Briefly, during the task, a multicolored hexagonal cue was presented for 200 ms and participants were instructed to focus upon the edge corresponding to a given color. The orientation of this edge predicted in which of 6 radially located discs the upcoming target would be displayed after a delay of 500–2000 ms. In 80% of trials the target appeared in the cued location. Targets were presented for 200 ms after which participants had to indicate the target's shape (either '+' or 'x'). The dataset from these studies is publicly available (BNCI-Horizon-2020.eu) and consists of eight subjects' pre-processed recordings for 60 EEG channels, 2 EOG channels. Data was recorded at 1 kHz but the published dataset has been down-sampled to 200 Hz. Our only further modification to the data was a single wavelet reconstruction to estimate derivatives given noisy data. For this step, we simply used a 1-D Daubechies 7-tap decomposition and level 3 reconstruction, although more specialized methods exist ([52]). Statistical analyses between previous and current results are based upon paired t-tests for each case. The significance of pair-wise accuracy was assessed with 1-sample t-tests for accuracy greater than chance (50%) and Bonferoni corrections which mimicked the original analyses with this dataset ([14]). For the 15 possible pairwise-comparisons this approach typically resulted in corrected significance thresholds of roughly 65% for $p < .05$, although the precise threshold varied with the number of trials (4).

5.4.1. Visualizing Conic Projections—Plots in Figure 5 display conic classification of subject 4's data for the lower upper right vs. lower left contrast for the 21st, 24th, 27th, and 30th trial of each of the two locations. Time series in Figure 5B are the concatenated time series of the largest eigencomponent for each location across the four classification periods each. We chose subject 4 to illustrate conic projections as the median-performing subject (in terms of conic-decoding performance). For context, we also include identical plots for the best-performing (#3; 8) and worst-performing subjects (#8; 9).

5.4.2. Individualized Results for n-Class Decoding—In the main text we compare performance of the conic classifier with that of other methods on the same benchmark dataset. Namely, we presented pairwise classification accuracy for comparison to the original analyses by Treder and colleagues ([14]) and full 6-way classification for comparison to the more recent work by Samaha and colleagues ([24]). However, one of the main advantages of the current classifier is that it is inherently multi-class as opposed to most linear or even quadratic classifiers which are generally pairwise, necessitating adhoc methods such as voting schemes to perform multi-class decoding. To better illustrate the ability of the conic classifier on multi-class problems we therefore include results comparing the full range of possible comparisons (2-way through 6-way) for each individual subject (Fig. 10). Results are plotted for the mean performance across all possible combinations of n-locations.

5.5. Effects of Artifact Rejection and Eigenvalue Estimation

For results displayed in the main text we included all trials meeting the trial-length criteria without consideration of possible eyeblink and motion artifact. Moreover, we censored eigenvalues with relative magnitudes less than 1/10, 000 as calculation of very small

eigenvalues can become numerically unstable. To ensure that these choices did not substantially affect conclusions we present the results corresponding to Fig. 3 for the cases of either motion artifact removed based upon the exclusion criteria adopted by Samaha and colleagues ([24]) or with alternative conditions for eigenvalue censoring. We considered two alternative schemes for calculating the eigenvalues: either including all eigenvalues or, rather than removing small components, to remove the largest posterior component via Principal Component Analysis (PCA). PCA was performed upon the full data (all time points) in the native space (before derivative calculation/wavelet smoothing) and the largest principal component was removed. The eigenvalue and artifact conditions were fully crossed (Fig. 11) and we computed the number of significant pairs (Fig. 11A,B) and pairwise accuracies (Fig. 11C) for each subject. Group-level performance was very similar regardless of these preprocessing conditions (Fig. 11B). Likewise, individual performance as measured by decoding accuracy for each pair of locations was generally consistent across conditions (Fig. 11C). Classification accuracy for subject 8 slightly increased when removing artifact (averaged across eigenvalue conditions: $\text{paired-}t(14)=2.16$, $p=.0485$, $\mu = .027 \pm .0485$) while accuracy for subject slightly 6 decreased ($\text{paired-}t(14)=-4.635$, $p=.0004$, $\mu = -.021 \pm .0184$). We did not observe any significant changes due to eigenvalue conditions. These changes in pairwise accuracy had a corresponding increase/decrease in the number of significantly decoded pairs (Fig. 11A). In contrast, several high-performing subjects actually had the number of significantly classified pairs decrease when artifactual trials were removed (Fig. 11A). However, these decrements do not necessarily imply that performance worsened: since the significance threshold (in terms of accuracy) decreases with the number of classifications, removing trials due to artifact effectively increased the accuracy threshold for significance. Thus, although the changes in pair-wise decoding accuracy due to preprocessing conditions were minor for most subjects, artifact-based censoring can lead to fewer significantly classified pairs by decreasing the sample size. The main take-away from this analysis is therefore that artifact had relatively little impact upon decoding performance for most subjects (Fig. 11C) but removing artifact did increase performance for the worst-performing subject (# 8) and negated the previously observed spectral advantage in decoding for this subject.

