

REGULARIZED BRAIN READING WITH SHRINKAGE AND SMOOTHING

BY LEILA WEHBE^{*,1}, AADITYA RAMDAS^{*,2},
REBECCA C. STEORTS^{†,3} AND COSMA ROHILLA SHALIZI^{‡,4}

University of California, Berkeley^{}, Duke University[†]
and Carnegie Mellon University[‡]*

Functional neuroimaging measures how the brain responds to complex stimuli. However, sample sizes are modest, noise is substantial, and stimuli are high dimensional. Hence, direct estimates are inherently imprecise and call for regularization. We compare a suite of approaches which regularize via *shrinkage*: ridge regression, the elastic net (a generalization of ridge regression and the lasso), and a hierarchical Bayesian model based on small area estimation (SAE). We contrast regularization with *spatial smoothing* and combinations of smoothing and shrinkage. All methods are tested on functional magnetic resonance imaging (fMRI) data from multiple subjects participating in two different experiments related to reading, for both predicting neural response to stimuli and decoding stimuli from responses. Interestingly, when the regularization parameters are chosen by cross-validation independently for every voxel, low/high regularization is chosen in voxels where the classification accuracy is high/low, indicating that the regularization intensity is a good tool for identification of relevant voxels for the cognitive task. Surprisingly, all the regularization methods work about equally well, suggesting that beating basic smoothing and shrinkage will take not only clever methods, but also careful modeling.

1. Introduction. A major goal of functional brain imaging is to relate activity levels in various parts of the brain to differences in stimuli. Typical fMRI experiments measure activity in tens of thousands of volume elements called voxels (i.e., “volume pixels”) within the brain, over 10^2 – 10^3 time steps, while realistic stimuli vary on hundreds or thousands of dimensions (see Section 1.1). Moreover, neuroscientists want to study heterogeneity across the brain in responses to stimuli, discounting noisy variations. Taking each voxel on its own, estimates of response functions are inherently imprecise due to the level of the noise and the high dimensionality of the problem. In statistics, such estimation problems are addressed

Received January 2014; revised December 2014.

¹Supported by NIH Grant 5R01HD075328-02 and by DARPA Grant FA8750-13-2-0005.

²Supported by NSF Grant IIS-1247658.

³Supported by NSF Grants SES1130706 and DMS-10-43903 as well as the John Templeton Foundation.

⁴Supported by NIH Grant #2 R01 NS047493 and by NSF Grant DMS-12-07759.

Key words and phrases. fMRI, small area estimation, regularization, shrinkage, spatial smoothing.

by regularization, especially shrinking estimates toward a reference value (e.g., 0). While shrinkage statistically stabilizes parameters estimates, this may or may not help achieve the scientific aim of better understanding of the organization of the brain. Due to this, we examine whether, and how, common regularization techniques serve the inferential goals of cognitive neuroscience.

Because such an investigation cannot be done abstractly, we study the behavior of four different methods of regularizing linear regression, in two experiments related to different aspects of reading. Three of our methods regularize by shrinkage: ridge regression; the elastic net, which generalizes both ridge and the lasso; and a hierarchical Bayesian (HB) model, developed for small area estimation (SAE). Ridge regression, the lasso, and the elastic net exemplify modern high-dimensional frequentist statistics, based on penalized optimization. The SAE model is an instance of the Bayesian approach increasingly used in neuroscience [Genovese (2000), Lee et al. (2011), Park et al. (2013)], where a hierarchical process generates the parameters. Our fourth method of regularization smooths the data over spatial regions. We also consider combinations of shrinkage and spatial smoothing, developing a novel decision-theoretic method for smoothed SAE in the spirit of Louis (1984) and Datta et al. (2011). All methods were compared to the performance of unregularized ordinary least squares (OLS) regression.

As mentioned, we evaluated our methods on two experiments: one studying the representation of the meaning of individual word-picture pairs (E1, Section 1.2), and the second studying story comprehension (E2, Section 1.3). The two experiments differed in their subject pool, in the nature of the stimulus (independent, randomly presented word picture-pairs vs. consecutive words of a real story that requires the maintenance of a complex context representation), in how long each stimulus was presented (10 s vs. 0.5 s), and in the nature of the appropriate analysis (static vs. time series). Findings about regularization methods common to both experiments are unlikely to be artifacts of just one experiment.

Surprisingly, despite their different rationales and inner workings, *all* of our methods of regularization gave very similar out-of-sample performances in both experiments. All achieved low mean-squared-errors in predicting neural responses to stimuli, and high accuracy in classifying novel stimuli based on neural response (Section 3.1). They improved in both respects over unregularized OLS, though only slightly. They produced very similar parameter estimates, especially ridge regression and SAE, a connection explained in Section 2.3. They showed a consistent pattern of how much estimates in different parts of the brain were regularized—they imposed more shrinkage or more smoothing in areas of low signal strength. This suppression of “noisy” brain regions is perhaps their greatest advantage over OLS (Section 3.2.2). This indicates that single-voxel regularization could be viewed as a detector for informative brain regions because it allows predicted brain activity to be different from zero only in the informative, less noisy regions. The near-equal performance of all regularizers means that choices between them must be based on considerations such as computational cost (Section 4.2)

and/or biological plausibility. Improving on these outcomes must come from better biological modeling and not more clever general-purpose statistical methods.

The rest of this paper proceeds as follows. We present the necessary neuroscientific background in Section 1.1 and summarize the two data sets used in this paper (Sections 1.2 and 1.3). We describe the details of our methods in Section 2, provide in-depth results in Section 3, and conclude with a discussion in Section 4.

1.1. *Neuroscientific background.* Cognitive neuroscientists use functional magnetic resonance imaging to study how the brain implements cognition. fMRI specifically measures “hemodynamic response,” the change in blood oxygen levels after neural activity, as a proxy for information processing [Ashby (2011)]. Let y_{vt} be the measured activity at voxel v and discrete time-point t . While these measurements are often smoothed spatially to reduce noise, most analyses involve running a separate regression of each voxel against stimulus features [Ashby (2011), Chapter 5].

In experiments with static stimuli that are presented with enough time in between different instances to make them effectively independent, one could make the design matrix of the regression simply contain stimulus features, and use an appropriate time window to obtain a single average response y_{vt} for each stimulus t . In such cases, for precision there are usually multiple repetitions of each stimuli. On the other hand, in experiments with dynamic stimuli, the contents of the design matrix for such regressions are dictated by the fact that the hemodynamic response has a long time latency, typically peaking about six seconds after stimulus onset. The time-courses of stimulus features are convolved with a kernel function modeling the hemodynamic response, resulting in a time-varying set of covariates in the design matrix. For each v , y_{vt} is regressed against these covariates.

Statically or dynamically, it is typically assumed that voxels which are found to have statistically significant regression coefficients are actually involved in processing the stimuli. (We return to this assumption in Section 4.3.) The focus on statistical significance, and the fact that there are usually more points than covariates, explains the popularity of unregularized OLS [Ashby (2011), Chapter 5]. Spatial information is not often utilized, but region information is sometimes used to threshold significance maps, searching for contiguous blobs of significant voxels [Smith (2004)].

More recently, multivariate pattern analysis has used information from multiple voxels to decode underlying cognitive states. In discriminative models, fMRI images are fed as input to a classifier, which attempts to “reverse infer” the stimulus or the state of the brain. Some of these methods, especially discriminative Bayesian models, take advantage of the spatial smoothness of the fMRI image [Norman et al. (2006), Park et al. (2013), Pereira, Mitchell and Botvinick (2009)]. Our interest, however, lies in generative models that can both predict the fMRI images that arise in response to new stimuli and decode stimuli from responses.

Discriminative models can only decode, whereas generative models aim for a more complete understanding of neural dynamics [Naselaris et al. (2011)].

We use fMRI data from two experiments, the first static and the second dynamic, for both predicting neural response to stimuli as well as decoding stimuli from responses. While both involve reading, the two experiments probe reading differently, which we briefly describe in Sections 1.2 and 1.3. In both experiments, the data is analyzed by using feature representations of the stimuli and then expressing brain activity as a function of these stimulus features; this idea was introduced in Mitchell et al. (2008). The use of feature representations allows the experimenter to predict brain activity for a novel *unseen* stimulus, by multiplying regression coefficients learned on old *seen* stimuli by the feature representations of the new stimulus. Thus, the model that is learned can be assessed in terms of how well it generalizes its prediction to unseen stimuli.

1.2. *Experiment 1: Visual features of word-picture combinations.* The first experiment (E1) scanned native English speakers as they looked at word-picture combinations, specifically sixty concrete nouns (e.g., “apple”, “car”), accompanied by black-and-white line drawings of those objects [Mitchell et al. (2008)]. All nine subjects were exposed six times each to all sixty word-picture stimuli, varying in order. Here the latency of the hemodynamic response was handled by averaging the activity acquired 4–8 seconds after stimulus onset, resulting in a single brain image per subject per stimulus per exposure. The six repetitions of each stimulus are themselves averaged together (within subjects) in the data set.⁵

Each voxel was $3.125 \text{ mm} \times 3.125 \text{ mm} \times 6 \text{ mm}$, and every subject’s brain contained $\approx 21,000$ voxels. The subjects’ brains were morphed into the same anatomical space, although exact overlap is not achieved due to anatomical differences. The voxels are divided into 90 “regions of interest” (ROIs), generally believed to be anatomically and functionally distinct [Tzourio-Mazoyer et al. (2002)]. The ROIs vary greatly in size, from about 20 to about 800 voxels. For ROIs covering a large volume of the brain, the spatial smoothness we hope to exploit is washed out. To counter this, and achieve uniformity of size, we divided ROIs that had more than 200 voxels in half along their largest dimension (x , y , or z coordinate). This was repeated as necessary until all regions had 200 voxels or less. After this, we had 191 ROIs.

