

Combinatorial Neural Codes from a Mathematical Coding Theory Perspective

Carina Curto

ccurto2@math.unl.edu

Vladimir Itskov

vitskov2@math.unl.edu

Katherine Morrison

kmorris2@gmail.com

Zachary Roth

s-zroth1@math.unl.edu

Judy L. Walker

judy.walker@unl.edu

Department of Mathematics, University of Nebraska-Lincoln, Lincoln, NE 68588, U.S.A.

Shannon's seminal 1948 work gave rise to two distinct areas of research: information theory and mathematical coding theory. While information theory has had a strong influence on theoretical neuroscience, ideas from mathematical coding theory have received considerably less attention. Here we take a new look at combinatorial neural codes from a mathematical coding theory perspective, examining the error correction capabilities of familiar receptive field codes (RF codes). We find, perhaps surprisingly, that the high levels of redundancy present in these codes do not support accurate error correction, although the error-correcting performance of receptive field codes catches up to that of random comparison codes when a small tolerance to error is introduced. However, receptive field codes are good at reflecting distances between represented stimuli, while the random comparison codes are not. We suggest that a compromise in error-correcting capability may be a necessary price to pay for a neural code whose structure serves not only error correction, but must also reflect relationships between stimuli.

1 Introduction ---

Shannon's seminal work (Shannon, 1948) gave rise to two distinct though related areas of research: information theory (Cover & Thomas, 2006) and mathematical coding theory (MacWilliams & Sloane, 1983; Huffman & Pless, 2003). While information theory has had a strong influence on theoretical neuroscience (Attick, 1992; Borst & Theunissen, 1999; Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1999; Quiroga & Panzeri, 2009), ideas

central to mathematical coding theory have received considerably less attention. This is in large part due to the fact that the neural code is typically regarded as a description of the mapping, or *encoding map*, between stimuli and neural responses. Because this mapping is not in general understood, identifying which features of neural responses carry the most information about a stimulus is often considered to be the main goal of neural coding theory (Bialek, Rieke, de Ruyter van Stevenick, & Warland, 1991; deCharms & Zador, 2000; Jacobs et al., 2009; London, Roth, Beeren, Häusser, & Latham, 2010). In particular, information-theoretic considerations have been used to suggest that encoding maps ought to maximize information and minimize the redundancy of stimulus representations (Attneave, 1954; Barlow, 1961; Adelesberger-Mangan & Levy, 1992; Attick, 1992; Rieke et al., 1999), although recent experiments point increasingly to high levels of redundancy in retinal and cortical codes (Puchalla, Schneidman, Harris, & Berry, 2005; Luczak, Barthó, & Harris, 2009).

In contrast, mathematical coding theory has been primarily motivated by engineering applications, where the encoding map is always assumed to be well known and can be chosen at will. The primary function of a “code” in Shannon’s original work is to allow accurate and efficient *error correction* following transmission across a noisy channel. “Good codes” do this in a highly efficient manner, so as to achieve maximal channel capacity while allowing arbitrarily accurate error correction. Mathematical coding theory grew out of Shannon’s challenge to design good codes, a question largely independent of either the nature of the information being transmitted or the specifics of the encoding map. In this perspective, redundancy is critical to the function of a code, as error correction is possible only because a code introduces redundancy into the representation of transmitted information (MacWilliams & Sloane, 1983; Huffman & Pless, 2003).

Given this difference in perspective, can mathematical coding theory be useful in neuroscience? Because of the inherent noise and variability that is evident in neural responses, it seems intuitive that enabling error correction should also be an important function of neural codes (Schneidman, Berry, Segev, & Bialek, 2006; Hopfield, 2008; Sreenivasan & Fiete, 2011). Moreover, in cases where the encoding map has become more or less understood, as in systems that exhibit robust and reliable receptive fields, we can begin to look beyond the encoding map and study the features of the neural code itself. An immediate advantage of this new perspective is that it can help clarify the role of redundancy. From the viewpoint of information theory, it may be puzzling to observe so much redundancy in the way neurons are representing information (Barlow, 1961), although the advantages of redundancy in neural coding are gaining appreciation (Barlow, 2001; Puchalla et al., 2005). Experimentally, redundancy is apparent even without an understanding of the encoding map from the fact that only a small fraction of the possible patterns of neural activity is actually observed in both stimulus-evoked and spontaneous activity (Luczak et al., 2009). On the other hand, it is generally assumed that redundancy in neural responses, as in good

codes, exists primarily to allow reliable signal estimation in the presence of noisy information transmission. This is precisely the kind of question that mathematical coding theory can address: Does the redundancy apparent in neural codes enable accurate and efficient error correction?