5.5.1 . Decomposing the Classifier into Variance and Correlation—Just as spectral relationships may be decomposed into power and phase, the covariance matrix may be decomposed into contributions of variance and correlation to assess local (variance) vs. distributed (correlation) contributions to task-related activity. Simply taking the variance of each channel as input to a standard classifier does not remove multivariate information as the classifier will still determine boundaries based upon the covariance of observations (the covariance of variances). However, by computing the cone using only the diagonal (variance) terms all multivariate information is gone and the resultant classifier is actually identical to a weighted sum of univariate classification rules:

$$\arg \min_k \left(\sum_{t \in T_l} \left(\frac{\sum_{c=1}^n \dot{x}_c(t)^2 \sigma_{\dot{x}_{k_c}}^{-2}}{\prod_{c=1}^n \sigma_{\dot{x}_{k_c}}^{-2} / n} \right) \right) \quad (35)$$

where T_l denotes the l^{th} bin and $\sigma_{\dot{x}_{k_c}}^2$ denotes the variance of the derivative signal for class k

at channel c . To generate the classifier based upon correlation alone, we used the regular conic classifier ((2)) but z-scored the data channel-wise for every channel. Thus, for every trial, each transformed channel had a variance of one so that the covariance was identical to the correlation.

5.5.2. Testing the Influence of Anterior Channels—In addition to the 13 most posterior channels we also tested the contribution of anterior channels. Using the first 11 anterior channels (first two rows) produced decoding that was weak ($\mu = .201 \pm .0275$), but significantly above chance ($t(7) = 3.54, p = .0095, 2-tailed$) for the sixway classification. When we censored high motion trials using Samaha and colleague's ([24]) criteria the anterior classifier did not reach significance ($\mu = .1911 \pm .0296; t(7) = 2.33, p = .0525, 2-tailed$). Performance for the deconstructed conic classifiers (variance and correlation) was even weaker. These results confirm Samaha and colleague's ([24]) conclusion that activity in anterior channels does not reflect the attended location. Moreover, combining the very anterior and very posterior channels together did not improve classification accuracy above the posterior channels alone ($t(7) = -1.85, p = .106, 2-tailed$). In fact, the accuracy tended to decrease ($\mu = -.019 \pm .029$) indicating that the anterior channels were simply adding noise. Thus, results using the conic method agree with the previous findings that anterior-posterior coupling does contribute additional information regarding the attended location from what is found using the posterior channels alone.

5.5.3. Testing Conic Time Series—Conic time series were formed by taking the classification accuracy over moving windows of 6 data points (30 ms) to form predictions. Finer resolutions (i.e. single data points) are also possible, but we chose a (very small) sliding window instead to aid visualization. Accuracy is for leave-one-out cross-validation. Classifiers were trained on the interval 750-1900 ms post-cue and tested on the full trial interval. Thus the first 750 ms was not present in the training data for any trial. The resultant accuracies were averaged across class for each subject and smoothed with length 7 (35 ms), unit variance gaussian kernel. The kernel was scaled to have an area of exactly one so it didn't bias overall performance.

In order to test the significance of the conic time series we performed permutations based upon scrambled class labels (spatial location) with all else identical. The number of observations per scrambled class was the same as in the original data. We performed 10,000 permutations for each combination of subject x classifier. Using the individual subject distributions we then constructed the group-level distribution from the means of subject-wise permuted data repeated 350,000 times for each classifier. None of the resultant significance

thresholds displayed substantial variation with respect to classifier choice or time post-cue (all $\sigma^2 < .05\%$) so we set them constant at the mean value.

5.5.4. Ruling Out Visually-Evoked and Central/Parietal Influences to

Posterior Decoding—To control for the possibility that our earlier onset of significant posterior decoding was due to visually-evoked activity we considered the similarity between activity during the window of primary visual responses (i.e. the first 250 ms after stimulus onset=50 ms post-cue; [53]) and the delay period. This activity does not resemble the learned patterns during delay as reflected in insignificant decoding accuracy (Fig. 6A) so it is unlikely that the onset of significant posterior decoding is related to visually-evoked responses. Another possibility is that the earlier onset of significant decoding reflects later processing of the cue such as its semantic significance as opposed to reflecting spatial attention. We do not believe this explanation is likely either due to the spatial profile of decoding—namely a mismatch between central/parietal and posterior decoding time courses. Event-related activity associated with semantic processing has been previously described (in other experiments) over the relevant interval (300–500 ms post-cue) but is centered over central and parietal electrodes ([54],[55]) while we only considered occipetal and parieto-occipetal (the 13 most posterior). To test whether our early decoding onset reflected processes linked to the parietal lobes, we considered the parietal electrodes (an additional 11) both alone and in addition to posterior electrodes. In both cases classification accuracy was above chance. However, the addition of parietal electrodes did not augment the early decoding accuracy and parietal electrodes alone did not exhibit a sustained response starting at that interval. Moreover, combining posterior and parietal electrodes did not improve mean classification accuracy relative the posterior electrodes alone ($\mu = -.0298 \pm .0614$; $t(7) = -1.375, p = .212$). As parietal electrodes did not add additional information for decoding (as reflected in accuracy), posterior conic decoding is unlikely to be driven by parietal processes. However, parietal electrodes did reveal a different pattern of results with an even earlier, but weak and transient, component which reached significance ($p < .05$) beginning at 170 ms and lasted approximately 100 ms. This component was driven by the correlation of derivatives between parietal electrodes (as opposed to variance) and demonstrated a different time course than the decoding performance at posterior electrodes (namely a lack of sustained response). Thus, our earlier onset of significant posterior decoding cannot be explained by processes linked to parietal sites that tend to occur near the time interval of interest (i.e. semantic event-related activity [54],[55]).