We used eleven features related to the visual properties of the stimuli (e.g., “amount of white pixels on the screen”, “2D aspect ratio”). These annotations were provided to us by the authors of Sudre et al. (2012), who used the same stimulus set for a different experiment. The original experiment reported these features as ordinal variables on a five-point scale. We selected these features since

⁵Data were obtained from <http://www.cs.cmu.edu/~fmri/science2008/data.html>, accessed in November 2013.

they represent a fairly coherent set of precisely measured aspects of the stimuli, ones whose processing is well understood neurobiologically [Shepherd (1994)]. For the same reasons, we did not use the many other features also measured in the experiment which are related to semantic or physical properties of the stimuli (e.g., “Is it manmade?”, “Can I hold it in one hand?”), as manually rated on the same five-point scale by workers on Amazon’s Mechanical Turk crowdsourcing system [Sudre et al. (2012)].

In summary, the data (E1) consists of sixty words, represented by eleven features each, and their associated average voxel activity across nine subjects.

1.3. *Experiment 2: Textual features in narrative comprehension.* Our second experiment deals with the response to dynamic textual features in a narrative comprehension task [Wehbe et al. (2014)]. Eight subjects read Chapter 9 of *Harry Potter and the Sorcerer’s Stone* [Rowling (2012)] while in the fMRI scanner. In order to know exactly when each word was processed by the subjects, only one word from the text was shown at a time, on the center of a screen, each word being projected for 0.5 seconds. The sampling rate of fMRI acquisition was 2 seconds per observation, hence, four consecutive words were read during the time it took to scan the whole brain once. The experiment lasted 2710 seconds in total, giving us 1355 full brain scans.⁶

Spatially, the voxels were $3\text{ mm} \times 3\text{ mm} \times 3\text{ mm}$, somewhat smaller than in E1, and every subject’s brain contained $\approx 29,000$ voxels. As done in E1, the brains were morphed into a common space and divided into ROIs, and we further subdivided excessively large ROIs.

As in E1, we again look at only features related to the visual properties of the stimuli. Since the visual stimulus being received by the subject at any time is just a word printed on a screen, standardized in color, font, etc., we focus on a single quantitative textual feature which is comparable across words, namely, their length in letters. Each observation spans four words and, hence, we used both the mean and the standard deviation of the length of the presented words as our features. To account for the latency and persistence of the hemodynamic response, the stimulus features at time t are used as regressors for the activity at times $t + 1$ through $t + 4$. As before, we discard many of the features from Wehbe et al. (2014) relating to different kinds of semantic properties of the stimuli, like parts-of-speech tags (noun, verb, etc.) as well as other aspects of the story (characters, suspense, etc.) to maintain consistency in the paper and comparability with E1.

In summary, the data (E2) consist of two time series: (1) the mean and standard deviation of word lengths in every two-second interval and (2) the associated time series of voxel activities across eight subjects.

⁶Data is available at <http://www.cs.cmu.edu/~fmri/plosone>.

Contrast between E1 and E2. E1 probes the processing of static visual stimuli in a rather simple (even artificial) reading task. In contrast, E2 deals with dynamic, textual features in a narrative comprehension task. While the serial presentation of words is rare outside of the laboratory, the words were presented at a comfortable rate, and the subjects were previously asked to practice reading in this serial fashion [Wehbe et al. (2014)], making the overall setting much closer to “ecological validity” than is E1. Common findings about the properties and performance of statistical methods across such different settings are very unlikely to be artifacts of a particular experiment. We now turn to the description of the methods applied to both E1 and E2.

2. Methods. In previous analyses of both experiments [Mitchell et al. (2008), Wehbe et al. (2014)], the neural response to reading a word was modeled as a linear combination of the word’s features. While such linear models are ubiquitous in fMRI data analyses [Ashby (2011)], they have little biological basis. Nevertheless, any smooth model can be locally approximated by a linear regression over a sufficiently small domain, where the range of the feature variables here is fairly small. Plotting actual responses against linear fits shows that the latter are reasonable in these experiments (Figure 4 of the Supplementary Article [Wehbe et al. (2015)]). Hence, we follow the existing literature in using linear models, and explore multiple ways of fitting and regularizing them—OLS, ridge regression, the elastic net, and a hierarchical Bayesian model from small area estimation (SAE). We then consider including the effects of combining these techniques with various forms of spatial smoothing. Section 2.5 outlines our evaluation criteria for models and their regularizations, by their ability to both predict neural activity from stimuli and to reconstruct stimuli from activity.

2.1. Notation and model specification. We introduce notation consistent throughout the paper and note that we refer to real-valued variables by lowercase letters without boldface, vectors as boldfaced lowercase letters, and matrices in boldfaced uppercase.

In the linear model for static experiment E1, the average hemodynamic response y_{vt} of voxel v (for $v = 1, \dots, V$ for $V \approx 21,000$, varies per subject) to the stimulus, a word, and its associated image, displayed at time t (for $t = 1, \dots, T$ for $T = 60$), is a linear combination of stimuli features denoted by the P -dimensional feature vector \mathbf{x}_t (for $P = 11$),

$$y_{vt} = \mathbf{x}_t^\top \boldsymbol{\beta}_v + \varepsilon_{vt},$$

where $\boldsymbol{\beta}_v$ is the P -dimensional regression coefficient vector of v and ε_{vt} is mean-zero noise for voxel v at time t , with variance σ_v^2 , combining measurement error corrupting our observation with fluctuations and the effects of specification error. Finally, we assume that the ε_{vt} has a Gaussian distribution. More succinctly, we

will stack the \mathbf{x}_t s into a $T \times P$ matrix \mathbf{X} , and for each voxel v , write its activity over the course of E1 as a T -dimensional vector \mathbf{y}_v .

For dynamic experiment E2, the activity y_{vt} of voxel v (with $V \approx 29,000$, varies per subject) at time t (with $T = 1355$) is modeled as a linear function of the history of the stimulus, a continuous story, whose visual features are represented here as a time-series of two-dimensional vectors \mathbf{x}_t ,

$$y_{vt} = \sum_{k=1}^h \mathbf{x}_{t-k}^\top \boldsymbol{\beta}_{v,k} + \varepsilon_{vt},$$

where h represents how long the hemodynamic response to a stimulus persists and $\boldsymbol{\beta}_{vk}$ captures how the hemodynamic response at voxel v depends on the k th previous set of four words. We note that the mean and standard deviation of the word length of the t th set of four words is presented during the t th brain scan. We do not include $k = 0$ because we assume the time window when a stimulus is presented is too early to see a significant response of the voxel. As noted in Section 1.3, we set $h = 4$ here, meaning that the voxel activity at any time t is only affected by the preceding 8 seconds (16 displayed words). We also require $\boldsymbol{\beta}_{v,0} = 0$, meaning that there is a lag of two seconds before the hemodynamic response is seen. At any time t , the latest set of four words (taking 2 seconds) captured by x_{vt} does not play a role in the latest activity y_{vt} .

This may be put in a form more similar to the static case by regressing \mathbf{y}_v on the vector obtained by concatenating $\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}, \mathbf{x}_{t-3}$ into a single P -dimensional feature vector $\bar{\mathbf{x}}_t$ (for $P = 8$). We can similarly concatenate the regression coefficients for this concatenated feature vector to get a P -dimensional regression vector $\bar{\boldsymbol{\beta}}_v$. We overload notation to refer to $\bar{\mathbf{x}}_t$ and $\bar{\boldsymbol{\beta}}_v$ as \mathbf{x}_t and $\boldsymbol{\beta}_v$, since from this point the methods apply to both static and dynamic settings.

In both cases, the residual sum of squares is

$$\text{RSS}_v = \sum_{t=1}^T (y_{vt} - \mathbf{x}_t^\top \boldsymbol{\beta}_v)^2 = \|\mathbf{y}_v - \mathbf{X}\boldsymbol{\beta}_v\|_2^2,$$

where $\|\cdot\|_2^2$ is the squared Euclidean norm. OLS estimates $\boldsymbol{\beta}_v$ by minimizing the in-sample RSS, giving $\hat{\boldsymbol{\beta}}_v = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y}_v$. The covariance of the estimates, in a fixed design, is $\sigma_v^2 (\mathbf{X}^\top \mathbf{X})^{-1}$.

2.2. Ridge regression and elastic net. We now review both ridge regression and elastic net, giving the Bayesian counterparts to both. Ridge regression stabilizes OLS estimates via a penalty term [Hoerl and Kennard (1970)]. Specifically, the ridge estimator solves

$$(1) \quad \hat{\boldsymbol{\beta}}_v^R = \underset{\boldsymbol{\beta}_v}{\operatorname{argmin}} \text{RSS}_v + \lambda_v \|\boldsymbol{\beta}_v\|_2^2.$$

Equivalently, β_v^R is constrained to be small, $\|\beta_v^R\|_2^2 \leq c$, for some $c > 0$. The tuning parameter λ_v controls the degree of regularization. The ridge approach has been used before in neuroimaging with the *same* λ for all voxels [Mitchell et al. (2008)]. Importantly, in Section 3, we show that tuning λ *separately for each voxel* improves classification and prediction and provides valuable information about neural organization.