To investigate this question, we take a new look at neural coding from a mathematical coding theory perspective, focusing on error correction in combinatorial codes derived from neurons with idealized receptive fields. These codes can be thought of as binary codes, with 1s and 0s denoting neurons that are “on” or “off” in response to a given stimulus, and thus lend themselves particularly well to traditional coding-theoretic analyses. Although it has been recently argued that the entorhinal grid cell code may be very good for error correction (Sreenivasan & Fiete, 2011), we show that more typical receptive field codes (RF codes), including place field codes, perform quite poorly as compared to random codes with matching length, sparsity, and redundancy. The error-correcting performance of receptive field codes catches up, however, when a small tolerance to error is introduced. This error tolerance is measured in terms of a metric inherited from the stimulus space and reflects the fact that perception of parametric stimuli is often inexact. We conclude that the nature of the redundancy observed in receptive field codes cannot be fully explained as a mechanism to improve error correction, since these codes are far from optimal in this regard. On the other hand, the structure of receptive field codes does allow them to naturally encode distances between stimuli, a feature that could be beneficial for making sense of the transmitted information within the brain. We suggest that a compromise in error-correcting capability may be a necessary price to pay for a neural code whose structure not only serves error correction but must also reflect relationships between stimuli.

2 Combinatorial Neural Codes

Given a set of neurons labeled $\{1, \dots, n\} \stackrel{\text{def}}{=} [n]$, we define a *neural code* $\mathcal{C} \subset 2^{[n]}$ as a set of subsets of the n neurons, where $2^{[n]}$ denotes the set of all possible subsets. In mathematical coding theory, a *binary code* is simply a set of patterns in $\{0, 1\}^n$. These notions coincide in a natural way once we identify any element of $\{0, 1\}^n$ with its *support*,

$$\mathbf{c} \in \{0, 1\}^n \leftrightarrow \text{supp}(\mathbf{c}) \stackrel{\text{def}}{=} \{i \in [n] \mid c_i = 1\} \in 2^{[n]},$$

and we use the two notions interchangeably in the sequel. The elements of the code are called codewords: a *codeword* $\mathbf{c} \in \mathcal{C}$ corresponds to a subset of neurons and serves to represent a stimulus. Because we discard the details of the precise timing or rate of neural activity, what we mean by “neural code” is often referred to in the neural coding literature as a combinatorial code (Osborne, Palmer, Lisberger, & Bialek, 2008).

We will consider *parameters* of neural codes, such as size, length, sparsity, and redundancy. The *size* of a code \mathcal{C} is simply the total number of codewords, $|\mathcal{C}|$. The *length* of a code $\mathcal{C} \subset 2^{[n]}$ is n , the number of neurons. The (Hamming) *weight* $w_H(\mathbf{c})$ of a codeword $\mathbf{c} \in \mathcal{C}$ is the number of neurons in \mathbf{c} when viewed as a subset of $[n]$ or, alternatively, the number of 1s in the word when viewed as an element of $\{0, 1\}^n$. We define the *sparsity* s of a code as the average proportion of 1s appearing among all codewords,

$$s = \frac{1}{|\mathcal{C}|} \sum_{\mathbf{c} \in \mathcal{C}} \frac{w_H(\mathbf{c})}{n}.$$

Closely related to the size of a code \mathcal{C} is the code's *redundancy*, which quantifies the idea that typically more neurons are used than would be necessary to encode a given set of stimuli.¹ Formally, we define the redundancy ρ of a code \mathcal{C} of length n as

$$\rho = 1 - \frac{\log_2(|\mathcal{C}|)}{n}.$$

For example, the redundancy of the repetition code $\mathcal{C} = \{\emptyset, [n]\}$ of length n , consisting only of the all-zeros word and the all-ones word, is $\rho = \frac{n-1}{n}$; this may be interpreted as saying that all but one of the n neurons are extraneous. At the other end of the spectrum, the redundancy of the code $\mathcal{C} = 2^{[n]}$, consisting of all possible subsets of $[n]$, is $\rho = 0$. It is clear that ρ takes values between 0 and 1 and that any pair of codes with matching size and length will automatically have the same redundancy.²

2.1 Receptive Field Codes. Neurons in many brain areas have activity patterns that can be characterized by receptive fields. Abstractly, a *receptive field* is a map $f_i : X \rightarrow \mathbb{R}_{\geq 0}$ from a space of stimuli X to the average (nonnegative) firing rate of a single neuron, i , in response to each stimulus. Receptive fields are computed by correlating neural responses to independently measured external stimuli. We follow a common abuse of language, where both the map and its support (i.e., the subset of X where f_i takes on positive values) are referred to as receptive fields. *Convex* receptive fields are convex subsets of X . The main examples we have in mind pertain to orientation-selective neurons and hippocampal place cells. Orientation-selective neurons have tuning curves that reflect a neuron's preference for a particular

¹See Puchalla et al. (2005) and Levy and Baxter (1996) for related notions.

