



# Interpreting encoding and decoding models

Nikolaus Kriegeskorte<sup>1</sup> and Pamela K Douglas<sup>2</sup>

Encoding and decoding models are widely used in systems, cognitive, and computational neuroscience to make sense of brain-activity data. However, the interpretation of their results requires care. Decoding models can help reveal whether particular information is present in a brain region in a format the decoder can exploit. Encoding models make comprehensive predictions about representational spaces. In the context of sensory experiments, where stimuli are experimentally controlled, encoding models enable us to test and compare brain-computational theories. Encoding and decoding models typically include fitted linear-model components. Sometimes the weights of the fitted linear combinations are interpreted as reflecting, in an encoding model, the contribution of different sensory features to the representation or, in a decoding model, the contribution of different measured brain responses to a decoded feature. Such interpretations can be problematic when the predictor variables or their noise components are correlated and when priors (or penalties) are used to regularize the fit. Encoding and decoding models are evaluated in terms of their generalization performance. The correct interpretation depends on the level of generalization a model achieves (e.g. to new response measurements for the same stimuli, to new stimuli from the same population, or to stimuli from a different population). Significant decoding or encoding performance of a single model (at whatever level of generality) does not provide strong constraints for theory. Many models must be tested and inferentially compared for analyses to drive theoretical progress.

## Addresses

<sup>1</sup> Department of Psychology, Department of Neuroscience, Department of Electrical Engineering, Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY, United States

<sup>2</sup> Center for Cognitive Neuroscience, University of California, Los Angeles, CA, United States

Corresponding author: Kriegeskorte, Nikolaus ([nk2765@columbia.edu](mailto:nk2765@columbia.edu))

Current Opinion in Neurobiology 2019, 55:167–179

This review comes from a themed issue on **Machine learning, big data, and neuroscience**

Edited by **Jonathan Pillow** and **Maneesh Sahani**

For a complete overview see the [Issue](#) and the [Editorial](#)

Available online 28th April 2019

<https://doi.org/10.1016/j.conb.2019.04.002>

0959-4388/© 2019 Elsevier Ltd. All rights reserved.

## Encoding and decoding: concepts with caveats

The notions of encoding and decoding are rooted in the view that the brain processes information. Information about the world enters our brains through the senses, is processed and potentially stored, and informs our behavior over a large range of time scales. To understand brain function, then, we must understand the information processing: *what information* is processed, *in what format* it is coded in neural activity, and *how it is re-coded* across stages of processing, so as to ultimately contribute to behavior.

The view of the brain as an information-processing device (or, equivalently, a computational device) is closely linked to the concept of *representation*. A pattern of neural activity *represents* information about the world if it *serves the function* to convey this information to downstream regions, which use it to produce successful behavior [1].

When we talk about encoding and decoding, we focus on a particular brain region *X* whose function we are trying to understand. To simplify the problem, we divide the process into the encoder (which produces the code in region *X* from the sensory signals) and the decoder (which uses the code in region *X* to enable successful behavior). This bipartite division often provides a useful simplification. However, we have to consider three caveats to avoid confusion.

First, encoder and decoder do not inherently differ with respect to the nature of the processing. Where the encoder ends and the decoder begins depends on our region of interest *X*. When we move our focus to the next stage of representation in region *Y*, the processing between *X* and *Y*, which was part of the decoder, becomes part of the encoder. Whether a particular processing step in the brain is to be considered encoding or decoding, thus, is in the eye of the beholder.

Second, our region of interest *X* may not be part of all causal paths from input to output. The division about region *X* into encoder and decoder, then, misses a portion of the causal graph. Brain regions do not in general form chains of processing stages without skipping connections or recurrent signaling. The primate visual hierarchy is a case in point, where cortical areas interact in a network with about a third of all possible pairwise inter-area connections [2,3]. Encoding and decoding models are

nevertheless useful in this context, providing a partial view of the encoded information and its format.

Third, the terms encoding and decoding suggest that each is the inverse of the other: The encoder maps stimuli to brain responses; the decoder maps brain responses to stimuli. However, if we would like to interpret our encoding and decoding models as models of brain computation, then both encoder and decoder ought to operate in the causal direction. The decoder should not map *back* to stimuli, but *on* to motor representations. An encoding model might predict neural activity elicited by images [4,5,6]. A decoding model might attempt to predict category labels from neural activity, thus explicating information that might be reflected in a behavioral response [7].

The notion that encoding is followed by its inverse begs the question why the original input is not simply copied. A possible motivation for a sequence of an encoder and its inverse is to temporarily compress the information, for example, for passage through a bottleneck such as the optic nerve. A more general notion is that information is *re-coded* through a sequence of stages, possibly compressing it for more efficient representation [8], but also discarding unneeded information and converting needed information into a format that enables it to be exploited to control behavior [9,10]. When decoding models are conceptualized as inverse encoders (mapping back to the stimuli, rather than on to downstream representations), they cannot be interpreted as process models of brain

function. However, they can still be useful tools, revealing information present in a brain region and giving us clues about the format of the code. For a careful discussion of the interpretation of causal and anti-causal encoding and decoding models, see Weichwald *et al.* [11\*].

Encoding and decoding are important concepts in both theoretical and experimental neuroscience. This paper addresses experimentalists and focuses on the interpretation of empirical results obtained by fitting encoding and decoding models to brain-activity data.

### Decoding models: revealing information and its format

Decoding is sometimes described as ‘reading the brain’ or ‘cracking the neural code’ [12–14]. The underlying concept is that of decoding as inverse encoding, where the goal is not to model brain information processing, but to reveal the content of the code. The sense of revealing a mystery and gaining insight into hidden information drives the imagination of scientists and lay people alike.

The intuition that something important has been learned when we can decode some perceptual or higher cognitive content from brain activity is correct. However, ‘decrypting’ the brain should not be equated with understanding its computational mechanism. Decoding reveals the products, not the process of brain computation (Box 1). However, it is a useful tool for testing whether a brain region contains a particular kind of information in a particular format [12–14,15\*,16,17\*,18–29].

#### Box 1 Decoding models: benefits and limitations

Decoding provides an intuitive and compelling demonstration of the presence of information in a brain region. Decoding models bring several benefits:

- They enable researchers to look into the regions and reveal whether particular information is encoded in the responses. Encoded information might be used by downstream neurons (i.e. the encoding might serve as a representation). Training and testing decoders have helped move the focus from single-neuron activity and regional-mean activation to the information being processed, making analyses more relevant to the computational function in question.
- They enable cognitive neuroscience to exploit the fine-grained multivariate information present in modern imaging data. Before the advent of decoding, dominant brain mapping techniques required smoothing of the data, treating fine-grained information as noise. The multivariate modelling of local patterns also brought locally multivariate noise models to the literature, which greatly enhance the information that can be drawn from neuroimaging data.
- Decoding analyses use independent test data to assess the presence of information. Demonstrating significant information with independent test data ensures that violations of model assumptions lead to conservative errors (making significant results less likely). This is an advantage over methods (such as univariate and multivariate linear regression) that use distributional assumptions rather than independent test sets to perform statistical inference.

Decoding models are limited in several ways:

- They cannot in general be interpreted as models of brain-information processing. In the context of sensory systems, decoding models operate in the wrong direction, capturing the inverse of the brain’s processing of the stimuli. If the decoded variable is a behavioral response, then direction of information flow matches between model and brain. However, such applications are rare. Moreover, the most successful applications of decoding in the literature rely on linear decoders, which are useful to reveal explicit information, but cannot capture the complex nonlinear computations we wish to understand.
- Decoder weight maps are difficult to interpret (Figure 3) because voxels uninformative by themselves can receive large weights when they help cancel noise and because the weights are codetermined by the data and the prior, and the fitting procedure might select one but not the other of two essentially equally informative voxels.
- Decoding results provide only weak constraints on computational mechanisms. Knowing that a particular kind of information is present provides some evidence for and against alternative theories about what the region does. However, it does not exhaustively characterize the representation, let alone reveal the underlying computations.

A simple example of a decoding model is a linear classifier [30] that takes a measured brain-activity pattern as input and outputs a class label (say 'cat' and 'dog'), revealing which of two stimuli (say a particular image of a cat and a particular image of a dog) elicited the response pattern (Figure 1). A linear classifier may achieve this by computing a weighted sum of the measured responses and imposing a threshold to decide which label to return (see Ref. [21] for an introduction in the context of neuroimaging). The weights are set to maximize the accuracy of the decoding on a training data set. If decoding succeeds reliably on a test data set, then the region must contain information about the decoded variable (the binary stimulus distinction here).