While ridge regression was developed from a frequentist perspective, it has a well-known Bayesian interpretation [Hastie, Tibshirani and Friedman (2001)]. By imposing a Gaussian prior on β_v with prior precision λ , we find

$$(2) \quad \begin{aligned} y_{vt} | \mathbf{x}_t, \beta_v &\stackrel{\text{i.i.d.}}{\sim} N(\mathbf{x}_t^\top \beta_v, \sigma_v^2), \\ \beta_v &\stackrel{\text{i.i.d.}}{\sim} N(0, 1/\lambda_v \mathbf{I}). \end{aligned}$$

Under the formulation in (2), the posterior mode coincides exactly with the solution to (1). The solution to both formulations has a closed form:

$$\hat{\beta}_{v, \lambda_v}^R = (\mathbf{X}^\top \mathbf{X} + \lambda_v \mathbf{I})^{-1} \mathbf{X}^\top \mathbf{y}_v.$$

The covariance is $\sigma_v^2 (\mathbf{X}^\top \mathbf{X} + \lambda_v \mathbf{I})^{-1} \mathbf{X}^\top \mathbf{X} (\mathbf{X}^\top \mathbf{X} + \lambda_v \mathbf{I})^{-1}$, in a fixed-design regression.

The *elastic net* of Zou and Hastie (2005) generalizes ridge regression and the lasso of Tibshirani (1996):

$$\hat{\beta}_v^{\text{EN}} = \underset{\beta_v}{\text{argmin}} \text{RSS}_v + \lambda_{1v} \|\beta_v\|_1 + \lambda_{2v} \|\beta_v\|_2^2.$$

Setting $\lambda_{1v} = 0$ recovers ridge regression, and $\lambda_{2v} = 0$ recovers the lasso. The L_1 penalty makes $\hat{\beta}_v^{\text{EN}}$ sparse, shrinking coefficients on superfluous variables to zero, while the L_2 penalty alone favors small but nonzero coefficients. Again, previous neuroimaging studies favor setting λ_1, λ_2 globally, but we find improved performance by varying them across voxels (Section 3), as chosen by cross-validation [implemented in the `glmnet` MATLAB package by Friedman et al. (2010)].

As with ridge regression, the elastic net estimate can be viewed as the MAP estimate of a Bayesian model. As shown by Kyung et al. (2010), the required prior is a gamma-scale mixture of Gaussians:

$$(3) \quad \begin{aligned} y_{vt} | \mu_v, \mathbf{x}_t, \beta_v, \sigma_v^2 &\sim N(\mu_v + \mathbf{x}_t^\top \beta_v, \sigma_v^2), \\ \beta_v | \sigma_v^2, \mathbf{D}_\tau^* &\sim N(0, \sigma_v^2 \mathbf{D}_\tau^*), \\ \tau_1^2, \dots, \tau_P^2 &\sim \prod_{j=1}^P \frac{\lambda_1^2}{2} e^{-\lambda_1^2 \tau_j^2 / 2} d\tau_j^2, \quad \tau_1^2, \dots, \tau_P^2 > 0, \end{aligned}$$

where $\mathbf{D}_\tau^* = \text{Diag}\{(\tau_i^{-2} + \lambda_2)^{-1}\}$ for all i .

2.3. *Hierarchical Bayesian small area model.* It is biologically plausible that voxels within the same ROI respond similarly to stimuli. Penalization methods, such as the elastic net, make estimates of regression coefficients more precise via stabilization but do not pool information from related voxels. In contrast, techniques for stabilizing parameter estimates by partially pooling information across, or borrowing strength from, related areas have been extensively developed in the literature on small area estimation (henceforth SAE) [Rao (2003)]. While not traditional in neuroscience, SAE is well known to be effective at shrinkage when there are multiple regions [Pfeffermann (2013)], here ROIs. Hence, we explore simple SAE methods for regularization which incorporate ROI-level effects, without completely pooling within ROIs.

The SAE literature typically accomplishes partial pooling using hierarchical Bayesian (HB) models, so we follow that precedent. As before, we model the activity y_{vt} in a voxel v as a linear combination of the stimulus features \mathbf{x}_t :

$$(4) \quad \begin{aligned} y_{vt} &= \mathbf{x}_t^\top (\mathbf{z}_v + \mathbf{u}_{A(v)}) + \varepsilon_{vt} \\ &= \mathbf{x}_t^\top \boldsymbol{\beta}_v^{SA} + \varepsilon_{vt}, \end{aligned}$$

where $A(v)$ is the ROI containing voxel v , \mathbf{u}_a is a coefficient vector common to all voxels in area a , and \mathbf{z}_v is the coefficient vector specific to voxel v . We have

$$\begin{aligned} y_{vt} | \boldsymbol{\beta}_v^{SA}, \sigma_v^2 &\sim \mathcal{N}(\mathbf{x}_t^\top \boldsymbol{\beta}_v^{SA}, \sigma_v^2), \\ \boldsymbol{\beta}_v^{SA} &= \mathbf{u}_{A(v)} + \mathbf{z}_v, \\ \mathbf{z}_v | v_v^2 &= \mathcal{N}(0, v_v^2 \mathbf{I}), \\ \mathbf{u}_a | \alpha_a^2 &\sim \mathcal{N}(0, \alpha_a^2 \mathbf{I}), \\ \sigma_v^2 &\sim \mathcal{IG}(a, b), \\ \alpha_a^2 &\sim \mathcal{IG}(c, d), \\ v_v^2 &\sim \mathcal{IG}(e, f), \end{aligned}$$

where a, b, c, d, e , and f are user-fixed hyperparameters, and \mathcal{IG} (shape, scale) is the inverse gamma distribution. The full conditional distributions of all parameters are straightforward (Appendix A of the Supplementary Article [Wehbe et al. (2015)]), so the model can be estimated effectively using partially parallelized Gibbs sampling.

Just as ridge and the elastic net have Bayesian interpretations, the MAP estimates of this Bayesian SAE model can be seen as a penalized least-squares estimate. Such an estimate is (surprisingly) close to the estimate delivered by ridge regression, for the following reason: the SAE model has a Gaussian prior distribution $\mathbf{z}_v | v_v^2 \sim \mathcal{N}(0, v_v^2 \mathbf{I})$ for the regression coefficients specific to voxel v , and the voxel-specific variance has an inverse gamma prior distribution, where $v_v^2 \sim \mathcal{IG}(e, f)$.

Due to this, the marginal prior distribution of \mathbf{z}_v is a scaled t -distribution, which is well approximated by a Gaussian for reasonable values of the hyperparameters (see Appendix B of the Supplementary Article [Wehbe et al. (2015)] for details). Section 4.1 revisits the statistical implication of this mathematical approximation, which is that the posterior mode of the HB model must actually be close to the ridge regression estimate.

2.4. *Spatial smoothing.* Neuroimaging data is extremely noisy, and estimates have high variance, even after shrinkage. Much of this noise occurs at high spatial frequencies [Ashby (2011), Chapter 4], and spatial smoothing can help reduce the variance. Since nearby voxels often tend to share activation patterns, spatial averaging may cancel out such noise but maintain signal. Biologically, nearby voxels should tend to respond similarly to stimuli, since recordings of individual cells show that many areas of the brain have a regular spatial organization in their responses to stimuli [Shepherd (1994)]. While the length scales over which individual neurons' responses vary do not coincide with the sizes of voxels, which generally contain many cells with heterogenous properties, it is still the case that nearby voxels should have correlated responses to stimuli. Since the noise in fMRI data is often at much higher spatial frequencies than the signal from voxels, it is reasonable to think that spatially smoothing the activity will enhance the signal-to-noise ratio. This is often done as a preprocessing step [Ashburner et al. (2008)], but we examine it here as a means of stabilizing parameter estimates.

We explore two kinds of spatial smoothing: *nearest-neighbor voxel-level* and *ROI area-level smoothing*. First we introduce these two forms of smoothing, and then consider smoothed OLS estimates.

2.4.1. *Nearest-neighbor voxel-level and ROI area-level smoothing.* *Nearest-neighbor voxel-level smoothing* replaces every voxel by the local average of its nearby voxels. This is done either for the activity levels \mathbf{y}_v or the parameter estimates $\boldsymbol{\beta}_v$. Lacking more anatomically-based metrics, we define "nearness" using standard ℓ_p distances of two vectors \mathbf{r}_1 and \mathbf{r}_2 :

$$\|\mathbf{r}_1 - \mathbf{r}_2\|_p \equiv (|r_{11} - r_{21}|^p + |r_{12} - r_{22}|^p + |r_{13} - r_{23}|^p)^{1/p}.$$

When $p = 2$, this is Euclidean distance and the ℓ_p ball around a voxel contains all other voxels whose centers fall within the given radius. However, when $p = 1$, the ℓ_p ball is a tetrahedral pyramid. We choose a smoothing range or radius separately for each voxel by cross-validation, and replace its value by the average over all voxels within the ℓ_p ball.⁷

ROI area-level smoothing is defined through solving an optimization problem. Taking the set of regression coefficients in one ROI A, $\mathbf{B}_A := \{\boldsymbol{\beta}_v\}_{v \in A}$, which is a

⁷For a given radius, the ℓ_1 ball contains fewer voxels than the ℓ_2 , and both are smaller than the ℓ_∞ ball. The latter did so poorly in trials that we only consider ℓ_1 and ℓ_2 .