²In the coding theory literature, the *rate* of a code \mathcal{C} of length n is given by $\frac{\log_2(|\mathcal{C}|)}{n}$, so that the redundancy as we have defined it is simply 1 minus the rate. Because "rate" has a very different meaning in neuroscience than in coding theory, we will avoid this term and use the notion of redundancy instead.

angle (Watkins & Berkley, 1974; Ben-Yishai, Bar-Or, & Sompolinsky, 1995). Place cells are neurons that have place fields: each cell has a preferred (convex) region of the animal’s physical environment where it has a high firing rate (O’Keefe & Dostrovsky, 1971; McNaughton, Battaglia, Jensen, Moser, & Moser, 2006). Both tuning curves and place fields are examples of receptive fields.³

The neural code is the brain’s representation of the stimulus space covered by the receptive fields. When a stimulus lies in the intersection of several receptive fields, the corresponding neurons tend to co-fire while the rest remain silent. The active subset σ of neurons is a neural codeword and is identified as usual with a binary codeword c such that $\text{supp}(c) = \sigma$,

$$c = (c_1, \dots, c_n) \in \{0, 1\}^n, \quad \text{where } c_i = \begin{cases} 1, & i \in \sigma \\ 0, & i \notin \sigma \end{cases}.$$

For a given set of receptive fields on a stimulus space X , the *receptive field code* (RF code) $C \subset \{0, 1\}^n$ is simply the set of all binary codewords corresponding to stimuli in X . The *dimension* of an RF code is the dimension of the underlying stimulus space.⁴ In the case of orientation tuning curves, the stimulus space is the interval $[0, \pi)$, and the corresponding RF code is one-dimensional. In the case of place fields for an animal exploring a two-dimensional environment, the stimulus space is the environment, and the RF code is two-dimensional. From now on, we will refer to such codes as 1D RF codes and 2D RF codes, respectively.

Figure 1 shows examples of receptive fields covering one- and two-dimensional stimulus spaces. Recall that $f_i : X \rightarrow \mathbb{R}_{\geq 0}$ is the receptive field of a single neuron, and let $f = (f_1, \dots, f_n) : X \rightarrow \mathbb{R}_{\geq 0}^n$ denote the population activity map, associating with each stimulus a firing rate vector that contains the response of each neuron as dictated by the receptive fields. For a given choice of threshold θ , we can define a binary response map, $\Phi : X \rightarrow \{0, 1\}^n$, from the stimulus space X to codewords by

$$\Phi_i(x) = \begin{cases} 1 & \text{if } f_i(x) \geq \theta \\ 0 & \text{if } f_i(x) < \theta \end{cases}.$$

The corresponding RF code C is the image of Φ . Notice that many stimuli will produce the same binary response; in particular, Φ maps an entire

³In the vision literature, the term *receptive field* is reserved for subsets of the visual field. Here we use the term in a more general sense that is applicable to any modality, as in Curto and Itskov (2008).

⁴Note that this is distinct from the notion of “dimension of a code” in the coding theory literature.

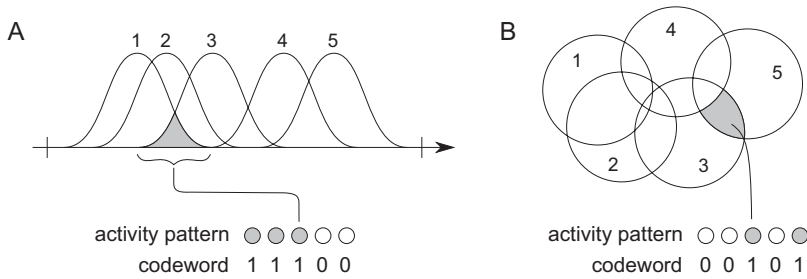


Figure 1: Receptive field overlaps determine codewords in 1D and 2D RF codes. (A) Neurons in a 1D RF code have receptive fields that overlap on a line segment (or circle). Each stimulus on the line corresponds to a binary codeword, with 1s corresponding to neurons whose receptive fields include the stimulus and 0s for neurons that are not active in response to the stimulus. (B) Neurons in a 2D RF code, such as a place field code, have receptive fields that partition a two-dimensional region into nonoverlapping intersection regions, such as the shaded area. All stimuli within one of these regions will activate the same set of neurons, and hence have the same corresponding codeword.

region of intersecting receptive fields to the same codeword, and so Φ is far from injective.

2.2 Comparison Codes. In order to analyze the performance of RF codes, we use two types of randomly generated comparison codes with matching size, length, and sparsity. In particular, these codes have the same redundancy as their corresponding RF codes. We choose random codes as our comparison codes for three reasons. First, as demonstrated by Shannon (1948) in the proof of his channel coding theorem, random codes are expected to have near-optimal performance. Second, the parameters can be tuned to match those of the RF codes; we describe the two ways in which we do this. Finally, random codes are a biologically reasonable alternative for the brain, since they may be implemented by random neural networks.