### Decoding provides a statistically advantageous way of testing for stimulus information

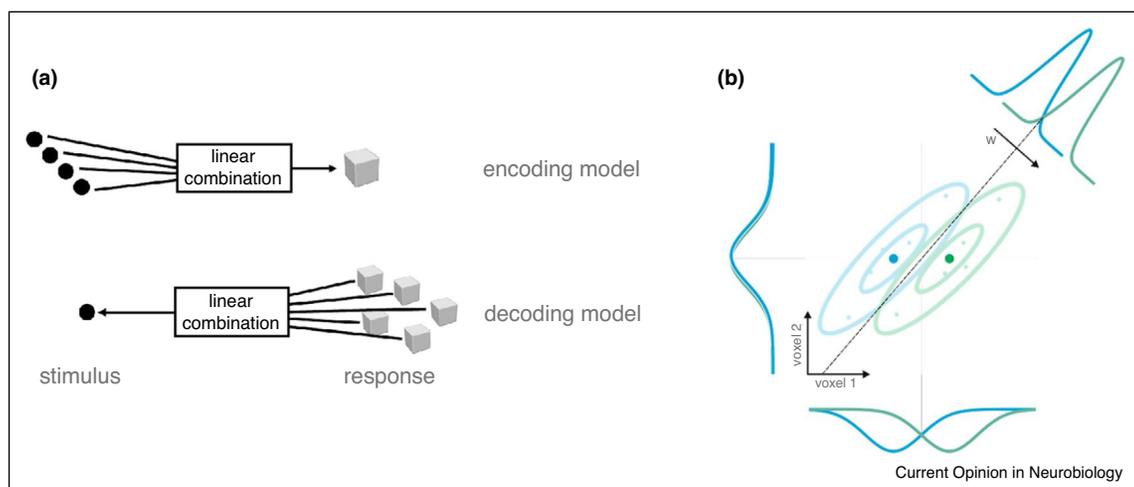
In the two-stimulus example, all the linear decoder shows is that the two images elicit distinct response patterns. This means that there is mutual information between stimulus and response. To demonstrate mutual information, we could have used an encoding model [31] or a multivariate test of difference between the response patterns elicited by the two stimuli (e.g. multivariate analysis of variance), instead of a decoding model. However, a univariate encoding model would in general have less sensitivity, because it does not account for the noise correlations between different response channels [32]. A multivariate analysis of variance would account for the noise correlations, but might have less specificity. In fact,

it might fail to control false-positives at the nominal level if its assumptions of multivariate normality did not hold (as is often the case), making it invalid as a statistical test [24]. Decoding provides a natural approach to modelling the noise correlations (e.g. using a multivariate normal model as in the Fisher linear discriminant), without relying on the model assumptions for the validity of the test: Violations of the decoding model assumptions will hurt decoding performance. We, thus, err on the safe side of concluding that there is no information about the stimuli in the responses. In sum, it is not the direction of the decoding model ('brain reading') that makes it a compelling test for information, but the statistical nature of the problem (noise correlations) and the fact that decoders are tested on independent data.

### Linear decodability indicates 'explicit' information

For decoding to succeed, the information must be present in the brain region in a format that the decoder can exploit. Linear decoders, the most widely used class, require that the distributions of patterns be linearly separable to some extent. This is a weakness in that we might fail to detect information encoded in a more complex format. However, it is a strength in that it provides clues to the format of the information we do detect. The simpler the decoder, the more focused its sensitivity will be. From the perspective of understanding the brain computations, it is attractive to use decoding operations that single neurons can implement. These include linear readout, but also simple nonlinear forms

Figure 1



Linear encoding and decoding models.

**(a)** Encoding and decoding model the relationship between stimuli and responses in opposite directions. An encoding model (top) predicts brain responses as a linear combination of stimulus properties (black circles). A decoding model (bottom) predicts stimulus properties as a linear combination of brain responses. **(b)** Example of a linear decoding model using a 2-dimensional feature space consisting of two voxels. Voxel 1 contains relevant information about which of two classes (green, blue) the stimulus belongs to. Voxel 2 contains no information about the stimulus class. The two dimensions jointly define the linear discriminatory boundary. Note that the weights assigned to each voxel are defined by the vector  $w$ , which is orthogonal to the decision boundary. Because the noise is correlated between the voxels, a linear decoder will assign significant negative weight to voxel 2, using this voxel (which contains only noise) to cancel the noise in the voxel which contains signal. As a result, interpreting the absolute weights of linear decoders requires care and additional analyses.

of readout such as radial basis function decoding [33]. Linear decodability indicates that a downstream neuron receiving input from a sufficient portion of the pattern, could read out the information in question [34]. Information amenable to direct readout by downstream neurons is sometimes referred to as ‘explicit’ in the code [24,35,36].

#### **The level of generalization beyond the training set must be considered when interpreting a decoding result**

Fitting a model always poses the risk of overfitting, that is, of optimizing the fit to the training data at the expense of predictive performance on independent data. Overfitting can lead to high decoding accuracy on the training set, even if the response patterns contain no information about the stimulus. Decoders, therefore, need to be tested for generalization to independent data [37,38]. In our example, we might test the decoder on an independent set of response measurements for the same two particular images of a cat and a dog. If decoding accuracy on this independent test set is significant, we can reject the null hypothesis that the response patterns contain no information about the stimulus [21,22].

However, detecting information about which of two images has been presented tells us almost nothing about the nature of the code. The two images must have distinct response patterns in the retina and V1 (low-level representations) as well as in the visuo-semantic regions of the ventral stream (high-level representations). We would, therefore, expect a linear decoder to work on new measurements for the same images in any of these regions. This reflects the fact that all the regions contain image information. In the retina, for example, we expect the two images to elicit distinct response patterns, while the manifolds of response patterns corresponding to the two categories are hopelessly entangled [34,40,41].

Given responses to just two images, we can demonstrate the presence of information, but have no empirical basis for characterizing the nature of the code (see Ref. [42] for a study limited by this drawback). In order to learn whether the region we are decoding from contains a low-level image encoding or a high-level categorical encoding, we can train the decoder on one set of cat and dog images and test it on another set of images of different cats and dogs [24,43,44].

To support the interpretation that ‘cats’ and ‘dogs’ are linearly separable in the representation (rather than the weaker claim that there is image information), it is not sufficient to increase the number of particular images of cats and dogs, while training and testing on the same images. The linear decoder has many parameters (one for each response channel) and is expected to overfit even to a larger set of particular images. Even for the retinal representation, we, therefore, expect a cat/dog decoder to generalize to new measurements performed on the

same images. We must test the linear decoder for generalization to different cats and dogs.

Note, however, that interpreting linear decodability as linear separability of the two classes in the neuronal representational space would further require the decoding accuracy to be so high that errors can be attributed to the measurement noise rather than the neural representation. In practice, we typically face ambiguity. For example, decoding accuracy may be significantly above chance, but far from perfect. This indicates that the code contains some linearly decodable information, but claims of linear separability may be difficult to evaluate as it would require attributing the substantial proportion of errors to limitations of the measurements (noise and subsampling), rather than to a lack of linear separability of the neuronal activity patterns.

If we managed to show that cats and dogs fall in linearly separable regions in representational space, there would still be a question about the nature of the features that support the encoding. There may be many different visual features that can be computed from images and that lend themselves to separating cats from dogs (e.g. the shape of the ears or the shape of the eyes). Revealing which particular features the encoding in a given brain region employs would require further experiments. For example, we could design artificial stimuli that contain subsets of the distinctive features of cats and dogs. We could then test whether a linear decoder trained on responses to natural images of cats and dogs generalizes to responses to the artificial stimuli. Successful decoding would support the hypothesis that the brain region employs some of the features present in the artificial images.

More generally, we can go beyond using different exemplars from the same categories in the test set. Testing decoders for cross-generalization between different domains is useful to address a host of questions about the invariances of the encoding, for example its consistency between imagery and perception [45,46] and its stability across latencies after stimulus onset [47].