5.5.5. Interpreting Local vs. Distributed Contributions to Decoding in EEG

Data—In general, dissociating the source of EEG signals is a nontrivial endeavor, although several methods have been developed to localize source signals (i.e. [56]). The main confound is volume conductance, by which changes in focal cortical activity can lead to changes in correlations between regions. For instance, scalp regions which have high conductance paths to the site of local activation will appear more correlated with each other. Therefore, simply identifying differences in decoding accuracy between variance and correlation is insufficient to suggest that the neural generators are more focal vs. distributed. However, in the present case we can separate these possibilities as the variance and correlation decoding time-courses are anti-correlated during this early interval. In contrast, a

focal neural generator will always have a positively correlated impact on variance and correlation as they both depend monotonically upon the generative signal's amplitude. We consider the early period of significant decoding to be from 300-500ms corresponding to the first 200ms after the start of significant decoding performance. Over this interval the group average time series for correlation-based and variance-based decoding accuracy are negatively correlated ($r(39) = -.324, p = .0108, 2\text{-tailed}$) and the same holds for individual subjects after detrending: (mean $r = -.22 \pm .28$ after detrending). Thus, the time courses of variance-based and correlation-based decoding accuracies during this early interval are inconsistent with a common focal neural source.

5.5.6. Generating Spatial Maps—All spatial maps displayed were generated using the group-averaged covariances of the derivative for each class. Data in Figure 4 is displayed from a posterior coronal view (as indicated in center). For each location we considered the cone generated by the contrast location “x” vs. the hemifield opposite “x” (the average covariance across 3 locations). The method in which we visualized the conic matrix is by conversion to a vector which forms a static head-map. As with functional connectivity, the more natural way to visualize the matrix is as a weighted graph between nodes, however, we chose to use a headmap to better visualize the evident retinotopic organization. To do this, we defined each channels contribution to the conic matrix as the sum of its squared contributions to each eigenvector v_j weighted by the corresponding eigenvalue λ_j . Thus the i^{th} channel's weight (w_i) was:

$$w_i = \sum_{j=1}^n \lambda_j [v_j]_i^2 \quad (36)$$

Smooth maps were created by interpolating the channel-wise measurements after coronal projection. Linear triangulation-based interpolation was performed with the built-in MATLAB function ‘griddata’.

5.5.7. Eigenlengths and Eigenvectors—In the two class case (labels = ± 1), the conic boundary corresponds to $\partial C(P) := \{\dot{x} \mid \dot{x}^T P \dot{x} = 0\}$ with P the weighted inverse of covariance matrices:

$$P := \frac{\Sigma_{(\dot{x}_{[-1]})}^{-1}}{\det |\Sigma_{(\dot{x}_{[-1]})}^{-1}|^{1/n}} - \frac{\Sigma_{(\dot{x}_{[+1]})}^{-1}}{\det |\Sigma_{(\dot{x}_{[+1]})}^{-1}|^{1/n}} \quad (37)$$

The class assignment function is $i(t) = \text{sgn}(\dot{x}(t)^T P \dot{x}(t))$. As P is real symmetric, it has a spectral decomposition UDU^H . Splitting the diagonal matrix of eigenvalues (D) into positive and negative components: $D = D^+ - D^-$ produces:

$$\dot{x}^T P \dot{x} = \|\sqrt{D^+}(U^H \dot{x})\|_2^2 - \|\sqrt{D^-}(U^H \dot{x})\|_2^2 \quad (38)$$

As the right side of Equation 38 involves the difference of two norms, we refer to them as the positive (with D^+) and negative (with D^-) eigenlengths. These eigenlengths are useful for generating a 2-dimensional projection of the data by which to compare separability as in Figure 5C. Similarly we generate "eigencomponents" of the system by simply rotating it into the cone's coordinate-axes ($U^H \dot{x}(t)$). We call the eigencomponents positive or negative based upon the sign of the corresponding eigenvalue in D and these may be used to linearly decompose the signal into class-sensitive components as in Figure 5A,B. In general, the magnitude of each component's eigenvalue is related to how tuned that component is for a specific class.

5.5.8. Testing The Conic Sensitivity to Noise—To illustrate the invariance properties of our conic classifier we considered trials number 21,24,27, and 30 for the lower right and upper left locations each for subject 3. Using this data we calculated the discriminating cone and projected all subsequent data onto the cone's main eigenvectors for visualization (the first two corresponding to lower right and the first for upper left). Thus, all transformed data are plotted on the same axis generated from the original cone. We then considered three noise-transformations on the data: multiplicative noise, dynamic temporal warping (DTW), and multivariate box noise. To simulate multiplicative noise we simply multiplied the pre-derivative data with a random spline before computing derivatives. We also used a random spline for the dynamic temporal warping example in which the spline values served as the time values at which we evaluated the data through interpolation. Thus, the transformation under multiplicative noise between a univariate spline $s(t)$ and multivariate data $X(t)$ was calculated as $s(t)X(t)$, while the transformation under DTW was calculated as $X(s(t))$ with non-integer values assessed through interpolation. In the latter case splines were constrained to be positive-valued and monotone. Random splines was calculated by first generating a set of random time spacing between interpolation control points (distributed $\mathcal{N}(0,10ms)$ and $\mathcal{N}(0, 50ms)$ for multiplication and DTW, respectively. The x-values for interpolation control points were likewise randomly generated from $\mathcal{N}(0, .2)$ and $\mathcal{N}(0, .05)$, respectively. To generate a temporal scaling function, the x-values generated from DTW were converted to their cumulative sum to generate a monotone time sequence and then rescaled to match the data's temporal bounds. Box noise was applied additively using a randomly generated box function. To create these functions we first specified a fixed number of jump points (20) and a vector-valued variance for the jump amplitudes of each channel. We generated these variances from the distribution $2\mathcal{N}(0,1)$. The spacing between jumps was generated from the distribution $5ms + \mathcal{N}(0, 5ms)^2$ separately for each channel and values were rescaled and floored to produce integer values within the data range. The amplitudes of each jump were generated from a normal distribution with the standard deviation mentioned previously (randomly drawn from $2\mathcal{N}(0,1)$ for each channel). After noise was added (independently for each class) data was projected into the original conic eigenspace to determine whether the original geometry still held.