2.2.1 Shuffled Codes. Given a RF code \mathcal{C} , we generate a *shuffled code* $\tilde{\mathcal{C}}$ in the following manner. Fix a collection of permutations $\{\pi_{\mathbf{c}} \mid \mathbf{c} \in \mathcal{C}\}$ such that $(c_{\pi_{\mathbf{c}}(1)}, \dots, c_{\pi_{\mathbf{c}}(n)}) \neq (c_{\pi_{\mathbf{c}'}(1)}, \dots, c_{\pi_{\mathbf{c}'}(n)})$ for all distinct $\mathbf{c}, \mathbf{c}' \in \mathcal{C}$, and set $\tilde{\mathcal{C}} = \{(c_{\pi_{\mathbf{c}}(1)}, \dots, c_{\pi_{\mathbf{c}}(n)}) \mid \mathbf{c} = (c_1, \dots, c_n) \in \mathcal{C}\}$.⁵ The shuffled code $\tilde{\mathcal{C}}$ has the same length, size, and weight distribution (and hence the same sparsity and redundancy) as \mathcal{C} . In our simulations, each permutation $\pi_{\mathbf{c}}$ is chosen uniformly at random with the modification that a new permutation is

⁵If the same permutation were used to shuffle all codewords, the resulting permutation-equivalent code would be nothing more than the code obtained from a relabeling of the neurons.

selected if the resulting shuffled codeword has already been generated. This ensures that no two codewords of \mathcal{C} correspond to the same word in the shuffled code.

2.2.2 Random Constant-Weight Codes. Constant-weight codes are subsets of $\{0, 1\}^n$ in which all codewords have the same weight. Given an RF code \mathcal{C} on n neurons, we compute the average weight of the codewords in \mathcal{C} and round this to obtain an integer w . We then generate a constant-weight code by randomly choosing subsets of size w from $[n]$. These subsets give the positions of the codeword that are assigned a 1, and the remaining positions are all assigned zeros. This process is repeated until $|\mathcal{C}|$ distinct codewords are generated, and the resulting code is then a random constant-weight code with the same length, size, and redundancy as \mathcal{C} , and approximately the same sparsity as \mathcal{C} .

3 Stimulus Encoding and Decoding

3.1 The Mathematical Coding Theory Perspective. The central goal of this letter is to analyze our main examples of combinatorial neural codes, 1D and 2D RF codes, from a mathematical coding theory perspective. We draw on this field because it provides a complementary perspective on the nature and function of codes that is unfamiliar to most neuroscientists. We first discuss the standard paradigm of coding theory and then explain the function of codes from this perspective. Note that to put neural codes into this framework, we must discretize the stimulus space and encoding map so that we have an injective map from the set of stimuli to the code; we describe this in the next section.

Figure 2A illustrates the various stages of information transmission using the standard coding theory paradigm, adapted for RF codes. A stimulus $x \in \check{X}$ gets mapped to a neural codeword $\mathbf{c} \in \mathcal{C}$ under an (injective) encoding map $\varphi : \check{X} \rightarrow \mathcal{C}$, where \check{X} is the (discretized) stimulus space. This map sends each stimulus to a neural activity pattern that is considered to be the *ideal response* of a population of neurons. The codeword, viewed as a string of 0s and 1s, then passes through a noisy channel, where each 0/1 bit may be flipped with some probability. A $1 \mapsto 0$ flip corresponds to a neuron in the ideal response pattern failing to fire, while a $0 \mapsto 1$ flip corresponds to a neuron firing when it is not supposed to. The resulting word is not necessarily a codeword and is referred to as the *received word*. This noisy channel output is then passed through a *decoder* to yield an estimate $\hat{\mathbf{c}} \in \mathcal{C}$ for the original codeword \mathbf{c} , corresponding to an estimate of the ideal response. Finally, if an estimate $\hat{x} \in \check{X}$ of the original stimulus is desired, the inverse of the encoding map may be applied to the estimated codeword $\hat{\mathbf{c}}$. Because the brain has access only to neural activity patterns, we will consider the ideal response as a proxy for the stimulus itself; the estimated neural codeword