$P \times |A|$ matrix. We penalize large differences between regression coefficients of voxels in the same area. In the Bayesian setting, these are the voxel-wise Bayes estimates. Specifically, for each ROI A , define $\tilde{\mathbf{B}}_A$ as

$$\tilde{\mathbf{B}}_A = \underset{\tilde{\mathbf{B}} = \{\tilde{\mathbf{b}}_v\}_{v \in A}}{\operatorname{argmin}} \sum_{v \in A} \|\tilde{\mathbf{b}}_v - \boldsymbol{\beta}_v\|_2^2 + \gamma \sum_{i, j \in A} q_{ij}^A \|\tilde{\mathbf{b}}_i - \tilde{\mathbf{b}}_j\|_2^2,$$

with penalty factor γ and $|A| \times |A|$ similarity matrix \mathbf{Q}^A . Fixing $q_{ij}^A = 1$ for all $i, j \in A$, leads to more uniform smoothing. However, letting

$$q_{ij}^A = \exp\{-d(i, j)^2/h^2\}$$

if $i, j \in A$, where d is the Euclidean distance between the locations of voxels i and j and h is a bandwidth, allows closer voxels to be more influential. Since the above optimization problem splits across the dimensions of $\boldsymbol{\beta}_v$, we get P independent optimization problems. Denoting the p th row of $\tilde{\mathbf{B}}_A$ as $\tilde{\mathbf{b}}_p^A$, we find

$$\sum_{i, j \in A} q_{ij}^A (\tilde{\mathbf{b}}_{ip} - \tilde{\mathbf{b}}_{jp})^2 = \tilde{\mathbf{b}}_p^{A\top} \boldsymbol{\Omega}_A \tilde{\mathbf{b}}_p^A,$$

where $\boldsymbol{\Omega}_A := 2(D^A - Q^A)$ is twice the graph Laplacian formed using Q^A as the adjacency matrix and D^A as a diagonal matrix whose i th entry is $\sum_j Q_{ij}^A$ [von Luxburg (2007), Proposition 1]. Hence, $\tilde{\mathbf{B}}_A = (I + \gamma \boldsymbol{\Omega}_A)^{-1} \mathbf{B}_A$. Parameters γ and h are chosen by cross-validation.

2.4.2. Smoothed OLS. Since OLS estimates are linear in y_{vt} and covariates are identical across voxels, smoothing $\boldsymbol{\beta}_v$ is equivalent to smoothing y_{vt} . At any voxel v , let \mathcal{S}_v be the set of voxels which are combined with it in smoothing, with the weight of voxel $u \in \mathcal{S}_v$ in the smoothing for v being c_{uv} . These weights are functions of the radius of smoothing in the nearest neighbor version, or of γ and q for ROI-level smoothing. Then the smoothed estimate at v is

$$\begin{aligned} \hat{\boldsymbol{\beta}}_v &= \sum_{u \in \mathcal{S}_v} c_{uv} \hat{\boldsymbol{\beta}}_u \\ &= \sum_{u \in \mathcal{S}_v} c_{uv} (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y}_u \\ &= (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \sum_{u \in \mathcal{S}_v} c_{uv} \mathbf{y}_u \\ &= (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \bar{\mathbf{y}}_v, \end{aligned}$$

which is the OLS estimate with the smoothed response $\bar{\mathbf{y}}$.⁸

⁸We are certainly not the first to note that linear smoothing commutes with OLS estimation—see, for example, Friston et al. [(2010), page 12].

Despite the simplicity of the technique, smoothed OLS produces results quite comparable to regularization methods such as ridge regression (see Section 3).

The equivalence of smoothing parameter estimates and smoothing the activity does not hold with our other, nonlinear estimators. When we report results for combinations of smoothing with other forms of regularization, we are smoothing the parameter estimates.

2.5. Evaluation criteria. Typically, cognitive neuroscientists engage in two forms of predictive inference with fMRI: *forward inference*, from stimuli to configurations of activity over the brain, and *reverse inference*, from patterns of activity to stimuli. While these are often approached as two separate tasks with two distinct sets of models, we perform both forward and reverse inference, using a common model.

Forward inference is a regression problem, where the regression models reviewed above can be applied immediately. Our evaluation criterion for forward inference is the voxel-wise residual sum of squares, normalized by the total sum of squares.

Reverse inference is more delicate. If we were primarily interested in decoding stimuli from observed neural activity, we could follow the usual practice in fMRI data analysis of estimating “tailored” classifiers or discriminative models [Pereira, Mitchell and Botvinick (2009), Poldrack (2008), Yarkoni et al. (2011)]. These might be accurate for the particular conditions they were trained on, but by construction they cannot generalize to previously unseen stimuli, unless they predict as an intermediate step the individual features of the stimuli and then identify the correct stimuli based on the decoded features [Sudre et al. (2012)]. Moreover, discriminative models do not directly represent anything about how the brain processes information, which is the main point of scientific interest.⁹ As shown by Haufe et al. (2014), the parameters learned in a decoding model, corresponding to each voxel’s contribution in a decoding task, cannot be readily used to infer if a voxel is representing a task of interest. For example, some voxels that represent a background process unrelated to the task might receive a high regression weight that serves to subtract that process from the voxels that are informative to the task.¹⁰

As shown by Mitchell et al. (2008), it is possible to use a forward model to do reverse inference, and doing so provides an additional check on the forward model’s ability to represent how the brain processes stimuli. This is an instance of “zero-shot classification” [Palatucci et al. (2009)] adapted to the neural prediction task, where the model is trained as usual, but with the data for some stimuli held

⁹Symbolically, scientists want to know about $p(Y|X)$, while discriminative models at best give $p(X|Y)$, which, by Bayes’s rule, combines $p(Y|X)$ and the distribution of stimuli $p(X)$.

¹⁰Haufe et al. (2014) do suggest a method to enable a neurophysiological interpretation of the parameters of linear decoding models.

out. The trained model is faced with the y_{vt} for a held-out stimulus condition in a particular voxel v , and the two sets of features for the correct stimulus condition and another unseen stimulus condition chosen at random. Next, the trained model makes a prediction for both stimuli, and the data point is assigned to the stimulus whose predicted activity is closer to the observed y_{vt} . By design, chance performance for the balanced binary reverse-inference task is 50%.

This is easily extended to from one voxel v to the entire brain. We compute the distance between the observed y and the predictions of the forward model as a weighted sum of the voxel-wise distances, where the weights depend on the classification accuracy of the individual voxels on the training set. Each voxel is weighed by the inverse of its rank when the per-voxel classification accuracies are sorted in decreasing order. Now, each stimulus is not represented by the original features, but instead by the weighted error in its forward model's predictions of neural activity.¹¹ (As before, voxel-wise accuracies are determined through cross-validation.) One consequence of the weighting scheme is that whole-brain reverse inference is highly accurate if there are only a few high-accuracy voxels. Of course, whole-brain classification can also be accurate even if no one voxel has high accuracy.

Validation sets and cross-validation. We evaluate both forward and reverse inferences with nested 10-fold cross-validation. 10% of the data is held for testing. We then use the remaining training set (90%) to compute the different estimates. For E2, we throw out 5 images on the boundaries of the training set and the test set to insure that there is no signal leakage from the training to the test set due to the slow decay of hemodynamic responses, causing unintended correlations. If we choose not to smooth the estimates, then we proceed as follows with the training set.

For ridge regression, we use generalized cross-validation [Golub, Heath and Wahba (1979)] to approximate leave-one-out cross-validation error for different λ_v values at each voxel. For the elastic net, we use the ten-fold cross-validation option provided in Friedman et al. (2010) to choose the regularization parameters. Finally, for SAE, the level of regularization is determined by the posterior mean variance of \mathbf{z}_v , since high variance corresponds to the model being able to choose the parameter freely, that is, low regularization. The posterior mean variance of \mathbf{z}_v is determined automatically by the Gibbs sampler.

If we choose to smooth the estimates, then in order to pick the smoothing parameter for every voxel and every estimator, we run a nested cross-validation loop. That is, within the 90% training portion of the data, 80% is randomly selected as “inner-fold training” data, and 10% is randomly selected as a validation set. The

¹¹This is analogous to the way support vector machines and other kernel classifiers expand the dimension of the feature space by computing many nonlinear functions of the features [Cristianini and Shawe-Taylor (2000)], and to the use of generative model likelihoods to define discriminative kernels [Jaakkola and Haussler (1999)].

inner-fold training is done exactly as in the previous paragraph. Smoothing parameters are then set using the average single-voxel classification accuracy on the validation set.

After training, the out-of-sample performance of both unsmoothed and smoothed estimators is reported using the testing set. Thus, the parameters never adapt to the testing set, and we report valid estimates of out-of-sample performance.

3. Results. Our main findings are as follows: using cross-validation to pick tuning parameters separately for each voxel:

1. Regularization offers small but real gains in forward prediction;
2. Regularization does not seem to offer improvement in reverse prediction at the individual voxel level, or whole-brain reverse inference;
3. All forms of regularization work about equally well for prediction;
4. Regularization succeeds in making parameter estimates more precise;
5. The spatial pattern of regularization is highly informative: voxels where unregularized OLS is least accurate are precisely the ones which are more heavily smoothed or regularized under cross-validation.