Acknowledgments

MS was funded by NSF-DGE-1143954 from the US National Science Foundation. TB acknowledges R37 MH066078 from the US National Institute of Health. SC holds a Career Award at the Scientific Interface from the Burroughs-Wellcome Fund. We thank J Samaha for sending us his data and set of trials excluded from classification due to artifact. Portions of this work were supported by AFOSR 15RT0189, NSF ECCS 1509342 and NSF CMMI 1537015, from the US Air Force Office of Scientific Research and the US National Science Foundation, respectively.

References

- [1]. Andrzejak RG, Lehnertz K, Mormann F, Rieke C, David P, Elger CE, Indications of nonlinear deterministic and finite-dimensional structures in time series of brain electrical activity: Dependence on recording region and brain state, *Physical Review E* 64 (061907) (2001) 1–8. doi: 10.1103/PhysRevE.64.061907.
- [2]. Adeli H, Ghosh-Dastidar S, Dadmehr N, A wavelet-chaos methodology for analysis of eegs and eeg subbands to detect seizure and epilepsy, *IEEE Transactions on Biomedical Engineering* 54 (2) (2007) 205–211. doi:10.1109/TBME.2006.886855. [PubMed: 17278577]
- [3]. Adeli H, Ghosh-Dastidar S, Dadmehr N, A spatio-temporal wavelet-chaos methodology for eeg-based diagnosis of alzheimer's disease, *Neuroscience Letters* 444 (2) (2008) 190–194. doi: 10.1016/j.neulet.2008.08.008. [PubMed: 18706477]
- [4]. Sitges C, Bornas X, Llabrés J, Noguera M, Montoya P, Linear and nonlinear analyses of eeg dynamics during non-painful somatosensory processing in chronic pain patients, *International Journal of Psychophysiology* 77 (2) (2010) 176–183. doi:10.1016/j.ijpsycho.2010.05.010. [PubMed: 20538021]
- [5]. Iasemidis LD, Sackellares JC, Zaveri HP, Williams WJ, Phase space topography and the lyapunov exponent of electrocorticograms in partial seizures, *Brain Topography* 2 (3) (1990) 187–201. doi: 10.1007/BF01140588. [PubMed: 2116818]
- [6]. Güler NF, Übeyli ED, Güler , Recurrent neural networks employng lyapunov exponents for eeg signals classification, *Expert Systems with Applications* 29 (3) (2005) 506–514. doi:10.1016/j.eswa.2005.04.011.
- [7]. Takens F, Detecting strange attractors in turbulence, in: Rand D, Young L-S (Eds.), *Lecture Notes in Mathematics* 898, Springer-Verlag, Berlin, 1981, pp. 366–381.
- [8]. Sauer T, Yorke J, Casdagli M, Embedology, *Journal of Statistical Physics* 65 (1991) 579–616. doi: 10.1007/BF01053745.
- [9]. Tajima S, Yanagawa T, Fujii N, Toyozumi T, Untangling brain-wide dynamics in consciousness by cross-embedding, *PLoS Computational Biology* doi:10.1371/journal.pcbi.1004537.
- [10]. Smith RA, Existence of periodic orbits of autonomous ordinary differential equations, *Proceedings of the Royal Society of Edinburgh Section A: Mathematics* 85 (1–2) (1980) 153–172. doi:10.1017/S030821050001177X.
- [11]. Sanchez LA, Cones of rank 2 and the poincaré-bendixson property for a new class of monotone systems, *Journal of Differential Equations* 246 (2009) 1978–1990. doi:10.1016/j.jde.2008.10.015.
- [12]. Sanchez LA, Existence of periodic orbits for high-dimensional autonomous systems, *Journal of Mathematical Analysis and Applications* 363 (2) (2010) 409–418. doi:10.1016/j.jmaa.2009.08.058.
- [13]. Feng L, Wang Y, Wu J, Semiflows “monotone with respect to high-rank cones” on a banach space, *SIAM Journal on Mathematical Analysis* 49 (1) (2017) 142–161. doi: 10.1137/16M1064295.
- [14]. Treder MS, Bahramisharif A, Schmidt NM, van Gerven MA, Blankertz B, Brain-computer interfacing using modulations of alpha activity induced by covert shifts of attention, *Journal of NeuroEngineering and Rehabilitation* 8 (24). doi:10.1186/1743-0003-8-24.
- [15]. Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P, Distributed and overlapping representations of faces and objects in ventral temporal cortex, *Science* 293 (5539) (2001) 2425–2430. doi:10.1126/science.1063736. [PubMed: 11577229]