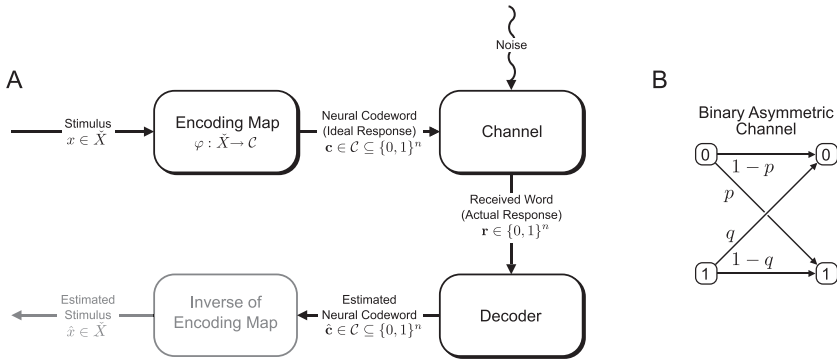


Figure 2: Stimulus encoding and decoding from a mathematical coding theory perspective. Here \check{X} is the discretized stimulus space and \mathcal{C} is the neural code. (A) A stimulus $x \in \check{X}$ is mapped, via an (injective) encoding map $\varphi: \check{X} \rightarrow \mathcal{C}$, to a neural codeword $\mathbf{c} \in \mathcal{C} \subseteq \{0, 1\}^n$. The effect of noise is modeled by passing the codeword through a noisy channel. The output of the channel is the received word $\mathbf{r} \in \{0, 1\}^n$, representing the actual response of the population to a particular presentation of the stimulus; typically the received word is not an element of the neural code. In order to estimate the ideal response, the received word is passed through a decoder to produce an estimated neural codeword $\hat{\mathbf{c}} \in \mathcal{C}$. Decoding is considered to be “correct” if this codeword matches the ideal response for the original stimulus. Finally, the inverse of the encoding map can be used to identify the estimated neural codeword with an estimate for the stimulus. (B) The binary asymmetric channel (BAC) acts independently on individual bits of a binary word in $\{0, 1\}^n$. The effect of noise is to flip 0/1 bits according to a false-positive probability p and a false-negative probability q .

thus represents the brain’s estimate of the stimulus, and so we can ignore this last step.

The mathematical coding theory perspective on stimulus encoding and decoding has several important differences from the way neuroscientists typically think about neural coding. First, a clear distinction is made between a code, which is simply a set of codewords (or neural response patterns) devoid of any intrinsic meaning, and the encoding map, which is a function that assigns a codeword to each element in the set of objects to be encoded. Second, this map is always deterministic, as the effects of noise are considered to arise purely from the transmission of codewords through a noisy channel.⁶ For neuroscientists, the encoding of a signal into

⁶In engineering applications, one can always assume the encoding map is deterministic. In the neuroscience context, however, it may be equally appropriate to use a probabilistic encoding map.

a pattern of neural activity is itself a noisy process, and so the encoding map and the channel are difficult to separate. If we consider the output of the encoding map to be the ideal response of a population of neurons, however, it is clear that actual response patterns in the brain correspond not to codewords but rather to received words. (The ideal response, on the other hand, is always a codeword and corresponds intuitively to the average response across many trials of the same stimulus.) In the case of RF codes, a natural encoding map sends each stimulus to the codeword corresponding to the subset of neurons that contain the stimulus in their receptive fields. In the case of the random comparison codes, an encoding map that assigns codewords to stimuli is chosen randomly (details are given in the next section).

Another important difference offered by the coding theory perspective is in the process of decoding. Given a received word, the objective of the decoder is to estimate the original codeword that was transmitted through the channel. In the case of neural codes, this amounts to taking the actual neural response and producing an estimate of the ideal response, which serves as a proxy for the stimulus. The function of the decoder is therefore to *correct errors* made by transmission through the noisy channel. In a network of neurons, this would be accomplished by network interactions that evolve the original neural response (the received word) to a closely related activity pattern (the estimated codeword) that corresponds to an ideal response for a likely stimulus.

This leads us to the coding theory perspective on the function (or purpose) of a code. Error correction is possible only when errors produced by the channel lead to received words that are not themselves codewords, and it is most effective when codewords are “far away” from each other in the space of all words, so that errors can be corrected by returning the “nearest” codeword to the received word. The function of a code, therefore, is to represent information in a way that allows accurate error correction in a high percentage of trials. The fact that there is redundancy in how a code represents information is therefore a positive feature of the code rather than an inefficiency, since it is precisely this redundancy that makes error correction possible.

3.2 Encoding Maps and the Discretization of the Stimulus Space. In the definition of RF codes above, the stimulus space X is a subset of Euclidean space, having a continuum of stimuli. Via the associated binary response maps, a set of n receptive fields partitions the stimulus space X into distinct overlap regions, such as the shaded regions in Figure 1. For each codeword $\mathbf{c} \in \mathcal{C}$, there is a corresponding overlap region $\Phi^{-1}(\mathbf{c})$, all of whose points map to \mathbf{c} . The combinatorial code \mathcal{C} therefore has limited resolution and is not able to distinguish among stimuli in the same overlap region. This leads to a natural discretization of the stimulus space, where we

assign a single representative stimulus, the center of mass,⁷ to each overlap region, and we write

$$\check{x}(\mathbf{c}) \stackrel{\text{def}}{=} \frac{\int_{\Phi^{-1}(\mathbf{c})} x \, dx}{\int_{\Phi^{-1}(\mathbf{c})} dx},$$

where $x \in X$ refers to a one- or two-dimensional vector, and the integral is either a single or double integral, depending on the context. In practice, for 2D RF codes we use a fine grid to determine the center of mass associated with each codeword (see appendix B.2).