We explain these points in turn.

3.1. Prediction. To summarize, while all models and methods had some predictive ability in both experiments, none of them clearly dominated the others. Model checking, discussed in Appendix C of the Supplementary Article [Wehbe et al. (2015)], shows that this was not because the models were grossly inappropriate, though they are somewhat misspecified.

Forward inference. All our methods had nontrivial ability to do forward prediction in both experiments for some of the voxels, which should be the voxels that are implicated in visual processing (Figure 1). All methods of regularizing OLS, including spatial smoothing, led to generally small but significant improvements. The improvement is seen in the noisy voxels: the high RSS in those voxels is greatly reduced when shrinkage or smoothing is used, effectively driving the prediction in the noisy voxels to zero. Since results for both neighborhood- and ROI-based smoothing were nearly identical, we report only those for smoothing over ℓ_2 balls (however, see Section 4.) Combining smoothing with shrinkage did not help forward inference; if anything, it often made it worse than either alone (Figure 2).

Reverse inference. The effect of regularization on single-voxel reverse inference is ambiguous (see Appendix E of the Supplementary Article [Wehbe et al. (2015)]): accuracy goes up in some voxels and down in others, with no change over all. The classification accuracy of the good voxels varies much less across the




FIG. 1. Effect of regularization on out-of-sample normalized RSS (RSS/σ^2) for E1 (top) and E2 (bottom). For each of the plots, the OLS RSS/σ^2 (horizontal axis) is contrasted with the modified RSS/σ^2 after OLS smoothing for ridge, elastic net, or small area shrinkage (vertical axis). For both experiments, the four methods result in smaller RSS/σ^2 on average. Furthermore, for all the methods, the predicted activity in the bad voxels (i.e., voxels where RSS/σ^2 is larger than 1) is pushed toward zero. This is visible by the RSS/σ^2 values being reduced toward 1. In other words, shrinkage and smoothing are forcing the estimated parameters to be almost zero if the voxel is noisy and there is nothing that can be predicted. Note that the scales of the axes are different for the two experiments.

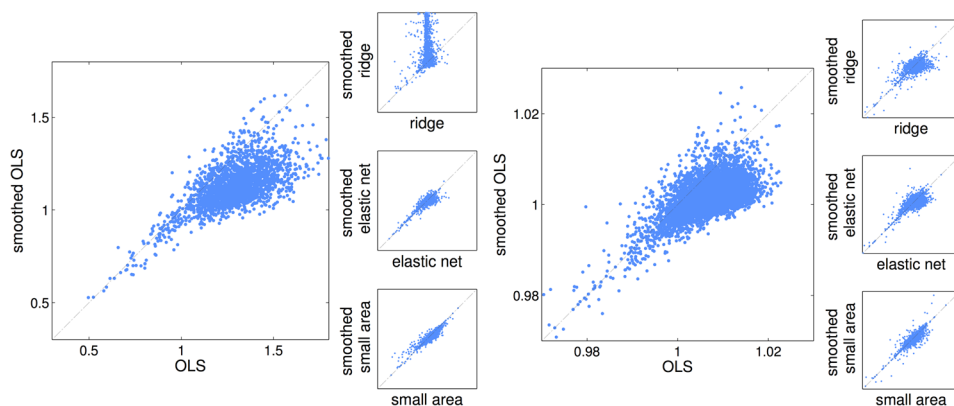


FIG. 2. Normalized RSS for unsmoothed and smoothed estimators for E1 (left) and E2 (right). The larger panels show voxel-wise normalized residuals (RSS/σ^2) for OLS before smoothing (horizontal axis) and after (vertical), showing the value of spatial smoothing for forward inference. The smaller panels consist of the same comparison for ridge regression (top), the elastic net (middle), and the small area model (bottom), showing that combining smoothing and shrinkage is if anything worse than shrinkage alone. The axes for the smaller panels have been omitted for clarity: they correspond to the larger panels axes.




FIG. 3. Whole-brain classification accuracy in experiments E1 (left) and E2 (right), averaging over subjects, for all combinations of estimators and smoothing. Regularization choice or the presence or absence of smoothing does not affect whole-brain classification accuracy.

different estimators than the accuracy of the bad voxels (see figures in Appendix E of the Supplementary Article [Wehbe et al. (2015)]).

Turning to whole-brain reverse inference, all methods, with and without smoothing, did much better than the chance rate of 50% in both experiments (Figure 3). However, the differences between methods are negligible, and certainly smaller than the fold-to-fold variability of cross-validation. This includes unregularized OLS.

All our methods predict equally well (up to experimental precision), which is surprising. We can rationalize the elastic net performing about as well as ridge regression on the grounds that the former extends the latter by adding an L_1 penalty, which might be unnecessary. Ridge regression is also linked to our hierarchical small area model via an approximation result (Appendix B of the Supplementary Article [Wehbe et al. (2015)]). However, such connections do not account for why all three forms of shrinkage perform about the same as smoothed OLS or unsmoothed OLS.

We do find a partial explanation from the way we do whole-brain classification (Section 2.5). Recall that we classify a pattern of activity as belonging to the stimulus whose predicted activity pattern is closest, but weight each voxel in this distance calculation depending on its individual classification accuracy. Thus, the weights are often dominated by a fairly small number of highly discriminative voxels. These voxels tend to also be ones where the forward model fits well, and cross-validation or Gibbs sampling selects little or no regularization for them. To support these claims, we examine the effects of regularization on the parameter estimates and the spatial patterns of regularization.

3.2. Regularization.

3.2.1. *Evidence of successful regularization.* In light of the surprising predictive equivalence of our different methods with each other and with OLS, it is worth verifying that our regularizers were in fact regularizing the estimation. From the




FIG. 4. *Left: Histograms of one regression coefficient's standard errors in E1, aggregating over all voxels, for both OLS and SAE. The sharp peaking of the latter histogram, to the left of the former, indicates that the typical parameter estimate has been made much more precise by the hierarchical model. Right: scatter-plot of the same standard errors. Most of the points fall below the diagonal, so most parameters are being estimated more precisely. Other coefficients and methods of regularization behaved similarly.*

standpoint of small area estimation theory, the crucial question is whether the parameter estimates are more precise than the “direct” estimates of OLS. That is, do the new estimates show smaller standard errors, or smaller coefficients of variation, than the direct estimates?

Results like in Figure 4 are typical across the coefficients and the regularizers. After regularization, most parameter estimates for most voxels had significantly smaller standard errors, sometimes much smaller. This was true even while using cross-validation to pick how much to regularize each voxel. (See Appendix D of the Supplementary Article [Wehbe et al. (2015)] for additional documentation.)

3.2.2. Spatial patterns of regularization and their implications. The strength of regularization chosen by cross-validation is not uniform or even random across the brain. It shows quite pronounced, and informative, spatial structure, closely connected to how well voxels predict without regularization.

Figures 5 and 6 depict the relationship between the degree of regularization imposed by our methods, and several measures of predictive accuracy. For two horizontal slices of the brain, these figures illustrate how classification accuracy varies, how strongly regularized each voxel is, and how well the regression model does in and out of sample. (The accuracy plot is omitted for the elastic net to show both penalty factors.) Appendix F of the Supplementary Article [Wehbe et al. (2015)] provides the corresponding plots for the entire brain. The plots provided are for two subjects, one from each experiment. The other subjects present a very similar pattern of correspondence between voxels with high performance and weak regularization.

As Figure 5(a)–(d) shows, there is an inverse relationship between predictive performance [subfigures (A) and (D)] and the degree of regularization [subfigures (B)] that was chosen by cross-validation, whether that is the smoothing radius