- [16]. Fuentemilla L, Penny WD, Cashdollar N, Bunzeck N, Düzel E, Theta-coupled periodic replay in working memory, *Current Biology* 20 (7) (2010) 606–612. doi:10.1016/j.cub.2010.01.057. [PubMed: 20303266]
- [17]. Jafarpour A, Horner AJ, Fuentemilla L, Penny WD, Duzel E, Decoding oscillatory representations and mechanisms in memory, *Neuropsychologia* 51 (4) (2013) 772–780. doi: 10.1016/j.neuropsychologia.2012.04.002. [PubMed: 22561180]
- [18]. Haxby JV, Connolly AC, Guntupalli JS, Decoding neural representational spaces using multivariate pattern analysis, *Annual Review of Neuroscience* 37 (1) (2014) 435–456. doi: 10.1146/annurev-neuro-062012-170325.
- [19]. Spampinato C, Palazzo S, Kavasidis I, Giordano D, Souly N, Shah M, 2017 IEEE Conference on Computer Vision and Pattern Recognition doi:10.1109/CVPR.2017.479.
- [20]. Wen H, Shi J, Zhang Y, Lu K-H, Cao J, Liu Z, Neural encoding and decoding with deep learning for dynamic natural vision, *Cerebral Cortex* (2017) 1–25 doi:10.1093/cercor/bhx268. [PubMed: 28365777]
- [21]. Garcia JO, Srinivasan R, Serences JT, Near-real-time feature-selective modulations in human cortex, *Current Biology* 23 (6) (2013) 515–522. doi:10.1016/j.cub.2013.02.013. [PubMed: 23477721]
- [22]. Foster JJ, Sutterer DW, Serences JT, Vogel EK, Awh E, The topography of alpha-band activity tracks the content of spatial working memory, *Journal of Neurophysiology* 115 (1) (2016) 168–177. doi:10.1152/jn.00860.2015. [PubMed: 26467522]
- [23]. Haufe S, Meinecke F, Görgen K, Dähne S, Haynes J-D, Blankertz B, Bießmann F, On the interpretation of weight vectors of linear models in multivariate neuroimaging, *Neuroimage* 87 (2014) 96–110. doi: 10.1016/j.neuroimage.2013.10.067. [PubMed: 24239590]
- [24]. Samaha J, Sprague TC, Postle BR, Decoding and reconstructing the focus of spatial attention from the topography of alpha-band oscillations, *Journal of Cognitive Neuroscience* 28 (8) (2016) 1090–1097. doi:10.1162/jocn_a_00955. [PubMed: 27003790]
- [25]. Smith HL, Monotone dynamical systems: An introduction to the theory of competitive and cooperative systems, no. 41, American Mathematical Society, 2008.
- [26]. Hopfield J, Neurons with graded response have collective computational properties like those of two-state neurons, *PNAS* 81 (1984) 3088–3092. doi:10.1073/pnas.81.10.3088. [PubMed: 6587342]
- [27]. Boser BE, Guyon IM, Vapnik VN, A training algorithm for optimal margin classifiers, *Proc. of the fifth annual workshop on Computational learning theory* (1992) 144–152 doi: 10.1145/130385.130401.
- [28]. Norman KA, Polyn SM, Detre GJ, Haxby JV, Beyond mindreading: multi-voxel pattern analysis of fmri data, *Trends in Cognitive Sciences* 10 (9) (2006) 424–430. doi:10.1016/j.tics.2006.07.005. [PubMed: 16899397]
- [29]. Ahnert K, Abel M, Numerical differentiation of experimental data: local versus global methods, *Computer Physics Communications* 177 (2007) 764–774. doi:10.1016/j.cpc.2007.03.009.
- [30]. Honey CJ, Sporns O, Cammoun L, Gigandet X, Thiran JP, Meuli R, Hagmann P, Predicting human resting-state functional connectivity from structural connectivity, *PNAS* 106 (6) (2009) 2035–2040. doi:10.1073/pnas.0811168106. [PubMed: 19188601]
- [31]. Lorenz EN, Deterministic nonperiodic flow, *Journal of the Atmospheric Sciences* (1963) 130–141 doi:10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2.
- [32]. Hindmarsh JL, Rose RM, A model of neuronal bursting using three coupled first order differential equations, *Proc. R. Soc. Lond. B* 221 (1984) 87–102. doi:10.1098/rspb.1984.0024. [PubMed: 6144106]
- [33]. Lloyd SP, Least squares quantization in pcm, *IEEE T IT* (2) (1982) 129–137. doi:10.1109/TIT.1982.1056489.
- [34]. Schmidt NM, Blankertz B, Treder MS, α -modulation induced by covert attention shifts as a new input modality for eeg-based bcis, *Systems Man and Cybernetics (SMC)*, 2010 IEEE International Conference on (2011) 481–487 doi:10.1109/ICSMC.2010.5641967.
- [35]. Etzel J, Zachs J, Braver T, Searchlight analysis: Promise, pitfalls, and potential, *NeuroImage* 78 (2013) 261–269. doi:10.1016/j.neuroimage.2013.03.041. [PubMed: 23558106]