#### **Stimulus reconstruction provides a richer view of the information present, but complex priors complicate the interpretation**

A decoder may discriminate more than two categories. It might decode a continuous variable (e.g. the orientation of the stimulus; [17\*]), or it might output a multidimensional description of the stimulus. The most ambitious variant of decoding is stimulus reconstruction, where the output is a detailed estimated rendering of the stimulus. A decoder providing a rendering of the stimulus can potentially reveal more of the information present in the neuronal encoding [48–53].

Stimulus reconstruction is a particularly impressive feat of decoding for two related reasons. First, the space of

outputs is much richer, so the decoder reveals more of the encoded information (discriminating among more possible stimuli). Second, a more challenging level of generalization is typically required for the decoder to discriminate among this richer set of possible stimuli, because there are usually severe limitations on the number of stimuli that the decoder can be directly trained on. Successful generalization suggests that the structure of the model captures something about the code that enables generalization far beyond the training set.

To the extent that stimulus reconstruction, when applied to novel stimuli that the model has not been trained on, successfully specifies close matches within a rich space of reconstruction targets (e.g. pixel images), the analysis indicates that the code contains rich information about the stimulus. An extreme example of this would be a decoder that can precisely reconstruct arbitrary natural or unnatural images from their representation in V1. This would indicate that V1 encodes the image precisely and that the encoding is not restricted to natural images. In fact, this has never been shown for V1 and would be puzzling, because we expect visual representations to be specialized for natural stimuli.

In order to improve the appearance of reconstructions of natural stimuli, it is attractive to use a prior over the possible outputs of the decoder. For example, we might constrain the model to output an image whose low-level or high-level statistics match natural images. We may constrain the decoder more strongly, by limiting it to output only actual natural images. This constraint could be implemented using a restricted set of target images [51] or a generative model of natural images [54]. In either case, the quality of the reconstruction will reflect a combination of the information provided by the brain region and the information contributed by the prior. A complicated prior, therefore, also complicates the interpretation of the reconstruction results: good looking reconstructions no longer indicate that all the detail they provide is encoded in the brain region. The reconstruction has to be compared to the presented stimulus, and the complexity of the output space (which is reduced by the prior over the outputs) needs to be considered in the interpretation.

An important question is what we can learn from stimulus reconstructions. The goal to learn about the content and format of the code may not be ideally served by striving for the most natural looking reconstruction.

#### **Decoding models predicting behavioral responses from brain responses can be interpreted as brain-computational models**

Decoding is usually used as a tool of analysis that reveals aspects of the content and format of the information encoded in a brain region. The decoding model, thus,

is not interpreted as a model of brain computation. In the context of sensory systems, a decoder maps from brain responses to stimuli. Since stimulus processing by the brain operates in the opposite direction, it is difficult to interpret a decoder as a model of brain information processing. However, if a decoder is used to predict behavioral responses, for example, judgments of categorical or continuous stimulus variables (possibly including errors and reaction times on individual trials), then the decoder can be interpreted as a model (at a high level of abstraction) of the brain computations generating the behavioral responses from the encoding of the stimuli in the decoded brain region [55–57].

#### **Encoding models: testing comprehensive representational predictions**

Encoding models attempt to predict brain response patterns from descriptions of the experimental conditions (Figure 1a; [5°,20,31,58°,59°,60–62,63]). Encoding models, thus, operate in the opposite direction as decoding models.

If our goal is merely to demonstrate that a brain region contains information about the experimental conditions, then the direction the model should operate in is a technical issue: One direction may be more convenient for capturing the relevant statistical dependencies (e.g. noise correlations among responses), but a model operating in either direction could support the inference that particular information is present in the code. If our goal is to test computational theories, however, then the direction that the model operates in matters, because it determines whether the model can be interpreted as a brain-computational model.

#### **Encoding models predicting brain responses from sensory stimuli can serve as brain-computational models**

Whereas a decoding model typically serves to test for the presence of particular information in a brain region, an encoding model can provide a process model, at some level of abstraction, of the brain computations that produce the neuronal code. An encoding model makes comprehensive predictions about the representational space ([61,64°]; Box 2) and, ideally, should fully explain neuronal responses in the region in question (up to the maximum achievable given the noise in the data).

In sensory neuroscience, the experimental conditions correspond to sensory stimuli, and so an encoding model maps sensory stimuli to their encodings in sensory regions of the brain [5°,58°,59°,65]. A sensory encoding model, thus, operates in the direction of the information flow: from stimuli to brain responses. It should take raw sensory stimuli as input (e.g. images or sounds; [5°,31]; Khaligh-Razavi and Kriegeskorte, 2014; [60,66]; for recent reviews see Refs. [39°,67°]) and predict the patterns of brain

**Box 2 Encoding models: benefits and limitations**

An encoding model predicts the response of each measurement channel (e.g. a neuron or fMRI voxel) on the basis of properties of the experimental condition (e.g. sensory stimulus). Continuous brain maps can be obtained by fitting such a model to each response in turn. Responses are typically predicted as linear combinations of the features, rendering this approach closely related to classical univariate brain mapping. However, encoding models have several important benefits:

- More complex feature sets are used, often comprising thousands of descriptive features. Fitting the models, therefore, requires regularization. Typically, models are fitted using a penalty on the weights.
- Encoding models can have nonlinear components, such as the location and size of a visual receptive field.
- Models are tested for generalization to different experimental conditions (e.g. a different sample of visual stimuli).
- When sensory systems are studied, encoding models operate in the direction of information flow. They can then be used as a general method for testing brain-computational models, that is, models that process sensory stimuli. A vision model, for example, can take a novel image as input and predict responses to that image in cortex.

Analyses using encoding models to predict each response channel separately are limited in two ways:

- The separate modeling of each response channel is inconsistent with the multivariate nature of neuronal population codes and noise, and is statistically suboptimal. Separate modeling of each response entails low power for testing and comparing models, for two reasons: (1) Single responses (e.g. fMRI voxels) may be noisy and the evidence is not combined across locations. (2) The analyses treat the responses as independent, thus forgoing the benefit exploited by linear decoding approaches to model the noise multivariately. This is important in fMRI, where nearby voxels have highly correlated noise. As spatial resolution increases, we face the combined challenge of more and noisier voxels. This renders mapping with proper correction for multiple testing very difficult. In addition, relating results between individuals is not straightforward.
- When model parameters are color-coded and mapped across the cortex, the resulting detailed maps are not straightforward to interpret. (1) Weight maps of linear models are problematic to interpret in encoding models for the same reason they are problematic to interpret for decoding models: because a predictor's weight depends on the other predictors present in the model (unless each predictor is orthogonal to all others). The required regularization further complicates weight interpretation. (2) Models are fit at many locations and voxels highlighted on the basis of model fits. Post-selection inference on parameters is not usually performed. Because of these complications, results of fine-grained mapping across voxels are difficult to substantiate by formal inference.

Influential studies have met these challenges by focusing on single-subject analyses, acquiring a large amount of data in each subject, and limiting formal inference with correction for multiple testing to the overall explained variance of a model, with colors serving exploratory purposes.

activity the stimuli elicit. We can adjudicate among computational theories by implementing them in encoding models and testing their ability to predict brain-activity patterns [68].

**Brain-computational encoding models can be tested by predicting raw measurements, representational dissimilarities, or the activity-profiles distribution**

A brain-computational encoding model could consist in a set of hand-engineered features computed from the stimuli or in a neural network model trained on some task. How can we evaluate a high-dimensional representation in a brain-computational encoding model with brain-activity data, when the units of the model may not correspond one-to-one to the measured channels of brain activity?

One approach is to predict the raw measurements [5<sup>\*</sup>,31,58<sup>\*</sup>,59<sup>\*</sup>,62]. To achieve this, we can fit a linear model that maps from the outputs of the nonlinear encoder thought to capture the brain computations (e.g. a layer of a neural network model processing the stimuli) to each measured response channel, for example, each neuron or fMRI voxel in the region we would like to understand. The linear model will capture which units of the nonlinear encoding model best predict each response channel. It can be interpreted as capturing the sampling and averaging that occurs in the measurement of brain activity. A neuronal recording array, for example, will capture a sample of neurons, and fMRI will give us

measurements akin to local spatiotemporal averages at regular grid locations. For each response channel, the linear model will have a parameter for each nonlinear feature (e.g. each unit of the neural network layer). Fitting these parameters requires substantial training data and normally also a prior on the parameters (e.g. a 0-mean Gaussian prior as is implicit to fitting with an L2 penalty).