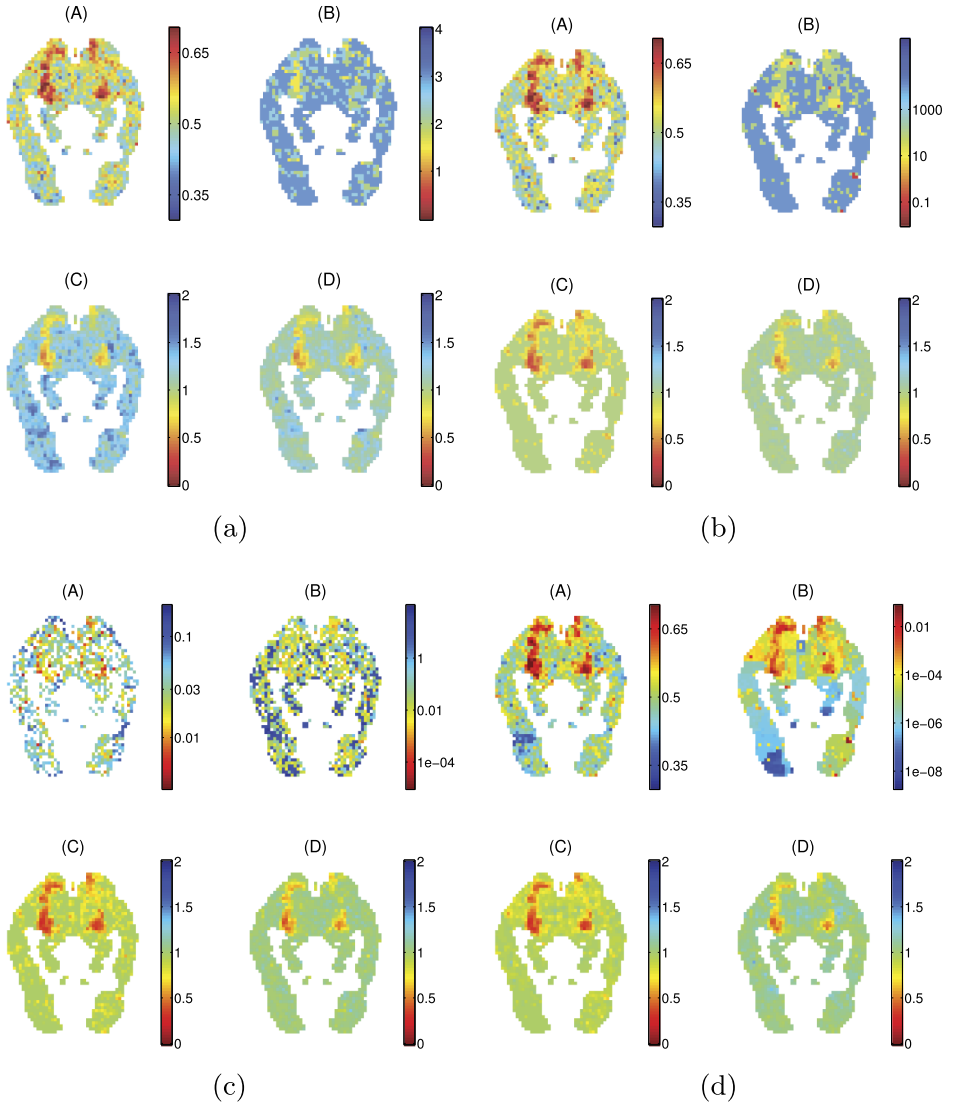


FIG. 5. Voxel-wise results for each method along one horizontal brain slice in experiment E1. Color schemes are flipped so that red always represents “good” and blue, “bad.” Note the similar patterns of classification accuracy in plots (a)-(A), (b)-(A), and (d)-(A). Also note how predictive performance [subfigures (A) and (D)] is inversely related to the degree of regularization in every case [the smoothing radius for OLS (a), the λ penalty for ridge (b), the λ_1 and λ_2 penalties for the elastic net (c), or the posterior mean variance in the small area model (d)—high variance means low regularization]. Finally, see that in many cases the in- and out-of-sample errors for “good” voxels are nearly the same. (a) OLS: (A) classification accuracy; (B) smoothing radius; (C), (D) normalized out-of-sample RSS pre- and post-smoothing. (b) Ridge: (A) classification accuracy; (B) λ parameter; (C), (D) normalized RSS in- and out-of-sample. (c) Elastic Net: (A) λ_1 (lasso penalty); (B) λ_2 (ridge penalty); (C), (D) normalized RSS in- and out-of-sample. (d) Small Area: (A) classification accuracy; (B) posterior mean variance of \mathbf{z}_v ; (C), (D) normalized RSS in- and out-of-sample.

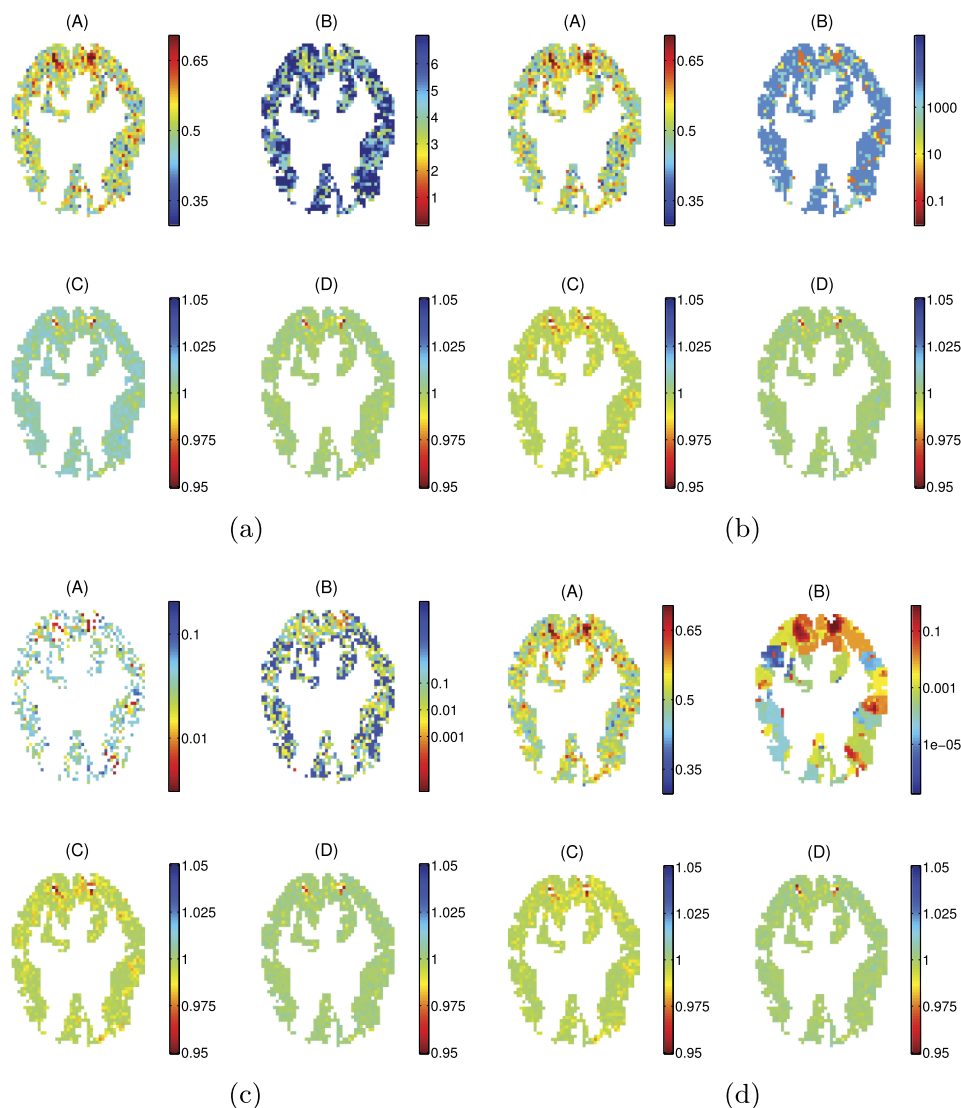


FIG. 6. Voxel-wise results for each method along one horizontal brain slice for experiment E2. Color schemes are flipped so that red always represents “good” and blue, “bad.” See Figure 5 for more details. As in Figure 5, predictive performance [subfigures (A) and (D)] is inversely related to the degree of regularization in every case. (a) OLS: (A) classification accuracy; (B) smoothing radius; (C), (D) normalized out-of-sample RSS pre- and post-smoothing. (b) Ridge: (A) classification accuracy; (B) λ parameter; (C), (D) normalized RSS in- and out-of-sample. (c) Elastic Net: (A) λ_1 (lasso penalty); (B) λ_2 (ridge penalty); (C), (D) normalized RSS in- and out-of-sample. (d) Small Area: (A) classification accuracy; (B) posterior mean variance of \mathbf{z}_v ; (C), (D) normalized RSS in- and out-of-sample.

for OLS (a), the λ penalty for ridge (b), the λ_1 and λ_2 penalties for the elastic net (c), or the small area model (d), where low regularization corresponds to a high variance parameter, that is, good voxels are allowed to pick their parameters freely. For the elastic net, good voxels have more lasso-like penalties, as they are voxels sensitive to *some* of the stimulus features. Smoothing acts as a regularizer for OLS, as seen by the reduced prediction in the bad voxels from subfigure (a)-(C) to subfigure (a)-(D). Thus, voxels with stronger signals (as reflected by higher accuracy) needed less regularization. Voxels with high accuracy [Figure 5(a), (b) and (d), part A] and especially voxels with low prediction error [subfigures (D)] are sparse and spatially clustered. Other voxels are by comparison noisy and more heavily regularized [subfigures (B)].

The correspondence between good classification accuracy and weak regularization explains the single voxel accuracy results mentioned in Section 3.1 and in Appendix E of the Supplementary Article [Wehbe et al. (2015)]. In good voxels, classification accuracy is not significantly affected by regularization since the penalty parameter is weak. In the bad voxels, the strong regularization forces the model to learn near-zero weights, and the leftover noise has a “random” effect on the single voxel classification accuracy, sometimes resulting in slight improvement, and sometimes in slight decrease.

For E1, the predictive voxels are clustered in the occipital cortex, which is well known to be heavily involved in visual processing [Shepherd (1994)]. For E2, the predictive voxels involve a smaller part of the occipital cortex, as well as some small clusters of voxels in more anterior regions associated with language comprehension (such as the left temporal lobe).