From now on, we will use the term *stimulus space* to refer to the discretized stimulus space:

$$\check{X} = \{\check{x}(\mathbf{c}) \mid \mathbf{c} \in \mathcal{C}\} \subset X.$$

Note that $|\check{X}| = |\mathcal{C}|$, so we now have a one-to-one correspondence between stimuli and codewords. The restriction of the binary response map Φ to the discretized stimulus space is the encoding map of the RF code,

$$\varphi = \Phi|_{\check{X}} : \check{X} \rightarrow \mathcal{C}.$$

Note that unlike Φ , the encoding map φ is injective, and so its inverse is well defined. This further supports the idea, introduced in the previous section, that the ideal response estimate returned by the decoder can serve as a proxy for the stimulus itself.

In the case of the comparison codes, we use the same discretized stimulus space \check{X} as in the corresponding RF code and associate a codeword to each stimulus using a random (one-to-one) encoding map $\varphi : \check{X} \rightarrow \mathcal{C}$. This map is generated by ordering both the stimuli in \check{X} and the codewords in the random code \mathcal{C} , and then selecting a random permutation to assign a codeword to each stimulus.

3.3 The Binary Asymmetric Channel. In all our simulations, we model the channel as a *binary asymmetric channel* (BAC). As seen in Figure 2B, the BAC is defined by a false-positive probability p , the probability of a 0 being flipped to a 1, and a false-negative probability q , the probability of a 1 being flipped to a 0. Since errors are always assumed to be less likely than faithful transmission, we assume $p, q < 1/2$. The channel operates on each individual bit, but it is customary to extend it to operate on full codewords via

⁷Although many of the overlap regions will be nonconvex, instances of the center of mass falling outside the corresponding region will be rare enough that this pathological case need not be considered.

the assumption that each bit is affected independently. This is reasonable in our context because it is often assumed (though not necessarily believed) that neurons within the same area experience independent and identically distributed noise. The BAC has as special cases two other channels commonly considered in mathematical coding theory: $p = q$ gives the binary symmetric channel (BSC), and $p = 0$ reduces to the Z-channel.

We will assume $p \leq q$, meaning that it is at least as likely that a 1 will flip to a 0 as it is that a 0 will flip to a 1. This is because the failure of a neuron to fire (due to, for example, synaptic failure) is considered more likely than a neuron firing when it is not supposed to. Recall that the sparsity s reflects the probability that a neuron fires due to error-free transmission. We will require $p < s$, as a false-positive response should be less likely than a neuron firing appropriately in response to a stimulus. Finally, since our neural codes are assumed to be sparse, we require $s < 1/2$. In summary, we assume

$$p \leq q < 1/2, \quad \text{and} \quad p < s < 1/2.$$

Note that the probability of an error across this channel depends on the sparsity of the code. For a given bit (or neuron), the probability of an error occurring during transmission across the BAC is $p(1 - s) + qs$, assuming that all codewords are transmitted with equal probability and all neurons participate in approximately the same number of codewords.

3.4 The ML and MAP Decoders. A decoder takes an actual response (or received word) $\mathbf{r} \in \{0, 1\}^n$ and returns a codeword $\hat{\mathbf{c}} \in \mathcal{C}$ that is an estimate of the ideal response (or sent word), $\mathbf{c} \in \mathcal{C}$. For each combination of code and channel, the decoder that is optimal, in the sense of minimizing errors, is the one that returns a codeword $\hat{\mathbf{c}}$ with maximal probability of having been sent, given that \mathbf{r} was received.⁸ This is called the maximum a posteriori (MAP) decoder, also known in the neuroscience literature as Bayesian inference (Ma, Beck, Latham, & Pouget, 2006) or an ideal observer decoder (Deneve, Latham, & Pouget, 1999):

$$\hat{\mathbf{c}}_{\text{MAP}} = \arg \max_{\mathbf{c} \in \mathcal{C}} P(\text{sent} = \mathbf{c} \mid \text{rec} = \mathbf{r}).$$

Although always optimal, this decoder can be difficult to implement in the neural context, as it requires knowing the probabilities $P(\text{sent} = \mathbf{c})$ for each codeword, which is equivalent to knowing the probability distribution of stimuli.

⁸In all of our decoders, we assume that ties are broken randomly, with uniform distribution on equally optimal codewords.