Instead of predicting the raw measurements, we can use a brain-computational encoding model to predict to what extent different stimuli are dissimilar in the representation, an approach called representational similarity analysis (RSA; [69<sup>\*</sup>,70]). The pairwise dissimilarities of the multivariate response patterns representing the stimuli can be summarized in a representational dissimilarity matrix (RDM). For example, for each pair of stimuli, the dissimilarity between the associated response patterns could be measured using the Euclidean distance. The RDM characterizes the geometry of the set of points in the multivariate response space that correspond to the stimuli. Noise correlations can be captured by estimating the covariance matrix of the residuals of the response-estimation model, and replacing the Euclidean distance by the Mahalanobis distance. To remove the positive bias associated with measuring distances between noisy data points, we can use the crossnobis (crossvalidated Mahalanobis) estimator [70,71]. The resulting crossnobis RDM provides a full characterization of the linearly decodable information in the representational space [64<sup>\*</sup>]. Comparing representations in models and brains at the level of

RDMs obviates the need for fitting a linear model to predict each measured response (thus reducing the need for training data) and enables the analysis to naturally handle noise correlations between responses (which are typically ignored when encoding models separately predict each of the measured response channels).

A third approach to the evaluation of encoding models is to predict the distribution of activity profiles. In pattern component modelling (PCM, [72<sup>\*</sup>]), this distribution is characterized by the second moment of the activity profiles. Like the RDM, this is a stimulus-by-stimulus summary statistic of the stimulus-response matrix. Each entry of the second-moment matrix corresponds to the inner product between two response patterns.

All three approaches can be construed as testing hypotheses about the representational space induced by the activity profiles [64<sup>\*</sup>]. Consider a linear encoding model using a Gaussian prior on the weights. Such a model predicts a Gaussian distribution of activity profiles. The predicted distribution of activity profiles is captured by its second moment. For representational similarity analysis, similarly, the RDM is a function of the second moment of the activity profiles. This core mathematical commonality between the methods notwithstanding, each is best suited for a particular set of questions.

Linear models predicting raw measurements lend themselves to univariate brain mapping, revealing which voxels or neurons are accounted for by a particular nonlinear encoding model. RSA lends itself to characterizing the geometry of the representational space, naturally handles noise correlations among responses, and reduces the need for training data. PCM can have greater sensitivity for adjudicating among models than the other two methods, at the expense of relying on stronger assumptions. The three methods are best viewed as part of a single toolbox of representational model analyses, whose elements can be combined as needed to address particular questions.

#### **The level of generalization beyond the training set must be considered when interpreting an encoding result**

Encoders, like decoders, are tested by evaluating how well they predict independent data, whether the predicted quantities are the raw brain-activity measurements, the representational dissimilarity matrix, or the second-moment matrix of the activity profiles. For encoders, as for decoders, the interpretation depends on both the prediction accuracy and the level of generalization beyond the training set that the model achieves.

Encoding models typically require the fitting of parameters, so overfitting needs to be accounted for in any inferential procedure. In the simplest type of a univariate linear encoding model, we can rely on Gaussian assumptions and perform inference without a separate

test set (e.g. [73]). However, more interesting models require independent test sets, for example when parameters are fitted using priors over the weights and when the model is a brain-computational model to be tested for generalization to new conditions.

A key consideration is how much flexibility to allow in fitting each model representation to a brain representation. One extreme is to allow no flexibility and assume that the model representation precisely predicts the geometry of the representational space [69<sup>\*</sup>]. This case is most naturally handled by RSA and PCM, but could also be implemented with linear encoders by using a prior that prevents any distortion of the representational geometry. The other extreme is to allow arbitrary linear remixing of the units of the nonlinear encoding model. This case is most naturally handled with linear encoding models, but can also be implemented with PCM and RSA ([72<sup>\*</sup>]; Khaligh-Razavi and Kriegeskorte, 2014; [74]). In practice, some compromise is desirable, which we can think of as a prior on the mapping from the brain-computational model to the measured brain responses. We might use a 0-mean Gaussian prior on the weights (e.g. [5<sup>\*</sup>]). Alternatively, we can limit flexibility more aggressively, by allowing each unit (or each feature map or layer) a single weight (not a separate weight for each response). Such weighted representational models (e.g. [63,74]) are naturally implemented with RSA and PCM. Each brain-computational model in this case predicts a superset of the features spanning the brain representational space (disallowing linear mixtures), but does not predict the prevalence of each of the features in the neuronal population.

The lowest level of generalization beyond the training set is generalization to new measurements for the same experimental conditions. This is sufficient, if the experimental conditions exhaustively cover the domain we would like to draw inferences about (consider the case of the representation of the five fingers in motor cortex: [75]). However, in a domain such a sensory systems, the goal is typically to evaluate to what extent a brain-computational model can predict brain representations of arbitrary stimuli. This requires a higher level of generalization beyond the training set. A vision model, for example, might be trained with responses to one sample of natural images and tested for generalization to responses to an independent (and nonoverlapping) sample from the same distribution of natural images. Because the set of all natural images is so rich, this is a challenging generalization task (as illustrated by the difficulty of computer vision). An even more stringent test of the assumptions implicit to a model is to train the model on a sample from one population of images and test it on a sample from a different population of images (e.g. [76]).

The prediction accuracy can be assessed in terms of whether it is significantly above chance level, whether it significantly differs from that of competing encoding models, and how

close it comes to the noise ceiling (the highest achievable accuracy given the noise in the data, [70]).

We can generalize claims about an encoding model to the extent that its predictions generalize. If we want to conclude that the model can predict responses for the stimuli presented, we need not test the model with different stimuli (only with different response measurements elicited by the same stimuli). If we want to conclude that the model can predict responses to arbitrary natural stimuli, we need to test it with new arbitrary natural stimuli. The population of conditions that the test set is a sample of defines the scope of the claims we can make [37,38,39\*].

We focus here on encoding and decoding models that are fitted to individual subjects' brains, so as to be able to exploit fine-grained idiosyncratic patterns of activity that are unique to each subject. Within-subject prediction accuracy may support generalization to a population of stimuli, but it doesn't support generalization to the population of subjects. In some fields, such as low-level vision, researchers draw on prior knowledge and *assume* that results that hold for a few subjects will hold for the population. If we were instead to use our data to generalize our inferences to the population of subjects, we would either need to predict results for held-out subjects (which would make it impossible to exploit individually unique fine-grained activity patterns) or perform inference on the within-subject prediction accuracies with subject as a random effect.

#### **The feature fallacy: interpreting the success of a model as confirmation of its basis set of features**

Linear encoding models predict each measured activity profile as a linear combination of a set of model features. When a model can explain the measured activity profiles, we might be tempted to conclude that the model features are encoded in the measured responses. This interpretation is problematic because the same linear space can be spanned by many alternative basis sets of model features.

The fact that multiple sets of basis vectors can span the same space is widely appreciated. However, it is not obvious whether the ambiguity is removed when a prior over the encoding weights is used. An encoding model with a 0-mean isotropic Gaussian prior on the weights (equivalent to ridge regression) predicts a Gaussian distribution of activity profiles (captured by the second moment of the activity profiles as the sufficient statistic) and a particular representational geometry (captured by the RDM). The use of a weight prior does reduce the ambiguity. In the absence of a weight prior, all models spanning the same space of activity profiles make equivalent predictions. With a weight prior, different models spanning the same space can predict distinct distributions of activity profiles. However, substantial ambiguity

remains (Figure 2) because there are infinitely many alternative ways to express the same distribution of activity profiles by different feature sets (each assuming a 0-mean isotropic Gaussian prior over the weights).

The freedom to choose among feature sets generating the same distribution of activity profiles is useful in that it enables us to implement a nonisotropic Gaussian prior on the weights of a given model (Tikhonov regularization) by re-expressing the model such that the features induce the same distribution of activity profiles in combination with a 0-mean isotropic Gaussian prior [77]. However, this freedom to express the same model by different feature sets needs to be kept in mind when interpreting results: Many alternative feature sets would have given the same results.