- [36]. Górecki T, Łuczak M, Using derivatives in time series classification, *Data Mining and Knowledge Discovery* 26 (2) (2013) 310–331. doi: 10.1007/s10618-012-0251-4.
- [37]. Górecki T, Łuczak M, Multivariate time series classification with parametric derivative dynamic time warping, *Expert Systems with Applications* 42 (5) (2015) 2305–2312. doi:10.1016/j.eswa.2014.11.007.
- [38]. Łuczak M, Hierarchical clustering of time series data with parametric derivative dynamic time warping, *Expert Systems with Applications* 62 (2016) 116–130. doi:10.1016/j.eswa.2016.06.012.
- [39]. Masterton RA, Abbott DF, Fleming SW, Jackson GD, Measurement and reduction of motion and ballistocardiogram artefacts from simultaneous eeg and fmri recordings, *NeuroImage* 37 (2007) 202–211. doi:10.1016/j.neuroimage.2007.02.060. [PubMed: 17582785]
- [40]. Jerbi K, Ossandón T, Hamamé CM, Senova S, Dalal SS, Jung J, Minotti L, Bertrand O, Berthoz A, Kahane P, Lachaux J-P, Task-related gamma-band dynamics from an intracerebral perspective: Review and implications for surface eeg and meg, *Human Brain Mapping* 30 (6) (2009) 1758–1771. doi:10.1002/hbm.20750. [PubMed: 19343801]
- [41]. Canolty RT, Edwards E, Dalal SS, Soltani M, Nagarajan SS, Kirsch HE, Berger MS, Barbaro NM, Knight RT, High gamma power is phase-locked to theta oscillations in human neocortex, *Science* 313 (2006) 1626–1628. doi:10.1126/science.1128115. [PubMed: 16973878]
- [42]. Monto S, Palva S, Voipio J, Palva JM, Very slow eeg fluctuations predict the dynamics of stimulus detection and oscillation amplitudes in humans, *Journal of Neuroscience* 28 (33) (2008) 8268–8272. doi: 10.1523/JNEUROSCI.1910-08.2008. [PubMed: 18701689]
- [43]. Torta ABL, Komorowski RW, Mannsf JR, Kopellc NJ, Eichenbaume H, Theta-gamma coupling increases during the learning of item-context associations, *PNAS* 106 (49) (2009) 20942–20947. doi:10.1073/pnas.0911331106. [PubMed: 19934062]
- [44]. van de Nieuwenhuijzen ME, Backus AR, Bahramisharif A, Doeller CF, Jensen O, van Gerven MAJ, Meg-based decoding of the spatiotemporal dynamics of visual category perception, *Neuroimage* (2013) 1063–1073. doi:10.1016/j.neuroimage.2013.07.075. [PubMed: 23927900]
- [45]. King J-R, Dehaene S, Characterizing the dynamics of mental representations: the temporal generalization method, *Trends in Cognitive Sciences* (4) (2014) 203–210. doi:10.1016/j.tics.2014.01.002. [PubMed: 24593982]
- [46]. Cohen JD, Perlstein WM, Braver TS, Nystrom LE, Noll DC, Jonides J, Smith EE, Temporal dynamics of brain activation during a working memory task, *Nature* 386 (1997) 604–608. doi: 10.1038/386604a0. [PubMed: 9121583]
- [47]. Chang C, Glover GH, Time-frequency dynamics of resting-state brain connectivity measured with fmri, *Neuroimage* 50 (1) (2010) 81–98. doi:10.1016/j.neuroimage.2009.12.011. [PubMed: 20006716]
- [48]. Kafashan M, Ching S, Palanca BJA, Sevoflurane alters spatiotemporal functional connectivity motifs that link resting-state networks during wakefulness, *Frontiers in Neural Circuits* 10 (2016) 107. doi:10.3389/fncir.2016.00107. [PubMed: 28082871]
- [49]. Riehl JR, Palanca BJ, Ching S, High energy brain dynamics during anesthesia-induced unconsciousness, *Network Neuroscience* doi:10.1162/NETN_a_00023.
- [50]. Wang F, Gelfand AE, Directional data analysis under the general projected normal distribution, *Statistical Methodology* 10 (1) (2013) 113–127. doi:10.1016/j.stamet.2012.07.005. [PubMed: 24046539]
- [51]. Friedman J, Hastie T, Tibshirani R, Sparse inverse covariance estimation with the graphical lasso, *Biostatistics* 9 (3) (2008) 432–441. doi:10.1093/biostatistics/kxm045. [PubMed: 18079126]
- [52]. Wang J, Wavelet approach to numerical differentiation of noisy functions, *Communications on Pure and Applied Analysis* 6 (3) (2007) 873–897. doi:10.3934/cpaa.2007.6.873.
- [53]. Makeig S, Westerfield M, Jung T-P, Enghoff S, Townsend J, Courchesne E, Sejnowski TJ, Dynamic brain sources of visually evoked responses, *Science* 295 (5555) (2002) 690–694. doi: 10.1126/science.1066168. [PubMed: 11809976]
- [54]. Kutas M, Hillyard S, Reading senseless sentences: brain potentials reflect semantic incongruity, *Science* 207 (4427) (1980) 203–205. doi: 10.1126/science.7350657. [PubMed: 7350657]