The maximum likelihood (ML) decoder,

$$\hat{\mathbf{c}}_{\text{ML}} = \arg \max_{\mathbf{c} \in \mathcal{C}} P(\text{rec} = \mathbf{r} \mid \text{sent} = \mathbf{c}),$$

is much more easily implemented. ML decoding is often used in lieu of MAP decoding because it amounts to optimizing a simple function that can be computed directly from the channel parameters. As shown in appendix A.1, on the BAC with parameters p and q we have

$$\hat{\mathbf{c}}_{\text{ML}} = \arg \max_{\mathbf{c} \in \mathcal{C}} \left[(\mathbf{c} \cdot \mathbf{r}) \ln \left(\frac{(1-p)(1-q)}{pq} \right) - w_H(\mathbf{c}) \ln \left(\frac{1-p}{q} \right) \right].$$

The ML decoder thus returns a codeword \mathbf{c} that maximizes the dot product $\mathbf{c} \cdot \mathbf{r}$ with the received word \mathbf{r} , subject to a penalty term proportional to its weight $w_H(\mathbf{c})$. In other words, it returns the codeword that maximizes the number of matching 1s with \mathbf{r} while minimizing the introduction of additional 1s.

For $p = q < 1/2$, as on the BSC, the maximization becomes (see appendix A.1)

$$\hat{\mathbf{c}}_{\text{ML}} = \arg \min_{\mathbf{c} \in \mathcal{C}} d_H(\mathbf{c}, \mathbf{r}),$$

where $d_H(\mathbf{c}, \mathbf{r}) = |\{i \in [n] \mid c_i \neq r_i\}|$ is the *Hamming distance* between two words in $\{0, 1\}^n$. This is the well-known result that ML decoding is equivalent to nearest neighbor decoding, with respect to Hamming distance, on the BSC.

3.5 An Approximation of MAP Decoding for Sparse Codes. In cases where all codewords are sent with equal probability, it is easy to see from Bayes' rule that $\hat{\mathbf{c}}_{\text{ML}} = \hat{\mathbf{c}}_{\text{MAP}}$ (see appendix A.2). When codewords are not equally likely, MAP decoding will outperform ML decoding, but it is impractical in the neural context because we cannot know the exact probability distribution on stimuli. In some cases, however, it may be possible to approximate MAP decoding, leading to a decoder that outperforms ML while being just as easy to implement. Here we illustrate this possibility in the case of sparse codes, where sparser (lower-weight) codewords are more likely.

For the BAC with parameters p and q and a code \mathcal{C} with sparsity s , we can approximate MAP decoding as the following maximization (see appendix A.2):

$$\hat{\mathbf{c}}_{\text{MAP}} \approx \arg \max_{\mathbf{c} \in \mathcal{C}} \left[(\mathbf{c} \cdot \mathbf{r}) \ln \left(\frac{(1-p)(1-q)}{pq} \right) - w_H(\mathbf{c}) \ln \left(\frac{(1-p)(1-s)}{qs} \right) \right].$$

Since we assume that $s < 1/2$, we see that the difference between this $\hat{\mathbf{c}}_{\text{MAP}}$ approximation and $\hat{\mathbf{c}}_{\text{ML}}$ is only that the coefficient of the $-w_H(\mathbf{c})$ penalty term is larger and now depends on s . Clearly, this decoder is no more difficult to implement than the ML decoder.

Figure 3 shows the results of two simulations comparing the MAP approximation to ML decoding on a 2D RF code. In the first case, Figure 3A, the probability distribution is biased toward sparser codewords, corresponding to stimuli covered by fewer receptive fields. Here we see that the MAP approximation significantly outperforms ML decoding. In the second case, Figure 3B, all codewords are equally likely. As expected, ML decoding outperforms the MAP approximation in this case, since it coincides with MAP decoding. When we consider a biologically plausible probability distribution that is biased toward codewords with larger regions $\Phi^{-1}(\mathbf{c})$ in the stimulus space, we find that ML decoding again outperforms the MAP approximation (see appendix A.2 and Figure 8), even though there is a significant correlation between larger region size and sparser codewords. Thus, we will restrict ourselves to considering ML decoding in the sequel; for simplicity, we will assume all codewords are equally likely.⁹

4 The Role of Redundancy in RF Codes

The function of a code from the mathematical coding theory perspective is to represent information in a way that allows errors in transmission to be corrected with high probability. In classical mathematical coding theory, decoding reduces to finding the closest codeword to the received word, where “closest” is measured by a metric appropriate to the channel. If the code has large minimum distance between codewords, then many errors can occur without affecting which codeword will be chosen by the decoder (Huffman & Pless, 2003). If the elements of a binary code are closely spaced within $\{0, 1\}^n$, errors will be more difficult to decode because there will often be many candidate codewords that could have reasonably resulted in a given received word.

When the redundancy of a code is high, the ratio of the number of codewords to the total number of vectors in $\{0, 1\}^n$ is low, and so it is possible to achieve a large minimum distance between codewords. Nevertheless, high redundancy of a code does not guarantee large minimum distance, because even highly redundant codes may have codewords that are spaced closely together. For this reason, high redundancy does not guarantee good error-correcting properties. This leads us to the natural question: Does the

⁹In cases where the distribution of stimuli is not uniform, our analysis would proceed in exactly the same manner with one exception: instead of using the ML decoder, which may no longer be optimal, we would use the MAP decoder or an appropriate approximation to MAP that is tailored to the characteristics of the codeword distribution.