In sum, the same distribution of activity profiles can be expressed in many ways (by different feature sets), so the fact that a model explains brain responses does not provide evidence in favor of the particular feature set chosen. Diedrichsen [78] termed the interpretation in terms of the particular features used the 'feature fallacy'. This fallacy is arguably somewhat persistent in the neuroscience literature [79].

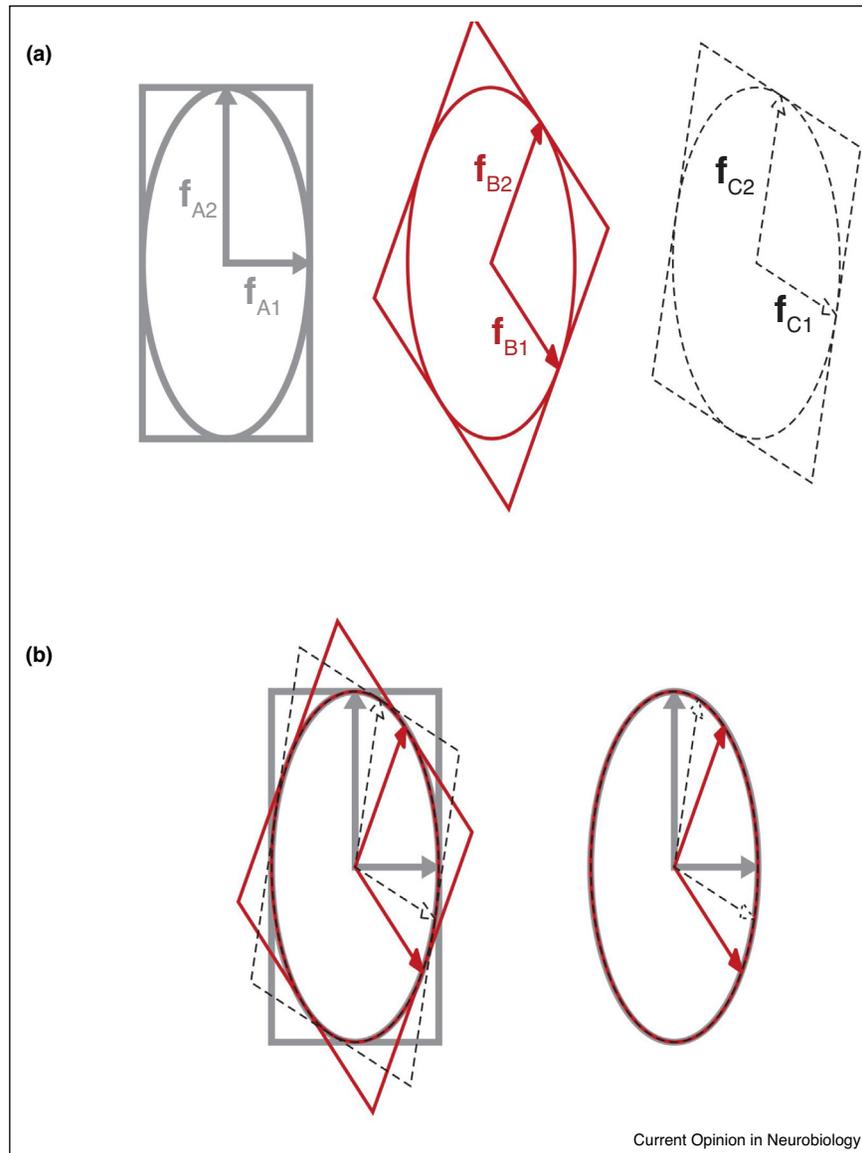
#### **Weights of linear models are not straightforward to interpret in either encoding or decoding models**

Beyond interpretation of the overall success of an encoding or decoding model, researchers often want to dig deeper and interpret the fitted parameters of their models. In the context of a decoding model, the weights assigned to the voxels would seem to tell us where the information the decoder uses is located in the brain. Similarly, in the context of an encoding model, the weights of different model features promise to reveal to what extent different model features are encoded in a brain region.

Unfortunately, the interpretation of the weights of linear models is not as straightforward as the simplicity of a linear relationship might suggest. A weight does not reflect the predictive power of an individual predictor (a measured brain response in decoding, a model features in encoding). Rather a weight reflects a predictor's contribution in the context of the rest of the model.

Uninformative predictors can receive large positive or negative weights. For example, an fMRI voxel that does not by itself contain any information can have a large absolute weight in a linear decoder if it improves decoding accuracy by cancelling noise that the voxel shares with other voxels that do contain stimulus information (Figure 1b; Figure 3; [80\*]). In an encoding model, similarly, a model feature might contain no information about the modeled response and still receive a large absolute weight. For example, an fMRI voxel that only responds to faces might be explained by a set of semantic