4. Discussion.

4.1. *Ridge and SAE.* We have shown that different forms of regularization predict about equally well. Moreover, they give similar parameter estimates, especially the SAE model of (4) and ridge regression. As already explained in Section 2.3, the marginal prior distribution of β_v is an inverse-gamma variance mixture of Gaussians, which is a t -distribution, where $\beta_v \sim t$. With even a moderate number of degrees of freedom in the t , the marginal prior on β_v is quite close to being Gaussian (Appendix B of the Supplementary Article [Wehbe et al. (2015)]). Similarly, the marginal prior on u_i is also a t -distribution. Since β_v and $u_{A(v)}$ are independent a priori, the prior on z_v is approximately Gaussian. Since the posterior mode under a Gaussian prior matches ridge regression, the z_v estimated from (4) will be close to the ridge regression estimates. We have not been able to find this approximation result in the literature, but suspect it is a rediscovery.

When we simulate from the SAE model, estimating that model shows better forward prediction than OLS or even ridge regression (Appendix C.1 of the Supplementary Article [Wehbe et al. (2015)]). The difference between SAE and ridge is small but systematic and significant. However, when the surrogate data from

the simulations is re-estimated with erroneous assignments of voxels to ROIs, the advantage of the SAE model over ridge regression vanishes. It *may* be that this is the way in which the SAE is misspecified, suggesting that a better choice of ROIs would lead to superior prediction. However, we have not been able to rule out other possible misspecifications.

4.2. *Computational costs.* While our four methods perform very similarly statistically, their computational costs differ by orders of magnitude (Table 1). Smoothed OLS and ridge stand out as the most attractive methods, with ridge pulling ahead due to its better behaved out-of-sample residuals.

Our simple and generic HB model is misspecified, not very firmly grounded in biology and, as Table 1 shows, computationally very costly. With considerable attention to the biology, well-specified models and priors might be crafted for specific applications, though at even greater computational expense. Due to this, we do not advocate the Bayesian approach, unless it could be combined with some way of *quickly* approximating posterior distributions, for example, variational methods [Broderick et al. (2013), Wainwright and Jordan (2008)] or consensus MCMC [Neiswanger, Wang and Xing (2013), Scott, Blocker and Bonassi (2013)]. Such extensions are beyond the scope of this paper.

4.3. *The detectability assumption.* As mentioned in Section 1.1, it is common in brain-imaging studies to do a separate regression for each voxel on the stimuli, and presume that only voxels with significant regression coefficients are involved in processing the stimulus. This assumption appears to have no neurobiological basis; we call it the “detectability assumption.” In practice, neuroscientists recognize this leads to some number of errors, both false positives and negatives. However, they often presume that these errors are random rather than systematic. Under the detectability assumption, methods of regularization might plausibly be seen as reducing the rate of false positives. If true, this is an important advantage

TABLE 1

Running times of the various procedures on the E1 data, using 8 Intel Xenon CPU E5-2660 0 cores (at 2.2 GHz), sharing 128 GB of RAM. Gibbs sampling for the SAE model was parallelized over the cores

	cpu time per fold per subject	Clock time per fold per subject	Total cpu time (with nested CV)
OLS	<1 s	<1 s	<1 min
Ridge	55 s	4 s	7.5 h
Elastic net	3120 s	390 s	429 h
Small area	5540 s	740 s	762 h
Smoothing, nested CV	40 s	20 s	5.5 h

for regularized estimates, even if they predicted no better than OLS. The spatial pattern of regularization, under this assumption, is an indicator of which voxels are involved in processing the stimulus features. This indication is strengthened by the similarity of these patterns under different methods of regularization (Section 3.2.2).

Since constant-bandwidth smoothing the data spatially is a common fMRI pre-processing step, usually followed by using OLS, it can be argued that existing analyses are already doing some regularization. However, as Figure 1 shows, our shrinkage methods reduce prediction error in the noisy voxels somewhat more than does smoothing OLS. Moreover, even if one preferred to use smoothed OLS rather than shrinkage, we have shown that good voxels do not need to be regularized as much as noisy voxels. Therefore, the current approach can be improved by choosing the smoothing parameter at every voxel.

Despite its ubiquity, it is hard to support the detectability assumption neurobiologically. A crucial component of it is the presumption that the hemodynamic response is systematically related to information processing. While it is true that increased spiking rates within a voxel will lead to a hemodynamic response, animal experiments show that neural information can be encoded in the time intervals between spikes rather than the spiking rate [Rieke et al. (1999)], and in the coordination of spiking across neurons, which may lie within the same voxel or be widely distributed across the brain [Abbott and Sejnowski (1998), Ballard, Zhang and Rao (2002), Engel, Fries and Singer (2001), Fries (2009)]. Further, many neural circuits work by inhibiting other neurons, and increasing inhibition may either increase or decrease energy demands and so hemodynamic responses, depending on fine-grained anatomical and physiological details [Logothetis (2008)].

Such considerations undermine the link between changes in local spiking rates, energy use by neurons, hemodynamic response, and actual neural computation or information-processing. If information is conveyed by timing, conveyed by synchrony, distributed across large spatial volumes, or works through a balance between excitation and inhibition, then much neural computation might be invisible in the hemodynamic signal which fMRI measures. This leads to systematic false negatives which are inevitable when working with fMRI.

Another difficulty with the detectability assumption is that it presumes that when a voxel's hemodynamic signal does respond to stimulus features, the regression coefficients are always relatively large. Usually, "relatively large" amounts to "statistically significant." This all runs together with the absolute magnitude of the regression coefficients, the sample size (including the duration of each experiment and the number of subjects), the variance of the stimulus features, and the extent to which the features are correlated with each other. With larger samples and higher-variance, less-correlated features, smaller regression coefficients become significant. That is, there is more power to detect small coefficients. The negative inference that certain voxels are not involved in processing stimulus features presumes that feature-sensitive voxels have coefficients large enough that the experiment has substantial power to detect them.

Regularization does not necessarily improve this situation. While it does avoid making a hard-and-fast decision based on significance, it is still true that the optimal amount of regularization generally declines with the sample size. Moreover, small coefficients, being hard to estimate, could be heavily penalized under cross-validation. Thus, while regularization may reduce the number of false positives, this may be more than counterbalanced by an increase in false negatives, unless all nonzero coefficients are fairly large.

In summary, the detectability assumption contains two parts—that neural information-processing always shows up in the hemodynamic response, and that associated regression coefficients are always either zero or large—which our current knowledge of neurobiology does not support. Nonetheless, without a feasible replacement, we hesitate to reject outright an assumption embraced by so much of the neuroscientific community.

4.4. Conclusion. Our main finding is that how we regularize, whether using shrinkage and or smoothing, is much less important for prediction than regularizing somehow (Sections 3.1 and 3.2.2). All regularization methods considered (ridge, elastic net, the small area HB model, and smoothed OLS) improved forward and backward predictions about equally. When we allowed the degree of regularization to vary across the brain, voxels with strong signals receive little regularization, while more noisy voxels are heavily regularized [Figure 5(a)–(d)]. Furthermore, very similar patterns emerged from all methods. Since the methods are similar predictively, we favor ridge and smoothed OLS on computational grounds. Ridge regression is already widely used, but smoothed OLS should be added to the fMRI toolkit.

None of our methods were designed for fMRI problems and none were informed by a deep understanding of the physics of measuring hemodynamic response or any type of neuropsychological model. However, we hope that better predictions can be obtained through regularization methods that express neurologically relevant forms of smoothness, sparsity, and similarity, rather than just being “off the shelf” priors or penalties. We do not mean to be dogmatic about whether neurobiological constraints should be expressed as objective functions or as stochastic processes, though we suspect that a penalized optimization approach is more computationally tractable than a Bayesian approach. It is hard to imagine a biologically sound Bayesian model leading to conjugate priors. If a Bayesian approach is taken, it should be biologically and neurologically sound and computationally efficient. Posterior approximation methods should play a crucial role, and we leave this for future exploration. Whether priors or penalties, the regularizers of the future must be neural models.

Acknowledgments. We would like to thank the referees and the Associate Editor for comments that led to major improvements of this paper. The views in this paper are of the authors alone and not of the funding agencies.

SUPPLEMENTARY MATERIAL

Supplementary Article: Appendix for “Regularized brain reading with shrinkage and smoothing” (DOI: [10.1214/15-AOAS837SUPP](https://doi.org/10.1214/15-AOAS837SUPP); .pdf). This supplement consists of six parts. It offers more details about: (A) our Small Area model and Gibbs sampler, (B) the Marginal Prior of the SAE Model, (C) model checking, (D) the effect of regularization on variability, and (E) the effect of smoothing and regularization on single voxel accuracy, as well as (F) whole brain plots of the experimental results that are portrayed in Figures 5 and 6 for a single slice.