- [55]. Kutas M, Federmeier KD, Thirty years and counting: Finding meaning in the n400 component of the event-related brain potential (erp), *Annual Review of Psychology* 62 (1) (2011) 621–647. doi: 10.1146/annurev.psych.093008.131123.
- [56]. Grech R, Cassar T, Muscat J, Camilleri KP, Fabri SG, Zervakis M, Xanthopoulos P, Sakkalis V, Vanrumste B, Review on solving the inverse problem in eeg source analysis, *Journal of NeuroEngineering and Rehabilitation* 5 (1) (2008) 25. doi:10.1186/1743-0003-5-25. [PubMed: 18990257]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Highlights for “Geometric classification of brain network dynamics via conic derivative discriminants”

- A new classifier decodes task states using the time evolution of neural data
- Information about time evolution is extracted from the derivatives of signals
- The classifier compares derivative covariance to decode cognitive states
- The classifier outperforms current methods in decoding spatial attention from EEG
- The method reveals retinotopy and new temporal markers of spatial attention

Fig 1.

The two variants of a recurrent continuous neural network (A,B). Line widths indicate connection strength and green/black indicates excitatory/inhibitory. The red markers in network B indicate a greater decay rate. Observed trajectories do not differentiate between generative systems in the original state space (C.1, C.2), but are readily distinguishable in the derivative space and its projection onto a spherical surface (D.1, D.2). Random intervals of each orbit are plotted for 150 initial conditions and two views (view 1: C.1,D.1; view 2: C.2,D.2) with derivatives projected onto S^{n-1} . In the original phase space (E) full trajectories of both systems approach a focus, however, they differ in the approach envelope which separates derivatives (F).

Fig 2.

Simulated comparison of blind conic vs. spectral classification. A) The system consists of an upstream chaotic attractor which provides binary downstream input to a two-layer recurrent system. Green cells are excitatory and red inhibitory. B) Example simulated voltages for the three cells in the bottom layer using median model parameters. Voltage traces are in black and blue bars indicate when the upstream input was “On”. C) The state of the upstream system is easily determinable by plotting the voltages in the normalized derivative space. D) Percent correct by decoding method and simulation parameters. Neither method was affected by the level of intrinsic system noise (IntVar), but conic decoding suffered with the addition of measurement noise (ExtVar). Conic decoding accuracy continued to increase with the upstream coupling strength (Stim), while spectral decoding showed a concave relationship. Both methods performed slightly better for moderate bin sizes (nWin) over very small ones but further increases did not improve performance. The shaded regions give standard deviations (n=136).

Fig 3.

Comparison of conic and alpha-band based decoding methods in terms of number of significantly decoded location-pairs. A) Both decoding methods possess large inter-subject variation in decoding performance. B) At the group-level conic decoding outperforms spectral. C) At the individual-level the techniques are negatively correlated and thus form complementary methods. For subject 8, the conic classifier did not significantly discriminate any locational pairs for the full uncensored data. However, censoring trials using the methods for eyeblink and motion artifact rejection adopted in ([24]) improved performance (see Materials and Methods).

Fig 4.

Anatomical weights distributions. Inverting conic weights to anatomically map classifier features produces a retinotopic map of covert spatial attention in posterior channels. Center panel indicates the corresponding region of head space (posterior coronal) while radial panels denote the weight mapping for subjects attending to that location vs. opposite locations with hot colors indicating increased prevalence in the discriminating spatiotemporal pattern.

Fig 5.

Conic Geometry for the median subject (#4). A) Plotting the first 2 negative components and the first positive component in derivative space reveals conic geometry. Data is displayed from a variety of viewpoints. B) Plotting the primary positive and negative eigenvectors of the cone for a 2-location contrast demonstrates that these components increase derivative amplitude during the corresponding cognitive state. C) Two dimensional projections onto the length in the "positive" and "negative" eigenspaces for the discriminative cone reveals robust classification of individual observations lasting only a few milliseconds.

Fig 6.

Decoding time series and performance for 6-way classification (all possible locations at once) using the conic method. A) The “instantaneous accuracy” time series produced by applying a classifier trained starting 750 ms post-cue (vertical dashed line) to the entire delay time series. Cone-like classifiers using the full covariance (red), the correlation-only (black), or the variance-only (blue) reveal an earlier onset of significant decoding than previously reported with the same data. Horizontal lines indicate the significance threshold for the permutation p-value indicated. Accuracy is for prediction over 30 ms windows and was smoothed with a Gaussian kernel. The shaded green bar indicates the early interval in which conic decoding was significant but spectral decoding was not ([24]). B) Decoding time series (same as A) but plotted to highlight the early anti-correlated behavior (outer bar) between information decodable by variance and information decodable by correlation. The inner bar is the same as in (A) and indicates that early decoding was primarily due to within-channel variance. Shaded regions indicate the standard error of mean accuracy across subjects. C) The conic classifier performed significantly better at the full 6-way spatial contrast than any other classification method. For support vector machines (SVM) using spectral power (cyan), only the alpha-band (8-13 Hz) was significantly different from chance (indicated by the dashed line) ([24]). All conic classifiers were significant and the restriction to either only variance or only correlation produced performance indistinguishable from that of alpha-power SVM.

Fig 7.