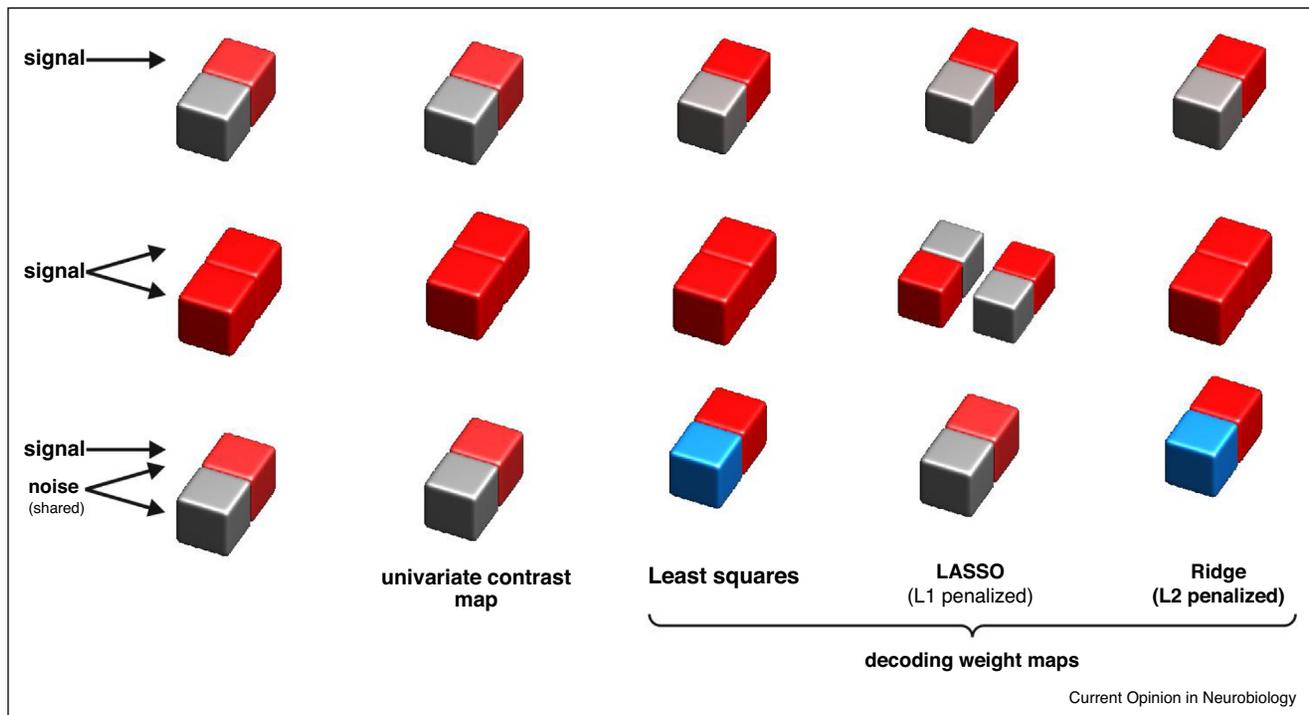
Figure 2



### The feature fallacy.

Different linear encoding models spanning the same space of activity profiles may not be distinguishable. There are many alternative sets of feature vectors  $\{f_1, f_2, \dots\}$  that span the same space of activity profiles. In the absence of a prior on the weights of the linear model, all these sets can equally explain a given set of brain responses. The ambiguity is reduced, but not resolved when a prior on the weights is assumed [78]. If we define a prior on the weights, then each model predicts a probability density over the space of activity profiles. This probabilistic prediction may be distinct for two sets of basis features, even if they span the same space. For example, if the weight prior is a 0-mean isotropic Gaussian, then each model assigns probabilities to different activity profiles according to a Gaussian distribution over the space of profiles. Two linear models may span the same space, but predict distinct distributions of activity profiles. However, even with a Gaussian weight prior, there are still (infinitely) many equivalent models that make identical probabilistic predictions. We illustrate this by example. **(a)** Three models (A, B, C) each contain two feature vectors as predictors (A:  $\{f_{A1}, f_{A2}\}$ , B:  $\{f_{B1}, f_{B2}\}$ , C:  $\{f_{C1}, f_{C2}\}$ ). The three models all span the same 2-dimensional space of activity profiles. For each model, we assume a 0-mean isotropic Gaussian weight prior. **(b)** All three models predict the same nonisotropic Gaussian probability density over the space of activity profiles (indicated by a single iso-probability-density contour: the ellipse). Model A (gray) predicts the density by modeling it with two orthogonal features that capture the principal-component axes, with features having different norms to capture the anisotropy. Model B predicts the same density by modeling it with two correlated features of similar norm. Model C falls somewhere in between, combining feature correlation and different feature norms to capture the same Gaussian density over the activity profiles. Note that there are many other models that span the same space, but will not induce the same probability density over activity profiles when complemented by a 0-mean isotropic Gaussian weight prior. A given linear encoding model's success at predicting brain responses provides evidence for the induced distribution of activity profiles, but not for the particular features chosen to express that distribution.

Figure 3



Weights of decoding and encoding models are difficult to interpret.

Three examples (rows) illustrate the difficulty of interpreting decoder weights for a pair of voxels. In the first example (top row), only the top voxel contains signal (stimulus information, red) and the two voxels have independent noise. This scenario is unproblematic: both univariate mapping (second column from the left) and decoder weight maps detect the informative voxel (red). In the second example (second row), both voxels contain the same signal. Here, univariate mapping and weight maps often work. However, the LASSO decoder, because of its preference for a sparse solution, may choose one of the voxels arbitrarily. In the third example (third row), only the top voxel contains signal and both voxels contain correlated noise. Univariate mapping correctly identifies the informative voxel. Linear decoders will give negative weight (blue) to the uninformative voxel, so as to cancel the noise.

descriptors unrelated to faces if the nonface semantic descriptors are capable, in combination, of capturing contextual variation that is correlated with the presence of faces.

Conversely, informative predictors may receive weights that are small in absolute value or zero. For example, in fMRI, a voxel might receive zero weight when other voxels suffice for decoding when a weight penalty (especially a sparsity encouraging weight penalty, for example defined by the L1 norm of the weights) leads the fitting procedure to select among equally informative voxels (Figure 3). Weight penalties can similarly suppress model feature weights in encoding models.

A related problem is performing statistical inference to test hypotheses about the weights [81]. Systems and cognitive neuroscience has yet to develop a proper set of methods for hypothesis testing in this context.

A simple remedy to the complications associated with the interpretation of fitted weights is to interpret only the

*accuracy* of decoding and encoding models (and its significance level) and not the parameters of the models. In the context of decoding, this makes sense for local regions of interest corresponding, for example, to cortical areas, which we can test one by one for the presence of particular information. It can be generalized to brain mapping by applying the decoder (or more generally any multivariate analysis) within a searchlight that scans the brain for the effect of interest [18]. Like classical univariate brain mapping, this approach derives its interpretability from the fact that each location is independently subjected to the same analysis. However, instead of averaging local responses, the evidence is summarized using local multivariate statistics. In the context of encoding models, similarly, we can focus on each model's overall performance and on inferential comparisons among multiple models.

### The single-model-significance fallacy

When our goal is merely to detect information in a brain region, we don't interpret the model as a model of brain computation. This lowers the requirements for the model: It need not operate in the direction of information

flow and it need not be neurobiologically plausible. The model is merely a statistical tool to sense a dependency. The choice of model in this scenario, will affect our sensitivity, and its structure is not entirely irrelevant to the interpretation. For example, decodability by a linear model tells us something about the format of the encoding. However, a single model will suffice to demonstrate the presence of information in a brain region.

When our goal is to gain insight into the computations the brain performs, a model has a more prominent role: it is meant to capture, at some level of abstraction, the computations occurring in the brain. We then require the model to operate in the causal direction and to be neurobiologically plausible (albeit abstracted). Examples of such models include encoding models of sensory responses and decoding models that predict behavioral responses from brain activity. Psychophysical models, which skip the brain entirely and predict behavioral responses directly from stimuli, also fall into this class. In all these cases, finding that a model explains significant variance is a very low bar and tells us little as to whether the model captures the computational process.

The single-model-significance fallacy is to interpret the fact that a single model explains significant variance as evidence in favor of the model. A simple example that is widely understood is linear correlation. A significant linear correlation does demonstrate a dependency between two variables, but it does not demonstrate that the dependency is linear. Similarly, the fact that a complex encoding model explains significant variance in the responses of a brain region to a test set of novel stimuli does demonstrate that the brain region contains information about the stimuli, but it does not demonstrate that the encoding model captures the process that computes the encoding.

Even a bad model can explain significant variance, especially if it has a large number of parameters fitted to the data. In order to learn about the underlying brain computations, we need to (a) consider multiple models, (b) assess what proportion of the explainable (i.e. non-noise) variance each explains at a given level of generalization, and (c) compare the models inferentially.

### Representational interpretations require additional assumptions

Decoding and encoding models are often motivated by the goal to understand how the brain represents the world, as well as the animal's decisions, goals, plans, actions, and motor dynamics. Significant variance explained by encoding and decoding models demonstrates the presence of information. Interpreting this information as a representation [82] implies the additional claim that the brain activity *serves the purpose* to convey the information to other parts of the brain [24,64\*,83]. This

functional interpretation is so compelling in the context of sensory systems that we sometimes jump too easily from findings of information to representational interpretations [84,85]. In addition to the presence of the information, its functional role as a representation implies that the information is read out by other regions, affecting downstream processing and ultimately behavior. Combining encoding and decoding models with stimulus-based and response-based experimentation can help disambiguate the causal implications [11\*]. Ideally, experimental control of neural activity should also be used to test whether activity has particular downstream or behavioral consequences [86,87]. To the extent that we rely on prior assumptions to justify a representational interpretation, it is important to reflect on these and consider if there is evidence from previous studies to support them.

### Conflict of interest statement

Nothing declared.

### References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest
- 1. Dechamps RC, Zador A: **Neural representation and the cortical code.** *Ann Rev Neurosci* 2000, **23**:613-647.
- 2. Felleman DJ, Van Essen D: **Distributed hierarchical processing in the primate cerebral cortex.** *Cerebral Cortex (New York, NY: 1991)* 1991, **1**:1-47.
- 3. Young MP: **Objective analysis of the topological organization of the primate cortical visual system.** *Nature* 1992, **358**:152.
- 4. Wu MCK, David SV, Gallant JL: **Complete functional characterization of sensory neurons by system identification.** *Annu Rev Neurosci* 2006, **29**:477-505.
- 5. Kay KN, Naselaris T, Prenger RJ, Gallant JL: **Identifying natural images from human brain activity.** *Nature* 2008, **452**:352-355 <http://dx.doi.org/10.1038/nature06713>.
- A seminal encoding and decoding study that helped the field envision how neuroimaging data can be used to test image-computable vision models.
- 6. Wen H, Shi J, Chen W, Liu Z: **Deep residual network predicts cortical representation and organization of visual features for rapid categorization.** *Sci Rep* 2018, **8**:3752.
- 7. Majaj NJ, Hong H, Solomon EA, DiCarlo JJ: **Simple learned weighted sums of inferior temporal neuronal firing rates accurately predict human core object recognition performance.** *J Neurosci* 2015, **35**:13402-13418.
- 8. Simoncelli EP, Olshausen BA: **Natural image statistics and neural representation.** *Ann Rev Neurosci* 2001, **24**:1193-1216.
- 9. Tishby N, Pereira FC, Bialek W: *The Information Bottleneck Method.* . arXiv preprint physics/0004057 2000.
- 10. Achille A, Soatto S: **Information dropout: learning optimal representations through noisy computation.** *IEEE Trans Pattern Anal Mach Intell* 2018, **40**:2897-2905.
- 11. Weichwald S, Meyer T, Özdenizci O, Schölkopf B, Ball T, Grosse-Wentrup M: **Causal interpretation rules for encoding and decoding models in neuroimaging.** *Neuroimage* 2015, **110**:48-59.