REFERENCES

- ABBOTT, L. F. and SEJNOWSKI, T. J. EDs. (1998). *Neural Codes and Distributed Representations: Foundations of Neural Computation*. MIT Press, Cambridge, MA.
- ASHBURNER, J., BARNES, G., CHEN, C.-C., DAUNIZEAU, J., FLANDIN, G., FRISTON, K., KIEBEL, S., KILNER, J., LITVAK, V., MORAN, R., PENNY, W., ROSA, M., STEPHAN, K., GITELMAN, D., HENSON, R., HUTTON, C., GLAUCHE, V., MATTOU, J. and PHILLIPS, C. (2008). SPM8 Manual. Functional Imaging Laboratory, Wellcome Trust Centre for Neuroimaging, Institute of Neurology, UCL.
- ASHBY, F. G. (2011). *Statistical Analysis of FMRI Data*. MIT Press, Cambridge, MA.
- BALLARD, D. H., ZHANG, Z. and RAO, R. P. N. (2002). Distributed synchrony: A probabilistic model of neural signalling. In *Probabilistic Models of the Brain: Perception and Neural Function* (R. P. N. Rao, B. A. Olshausen and M. S. Lewicki, eds.). *Neural Information Processing Series* 273–284. MIT Press, Cambridge, MA.
- BRODERICK, T., BOYD, N., WIBISONO, A., WILSON, A. C. and JORDAN, M. I. (2013). Streaming variational Bayes. In *Advances in Neural Information Processing Systems 26 [NIPS 2013]* (C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani and K. Q. Weinberger, eds.) 1727–1735.
- CRISTIANINI, N. and SHAWE-TAYLOR, J. (2000). *An Introduction to Support Vector Machines: And Other Kernel-Based Learning Methods*. Cambridge Univ. Press, Cambridge, MA.
- DATTA, G. S., GHOSH, M., STEORTS, R. and MAPLES, J. (2011). Bayesian benchmarking with applications to small area estimation. *TEST* **20** 574–588. [MR2864715](https://doi.org/10.1007/s11337-011-9111-1)
- ENGEL, A. K., FRIES, P. and SINGER, W. (2001). Dynamic predictions: Oscillations and synchrony in top-down processing. *Nat. Rev., Neurosci.* **2** 704–716.
- FRIEDMAN, J., HASTIE, T., TIBSHIRANI, R. and JIANG, H. (2010). Glmnet for Matlab. Statistics Department, Stanford Univ., Stanford.
- FRIES, P. (2009). Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annu. Rev. Neurosci.* **32** 209–224.
- FRISTON, K. J., ROTHSSTEIN, P., GENG, J. J., STERZER, P. and HENSON, R. N. (2010). A critique of functional localizers. In *Foundational Issues in Human Brain Mapping* (S. J. Hanson and M. Buzl, eds.) 3–24. MIT Press, Cambridge, MA.
- GENOVESE, C. R. (2000). A Bayesian time-course model for functional magnetic resonance imaging data. *J. Amer. Statist. Assoc.* **95** 691–703.
- GOLUB, G. H., HEATH, M. and WAHBA, G. (1979). Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics* **21** 215–223. [MR0533250](https://doi.org/10.1080/00141801.1979.11518609)
- HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2001). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, New York. [MR1851606](https://doi.org/10.1007/978-1-4939-9723-1)
- HAUFE, S., MEINECKE, F., GÖRGEN, K., DÄHNE, S., HAYNES, J.-D., BLANKERTZ, B. and BIESSMANN, F. (2014). On the interpretation of weight vectors of linear models in multivariate neuroimaging. *NeuroImage* **87** 96–110.

- HOERL, A. E. and KENNARD, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* **12** 55–67.
- JAAKKOLA, T. and HAUSSLER, D. (1999). Exploiting generative models in discriminative classifiers. In *Advances in Neural Information Processing Systems 11 [NIPS 1998]* (M. J. Kearns, S. A. Solla and D. A. Cohn, eds.) 487–493. MIT Press, Cambridge, MA.
- KYUNG, M., GILL, J., GHOSH, M. and CASELLA, G. (2010). Penalized regression, standard errors, and Bayesian lassos. *Bayesian Anal.* **5** 369–411. [MR2719657](#)
- LEE, K.-J., JONES, G. L., CAFFO, B. S. and BASSETT, S. S. (2011). Spatial Bayesian Variable Selection Models on Functional Magnetic Resonance Imaging Time-Series Data. Preprint.
- LOGOTHETIS, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature* **453** 869–878.
- LOUIS, T. A. (1984). Estimating a population of parameter values using Bayes and empirical Bayes methods. *J. Amer. Statist. Assoc.* **79** 393–398. [MR0755093](#)
- MITCHELL, T. M., SHINKAREVA, S. V., CARLSON, A., CHANG, K.-M., MALAVE, V. L., MASON, R. A. and JUST, M. A. (2008). Predicting human brain activity associated with the meanings of nouns. *Science* **320** 1191–1195.
- NASELARIS, T., KAY, K. N., NISHIMOTO, S. and GALLANT, J. L. (2011). Encoding and decoding in fMRI. *NeuroImage* **56** 400–410.
- NEISWANGER, W., WANG, C. and XING, E. (2013). Asymptotically exact, Embarrassingly Parallel MCMC. Preprint. Available at [arXiv:1311.4780](#).
- NORMAN, K. A., POLYN, S. M., DETRE, G. J. and HAXBY, J. V. (2006). Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends Cogn. Sci.* **10** 424–430.
- PALATUCCI, M., POMERLEAU, D., HINTON, G. E. and MITCHELL, T. M. (2009). Zero-shot learning with semantic output codes. In *Advances in Neural Information Processing Systems 22 [NIPS 2009]* (Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams and A. Culotta, eds.) 1410–1418. MIT Press, Cambridge, MA.
- PARK, M., KOYEJO, O., GHOSH, J., POLDRACK, R. A. and PILLOW, J. W. (2013). Bayesian structure learning for functional neuroimaging. In *16th International Conference on Artificial Intelligence and Statistics* (C. M. Carlvah and P. Ravikumar, eds.) 489–497.
- PEREIRA, F., MITCHELL, T. and BOTVINICK, M. (2009). Machine learning classifiers and fMRI: A tutorial overview. *NeuroImage* **45** S199–S209.
- PEFFERMANN, D. (2013). New important developments in small area estimation. *Statist. Sci.* **28** 40–68. [MR3075338](#)
- POLDRACK, R. A. (2008). The role of fMRI in cognitive neuroscience: Where do we stand? *Curr. Opin. Neurobiol.* **18** 223–227.
- RAO, J. N. K. (2003). *Small Area Estimation*. Wiley, Hoboken, NJ. [MR1953089](#)
- RIEKE, F., WARLAND, D., DE RUYTER VAN STEVENINCK, R. and BIALEK, W. (1999). *Spikes: Exploring the Neural Code*. MIT Press, Cambridge, MA. [MR1983010](#)
- ROWLING, J. K. (2012). *Harry Potter and the Sorcerer's Stone*. Pottermore Limited, London.
- SCOTT, S. L., BLOCKER, A. W. and BONASSI, F. V. (2013). Bayes and big data: The consensus Monte Carlo algorithm. Presented at the “EFaBBayes 250” conference, 16 December 2013, Duke Univ.
- SHEPHERD, G. M. (1994). *Neurobiology*, 3rd ed. Oxford Univ. Press, London.
- SMITH, S. M. (2004). Overview of fMRI analysis. *Br. J. Radiol.* **77** S167–S175.
- SUDRE, G., POMERLEAU, D., PALATUCCI, M., WEHBE, L., FYSHE, A., SALMELIN, R. and MITCHELL, T. (2012). Tracking neural coding of perceptual and semantic features of concrete nouns. *NeuroImage* **62** 451–463.
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. [MR1379242](#)

- TZOURIO-MAZOYER, N., LANDEAU, B., PAPATHANASSIOU, D., CRIVELLO, F., ETARD, O., DELCROIX, N., MAZOYER, B. and JOLIOT, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage* **15** 273–289.
- VON LUXBURG, U. (2007). A tutorial on spectral clustering. *Stat. Comput.* **17** 395–416. [MR2409803](#)
- WAINWRIGHT, M. J. and JORDAN, M. I. (2008). Graphical models, exponential families, and variational inference. *Found. Trends Mach. Learn.* **1** 1–305.
- WEHBE, L., MURPHY, B., TALUKDAR, P., FYSHE, A., RAMDAS, A. and MITCHELL, T. (2014). Simultaneously uncovering the patterns of brain regions involved in different story reading sub-processes. *PLoS ONE* **9** e112575.
- WEHBE, L., RAMDAS, A., STEORTS, R. C. and SHALIZI, C. R. (2015). Supplement to “Regularized brain reading with shrinkage and smoothing.” DOI:[10.1214/15-AOAS837SUPP](#).
- YARKONI, T., POLDRACK, R. A., NICHOLS, T. E., VAN ESSEN, D. C. and WAGER, T. D. (2011). Large-scale automated synthesis of human functional neuroimaging data. *Nature Methods* **8** 665–670.
- ZOU, H. and HASTIE, T. (2005). Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **67** 301–320. [MR2137327](#)

L. WEHBE
HELLEN WILLS NEUROSCIENCE INSTITUTE
UNIVERSITY OF CALIFORNIA, BERKELEY
BERKELEY, CALIFORNIA 94720
USA
E-MAIL: lwehbe@berkeley.edu

R. C. STEORTS
DEPARTMENT OF STATISTICAL SCIENCE
DUKE UNIVERSITY
DURHAM, NORTH CAROLINA 27708
USA
E-MAIL: beka@stat.duke.edu

A. RAMDAS
EECS DEPARTMENT
UNIVERSITY OF CALIFORNIA, BERKELEY
BERKELEY, CALIFORNIA 94720
USA
E-MAIL: aramdas@cs.berkeley.edu

C. R. SHALIZI
DEPARTMENT OF STATISTICS
CARNEGIE MELLON UNIVERSITY
5000 FORBES AVE.
PITTSBURGH, PENNSYLVANIA 15213
USA
E-MAIL: cshalizi@cmu.edu