Invariant properties of the Conic Classifier with respect to different noise forms. Left Column displays three different transformations applied to the same segment of EEG data from one trial type. Middle Column) the same forms of noise applied to a constant EEG segment from a different trial type. Right Column) Data for each trial projected onto a cone derived from the original (pre-noise) data (blue=lower-right condition, orange=upper-left). Line color indicates channel number. A) The combination of derivative filtering and conic scale invariance leads to robust performance for scalar multiplicative noise. B) Scale invariance of cones leads to classifier invariance with respect to dynamic temporal warping (DTW). C) Filtering properties of derivatives lead to robust performance in the face of multivariate box noise.

Fig 8.

Conic Geometry for the best-decoding subject (#3). A) Plotting the first 2 negative components and the first positive component in derivative space reveals conic geometry. Data is displayed from a variety of viewpoints. B) Plotting the primary positive and negative eigenvectors of the cone for a 2-location contrast demonstrates that these components increase derivative amplitude during the corresponding cognitive state. C) Two dimensional projections onto the length in the "positive" and "negative" eigenspaces for the discriminative cone reveals robust classification of individual observations lasting only a few milliseconds.

Fig 9.

Conic Geometry for the worst-decoding subject (#8). A) Plotting the first 2 negative components and the first positive component in derivative space reveals conic geometry. Data is displayed from a variety of viewpoints. B) Plotting the primary positive and negative eigencomponents of the cone for a 2-location contrast demonstrates that these components increase derivative amplitude during the corresponding cognitive state. C) Two dimensional projections onto the length in the "positive" and "negative" eigenspaces for the discriminative cone reveals robust classification of individual observations lasting only a few milliseconds. Note the large spike in this subject's data that is only visible in one of the eigenvectors.

Fig 10.

Individualized decoding accuracy for multi-class decoding. The x-axis (“n”) denotes the number of classes and performance was averaged over all possible combinations of “n” spatial locations using the same windows as the pairwise comparisons (starting 250 ms post-cue) and including all trials lasting atleast 1300ms (without artifact rejection).

Fig 11.

Effects of preprocessing steps (artifact-censoring and eigenvalue censoring) on classification performance. A) The number of significantly decoded pairs for some subjects varied with the pre-processing steps. Conic decoding performance for subject 8 is at least as good as for spectral decoding whenever artifactual trials are removed. B) Group-level performance was very similar across pre-processing choices. C) The pairwise decoding accuracy is generally unaffected by preprocessing choices (as opposed to the number of significant pairs). For subject 8, however, accuracy generally increased with artifact rejection.

Table 1.

Group-level comparisons between conic and alpha-band decoding as reported in [14],[24]. The second column (“Measure”) indicates the parameter tested, namely the number of significant pairs (locations) decoded, the maximal decoding accuracy among all possible pairs, and the mean accuracy when performing 6-way classification (all locations at once). The original and conic results are given in Mean(SD) format. We used independent samples t-tests and all p-values reported are two-tailed. Due to the much greater variance of conic decoding for the 6-way classification relative to [24], we corrected for unequal variance in the reported t-value. We use the ** superscript to denote uncorrected significance at $p < .01$

	Measure	Original	Conic	t	p	df
Treder 11	Sig. Pairs	3.5(2.7)	10.3(1.8)	3.558 **	.004 **	15
	Max Accuracy (2-way)	74.6(2.3)	76.4(3.2)	.572	.577	15
Samaha 16	Mean Accuracy (6-way)	23(4)	32.1(7)	3.119 **	.0088 **	12.1

Table 2.

The relationship between conic and spectral ([14]) decoding (number of significant pairs) trends negative. Statistics correspond to nonparametric correlation (Kendall), parametric correlation (Pearson), and parametric with the high leverage subject (# 8) removed. All p-values are two-tailed.

Statistic	Outlier	Correlation	p	df
Kendall (<i>k</i>)	yes	-.5879	.0738	6
Pearson (<i>r</i>)	yes	-.355	.3882	6
Pearson (<i>r</i>)	no	-.940	.0016	5

Table 3.

Statistical comparisons of the reduced cones (correlation/variance only) to chance, alpha-band decoding ([24]), and the full conic classifier as well as to each other. Testing compared to chance was performed via one-sample t-test and comparisons between conic classifiers and alpha-band decoding were performed with independent samples, unequal variance t-tests. Comparisons between full/reduced conic classifiers were performed with paired t-tests. All p-values are two-tailed and uncorrected for multiple comparisons.

	Mean	t	Corr p	df	t	Var p	df
Chance	16.67	3.343 *	.0124 *	7	6.855 **	.00024 **	7
Samaha 16	23(4)	.529	.6058	12.69	.672	.5141	12.19
Full Cone	32.1(7)	-7.985 **	.00009 **	7	-3.815 **	.0066 **	7
Var	24.13(3.08)	.081	.938	7		-	
Corr	24.3(6.46)		-			-	

*
= $p < .05$

**
= $p < .01$.

Table 4.

Number of trials for each subject meeting the trial length criteria by class (cued location).

Location		1	2	3	4	5	6
	1	70	71	72	66	58	64
	2	70	66	72	73	59	67
	3	66	72	63	64	64	62
Subject	4	68	74	69	68	64	69
	5	72	73	62	77	64	61
	6	47	53	49	51	49	48
	7	70	67	70	73	57	60
	8	73	67	59	73	64	69