This paper introduces a set of causal interpretation rules for encoding and decoding models on the basis of the technical literature on causal inference.

12. Cox DD, Savoy RL: **Functional magnetic resonance imaging (fMRI) “brain reading”**: detecting and classifying distributed patterns of fMRI activity in human visual cortex. *NeuroImage* 2003, **19**:261-270.
13. Haynes J-D, Rees G: **Decoding mental states from brain activity in humans**. *Nat Rev Neurosci* 2006, **7**:523-534 <http://dx.doi.org/10.1038/nrn1931>.
14. Tong F, Pratte MS: **Decoding patterns of human brain activity**. *Ann Rev Psychol* 2012, **63**:483-509.
15. Haxby JV *et al.*: **Distributed and overlapping representations of faces and objects in ventral temporal cortex**. *Science* 2001, **293**:2425-2430 <http://dx.doi.org/10.1126/science.1063736>.  
A seminal decoding study of ventral-stream category representations that helped the field conceptualize neuroimaging data as reflections of distributed neuronal population codes.
16. Carlson TA *et al.*: **Patterns of activity in the categorical representation of objects**. *J Cogn Neurosci* 2003, **15**:704-717.
17. Kamitani Y, Tong F: **Decoding the visual and subjective contents of the human brain**. *Nat Neurosci* 2005, **8**:679-685 <http://dx.doi.org/10.1038/nrn1444>.  
A seminal decoding study of primary visual representations of oriented gratings, which established that functional magnetic resonance imaging enables orientation decoding.
18. Kriegeskorte N, Goebel R, Bandettini P: **Information-based functional brain mapping**. *Proc Natl Acad Sci U S A* 2006, **103**:3863-3868 <http://dx.doi.org/10.1073/pnas.0600244103>.
19. Norman KA *et al.*: **Beyond mindreading: multi-voxel pattern analysis of fMRI data**. *Trends Cogn Sci* 2006, **10**:424-430.
20. Paninski L, Pillow J, Lewi J: **Statistical models for neural encoding, decoding, and optimal stimulus design**. *Progr Brain Res* 2007, **165**:493-507.
21. Mur M, Bandettini PA, Kriegeskorte N: **Revealing representational content with pattern-information fMRI—an introductory guide**. *Soc Cogn Affect Neurosci* 2009, **4**:101-109 <http://dx.doi.org/10.1093/scan/nsn044>.
22. Pereira F, Mitchell T, Botvinick M: **Machine learning classifiers and fMRI: a tutorial overview**. *NeuroImage* 2009, **45** (Suppl. 1):S199-209 <http://dx.doi.org/10.1016/j.neuroimage.2008.11.007>.
23. Pillow JW, Ahmadian Y, Paninski L: **Model-based decoding, information estimation, and change-point detection techniques for multineuron spike trains**. *Neural Comput* 2011, **23**:1-45.
24. Kriegeskorte N: **Pattern-information analysis: from stimulus decoding to computational-model testing**. *NeuroImage* 2011, **56**:411-421 <http://dx.doi.org/10.1016/j.neuroimage.2011.01.061>.
25. Hebart MN, Görgen K, Haynes JD: **The decoding toolbox (TDT): a versatile software package for multivariate analyses of functional imaging data**. *Front Neuroinform* 2015, **8**:88.
26. Haynes J-D: **A primer on pattern-based approaches to fMRI: principles, pitfalls, and perspectives**. *Neuron* 2015, **87**:257-270 <http://dx.doi.org/10.1016/j.neuron.2015.05.025>.
27. Hebart MN, Baker CI: **Deconstructing multivariate decoding for the study of brain function**. *Neuroimage* 2018, **180**:4-18.
28. Varoquaux G, Raamana PR, Engemann DA, Hoyos-Idrobo A, Schwartz Y, Thirion B: **Assessing and tuning brain decoders: cross-validation, caveats, and guidelines**. *NeuroImage* 2017, **145**:166-179.
29. Paninski L, Cunningham J: **Neural data science: accelerating the experiment-analysis-theory cycle in large-scale neuroscience**. *bioRxiv* 2017:196949.
30. Duda RO, Hart PE, Stork DG: *Pattern Classification*. John Wiley & Sons; 2012.
31. Naselaris T, Kay KN, Nishimoto S, Gallant JL: **Encoding and decoding in fMRI**. *NeuroImage* 2011, **56**:400-410 <http://dx.doi.org/10.1016/j.neuroimage.2010.07.073>.
32. Averbeck BB, Latham PE, Pouget A: **Neural correlations, population coding and computation**. *Nat Rev Neurosci* 2006, **7**:358.
33. Poggio T, Girosi F: **Networks for approximation and learning**. *Proc IEEE* 1990, **78**:1481-1497.
34. DiCarlo JJ, Cox D: **Untangling invariant object recognition**. *Trends Cogn Sci* 2007, **11**:334-341.
35. DiCarlo JJ, Zoccolan D, Rust NC: **How does the brain solve visual object recognition?** *Neuron* 2012, **73**:415-434.
36. Hong H, Yamins DLK, Majaj NJ, DiCarlo JJ: **Explicit information for category-orthogonal object properties increases along the ventral stream**. *Nat Neurosci* 2016, **19**:613-622 <http://dx.doi.org/10.1038/nn.4247>.
37. Hastie T, Tibshirani R, Friedman J: *The Elements of Statistical Learning*. New York, NY: Springer New York; 2009 <http://dx.doi.org/10.1007/978-0-387-84858-7>.
38. Kriegeskorte N: **Crossvalidation in brain imaging analysis**. *bioRxiv* 2015 . (No. biorxiv:017418v1). Retrieved from <http://bioRxiv.org/lookup/doi/10.1101/017418>.
39. Kriegeskorte N: **Deep neural networks: a new framework for modeling biological vision and brain information processing**. *Ann Rev Vision Sci* 2015, **1**:417-446.  
A review and perspective paper developing a framework for how to build neural network models of brain information processing with constraints from brain-activity data.
40. Chung S, Lee DD, Sompolinsky H: **Linear readout of object manifolds**. *Phys Rev E* 2016, **93**:060301.
41. Chung S, Lee DD, Sompolinsky H: **Classification and geometry of general perceptual manifolds**. *Phys Rev X* 2018, **8**:031003.
42. Kriegeskorte N *et al.*: **Individual faces elicit distinct response patterns in human anterior temporal cortex**. *Proc Natl Acad Sci U S A* 2007, **104**:20600-20605.
43. Anzellotti S, Fairhall SL, Caramazza A: **Decoding representations of face identity that are tolerant to rotation**. *Cereb Cortex* 2014, **24**:1988-1995 <http://dx.doi.org/10.1093/cercor/bht046>.
44. Freedman DJ, Riesenhuber M, Poggio T, Miller EK: **Visual categorization and the primate prefrontal cortex: neurophysiology and behavior**. *J Neurophysiol* 2002, **88**:929-941 <http://dx.doi.org/10.1152/jn.2002.88.2.929>.
45. Stokes M *et al.*: **Top-down activation of shape-specific population codes in visual cortex during mental imagery**. *J Neurosci* 2009, **29**:1565-1572.
46. Cichy RM *et al.*: **Imagery and perception share cortical representations of content and location**. *Cereb Cortex* 2012, **22**:372-380.
47. King JR, Dehaene S: **Characterizing the dynamics of mental representations: the temporal generalization method**. *Trends Cogn Sci* 2014, **18**:203-210.
48. Thirion B, Duchesnay E, Hubbard E, Dubois J, Poline JB, LeBihan D, Dehaene S: **Inverse retinotopy: inferring the visual content of images from brain activation patterns**. *NeuroImage* 2006, **33**:1104-1116.
49. Miyawaki Y, Uchida H, Yamashita O, Sato M, Morito Y, Tanabe HC, Sadato N, Kamitani Y: **Visual image reconstruction from human brain activity using a combination of multiscale local image decoders**. *Neuron* 2008, **60**:915-929 <http://dx.doi.org/10.1016/j.neuron.2008.11.004>.
50. Mesgarani N, David SV, Fritz JB, Shamma SA: **Influence of context and behavior on stimulus reconstruction from neural activity in primary auditory cortex**. *J Neurophysiol* 2009, **102**:3329-3339.
51. Naselaris T, Prenger RJ, Kay KN, Oliver M, Gallant JL: **Bayesian reconstruction of natural images from human brain activity**. *Neuron* 2009, **63**:902-915 <http://dx.doi.org/10.1016/j.neuron.2009.09.006>.
52. Santoro R, Moerel M, De Martino F, Valente G, Ugurbil K, Yacoub E, Formisano E: **Reconstructing the spectrotemporal**

- modulations of real-life sounds from fMRI response patterns.** *Proc Natl Acad Sci U S A* 2017, **114**:4799-4804.
53. Parthasarathy N, Batty E, Falcon W, Rutten T, Rajpal M, Chichilnisky EJ, Paninski L: **Neural networks for efficient bayesian decoding of natural images from retinal neurons.** *Advances in Neural Information Processing Systems*. 2017:6434-6445.
  54. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y: **Generative adversarial nets.** *Advances in Neural Information Processing Systems*. 2014:2672-2680.
  55. Shadlen MN, Britten KH, Newsome WT, Movshon JA: **A computational analysis of the relationship between neuronal and behavioral responses to visual motion.** *J Neurosci* 1996, **16**:1486-1510.
  56. Williams MA, Dang S, Kanwisher NG: **Only some spatial patterns of fMRI response are read out in task performance.** *Nat Neurosci* 2007, **10**:685.
  57. Walther DB *et al.*: **To err is human: correlating fMRI decoding and behavioral errors to probe the neural representation of natural scene categories.** In *Visual Population Codes - Toward a Common Multivariate Framework for Cell Recording and Functional Imaging*. Edited by Kriegeskorte N, Kreiman G. MIT Press; 2012:391-416.
  58. Dumoulin SO, Wandell BA: **Population receptive field estimates in human visual cortex.** *NeuroImage* 2008, **39**:647-660 <http://dx.doi.org/10.1016/j.neuroimage.2007.09.034>.  
This paper introduces nonlinear encoding models for early visual representations, which enabled more rigorous investigations of early vision with neuroimaging data.
  59. Mitchell TM *et al.*: **Predicting human brain activity associated with the meanings of nouns.** *Science* 2008, **320**:1191-1195.  
A seminal encoding study of semantic brain representations using semantic property descriptions to predict brain representations of novel stimuli.
  60. Yamins DLK *et al.*: **Performance-optimized hierarchical models predict neural responses in higher visual cortex.** *Proc Natl Acad Sci U S A* 2014, **111**:8619-8624.
  61. Naselaris T, Kay KN: **Resolving ambiguities of MVPA using explicit models of representation.** *Trends Cogn Sci* 2015, **19**:551-554 <http://dx.doi.org/10.1016/j.tics.2015.07.005>.
  62. van Gerven MA: **A primer on encoding models in sensory neuroscience.** *J Math Psychol* 2017, **76**:172-183.
  63. Khaligh-Razavi SM, Kriegeskorte N: **Deep supervised, but not unsupervised, models may explain IT cortical representation.** *PLoS Comput Biol* 2014, **10**:e1003915.
  64. Diedrichsen J, Kriegeskorte N: **Representational models: a common framework for understanding encoding, pattern-component, and representational-similarity analysis.** *PLoS Comput Biol* 2017, **13**:e1005508 <http://dx.doi.org/10.1371/journal.pcbi.1005508>.  
This paper explains the conceptual and mathematical relationships between linear encoding models, pattern-component modeling, and representational similarity analysis and how these methods provide complementary tools in a single toolbox for representational analyses.
  65. Huth AG *et al.*: **A continuous semantic space describes the representation of thousands of object and action categories across the human brain.** *Neuron* 2012, **76**:1210-1224.
  66. Kell AJE, Yamins DLK, Shook EN, Norman-Haignere SV, McDermott JH: **A task-optimized neural network replicates human auditory behavior, predicts brain responses, and reveals a cortical processing hierarchy.** *Neuron* 2018, **98**:630-644.e16 <http://dx.doi.org/10.1016/j.neuron.2018.03.044>.
  67. Yamins DLK, DiCarlo JJ: **Using goal-driven deep learning models to understand sensory cortex.** *Nat Neurosci* 2016, **19**:356-365.  
A review and perspective paper describing a framework for modeling sensory processing by training deep neural networks on tasks and testing them with brain-activity data.
  68. Kriegeskorte N, Diedrichsen J: **Inferring brain-computational mechanisms with models of activity measurements.** *Philos Trans R Soc B: Biol Sci* 2016, **371**:20160278 <http://dx.doi.org/10.1098/rstb.2016.0278>.
  69. Kriegeskorte N *et al.*: **Representational similarity analysis – connecting the branches of systems neuroscience.** *Front Syst Neurosci* 2008, **2**:4.  
The paper introducing representational similarity analysis.
  70. Nili H, Wingfield C, Walther A, Su L, Marslen-Wilson W, Kriegeskorte N: **A toolbox for representational similarity analysis.** *PLoS Comput Biol* 2014, **10**:e1003553.
  71. Walther A, Nili H, Ejaz N, Alink A, Kriegeskorte N, Diedrichsen J: **Reliability of dissimilarity measures for multi-voxel pattern analysis.** *Neuroimage* 2016, **137**:188-200.
  72. Diedrichsen J, Ridgway GR, Friston KJ, Wiestler T: **Comparing the similarity and spatial structure of neural representations: a pattern-component model.** *NeuroImage* 2011, **55**:1665-1678 <http://dx.doi.org/10.1016/j.neuroimage.2011.01.044>.  
This paper introduces pattern component modeling.
  73. Friston KJ, Holmes AP, Worsley KJ, Poline JP, Frith CD, Frackowiak RS: **Statistical parametric maps in functional imaging: a general linear approach.** *Hum Brain Mapp* 1994, **2**:189-210.
  74. Khaligh-Razavi S-M, Henriksson L, Kay K, Kriegeskorte N: **Fixed versus mixed RSA: explaining visual representations by fixed and mixed feature sets from shallow and deep computational models.** *J Math Psychol* 2016 <http://dx.doi.org/10.1016/j.jmp.2016.10.007>.
  75. Diedrichsen J, Wiestler T, Krakauer JW: **Two distinct ipsilateral cortical representations for individuated finger movements.** *Cereb Cortex* 2012, **23**:1362-1377.
  76. Eickenberg M, Gramfort A, Varoquaux G, Thirion B: **Seeing it all: convolutional network layers map the function of the human visual system.** *NeuroImage* 2017, **152**:184-194.
  77. Nunez-Elizalde AO, Huth AG, Gallant JL: **Voxelwise encoding models with non-spherical multivariate normal priors.** *bioRxiv* 2018:386318.
  78. Diedrichsen J: **Representational models.** In *The Cognitive Neurosciences*. Edited by Gazzaniga MS. 2018.
  79. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV: **Neural population dynamics during reaching.** *Nature* 2012 <http://dx.doi.org/10.1038/nature11129>.
  80. 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* 2014, **87**:96-110 <http://dx.doi.org/10.1016/j.neuroimage.2013.10.067>.  
This paper provides a careful mathematical discussion of the difficulty of interpreting the weights of linear models.
  81. Taylor J, Tibshirani RJ: **Statistical learning and selective inference.** *Proc Natl Acad Sci U S A* 2015, **112**:7629-7634.
  82. Dennett D: *The Intentional Stance*. MIT Press; 1987.
  83. Kriegeskorte N, Bandettini P: **Analyzing for information, not activation, to exploit high-resolution fMRI.** *Neuroimage* 2007, **38**:649-662.
  84. de-Wit L, Alexander D, Ekroll V, Wagemans J: **Is neuroimaging measuring information in the brain?** *Psychon Bull Rev* 2016, **23**:1415-1428.
  85. Ritchie JB, Kaplan DM, Klein C: **Decoding the brain: neural representation and the limits of multivariate pattern analysis in cognitive neuroscience.** *Br J Philos Sci* 2017.
  86. Afraz SR, Kiani R, Esteky H: **Microstimulation of inferotemporal cortex influences face categorization.** *Nature* 2006, **442**:692.
  87. Raizada RD, Kriegeskorte N: **Pattern-information fMRI: new questions which it opens up and challenges which face it.** *Int J Imaging Syst Technol* 2010, **20**:31-41.