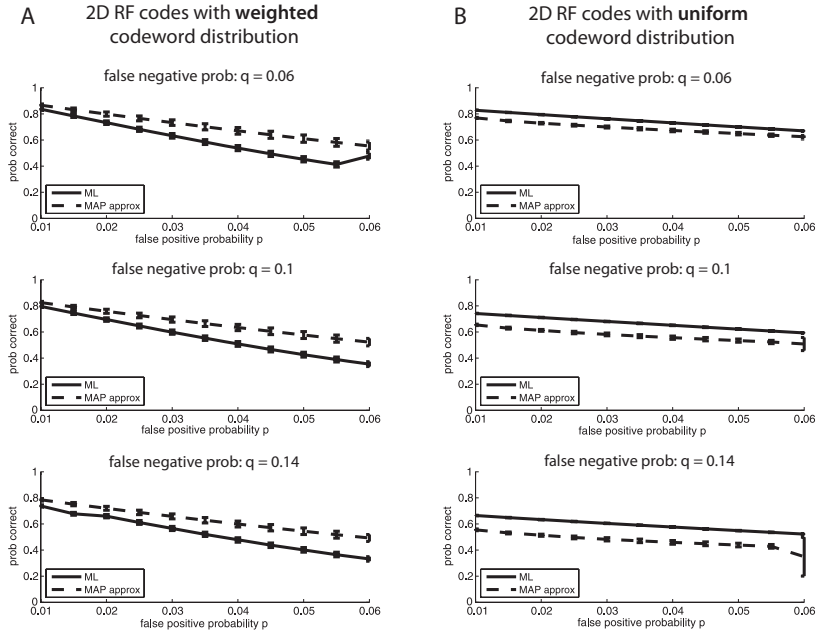


Figure 3: Approximate MAP decoding outperforms ML decoding for a weighted distribution of codewords. (A) With the false-negative probability q fixed at $q = 0.06$ (top), $q = 0.1$ (middle), and $q = 0.14$ (bottom), the false-positive probability p was varied in increments of 0.005 from 0.01 to 0.06 to produce different channel conditions for the BAC. On each channel, the performance of 100 2D RF codes of length 75 and mean sparsity $s = 0.069$ was assessed using both the standard ML decoder and our approximation to the MAP decoder. For each BAC condition and each code, 10,000 codewords were selected according to a weighted probability distribution, where the probability of sending codeword c was proportional to $s^{w_H(c)}(1-s)^{n-w_H(c)}$, as would be expected if 1s and 0s were sent through the channel with independent probabilities dictated by the sparsity. The fraction of correctly decoded words was then averaged across the 100 codes, with error bars denoting standard deviations. The MAP approximation consistently outperformed ML decoding for all channel conditions. (B) Same as in panel A, but this time codewords were selected according to a uniform probability distribution, with each codeword equally likely. In this case, ML decoding is equivalent to exact MAP decoding, which is always optimal. As expected, ML decoding outperformed approximate MAP decoding for each channel condition. Note that the error bars for ML decoding in this case are extremely small.

high redundancy of RF codes result in effective error correction? The answer depends to some extent on the particular decoder that is used. In the simulations that follow, we use ML decoding to test how well RF codes

correct errors. We assume that all codewords within a code are equally likely, and hence ML decoding is equivalent to (optimal) MAP decoding. It has been suggested that the brain may actually implement ML or MAP decoding (Deneve et al., 1999; Ma et al., 2006), but even if this decoder were not biologically plausible, it is the natural decoder to use in our simulations as it provides an upper bound on the error-correcting performance of RF codes.

4.1 RF Code Redundancy Does Not Yield Effective Error Correction.

To test the hypothesis that the redundancy of RF codes enables effective error correction, we generated 1D and 2D RF codes having 75 neurons each (see appendix B). For each RF code, we also generated two random comparison codes: a shuffled code and a random constant-weight code with matching parameters. These codes were tested on the BAC for a variety of channel parameters (values of p and q). For each BAC condition and each code, 10,000 codewords selected uniformly at random were sent across the noisy channel and then decoded using ML decoding. If the decoded word exactly matched the original sent word, the decoding was considered “correct”; if not, there was a failure of error correction.

Figure 4 shows the fraction of correctly decoded transmissions for fixed values of q and a range of p values in the case of 1D RF codes (see Figure 4A) and 2D RF codes (see Figure 4B), together with the performance of the comparison codes. In each case, the RF codes had significantly worse performance (less than 80% correct decoding in all cases) than the comparison codes, whose performances were near-optimal for low values of p . Repeating this analysis for different values of q yielded similar results (not shown).

As previously mentioned, in the case of the BSC, nearest neighbor decoding with respect to Hamming distance coincides with ML decoding. Thus, in the case of a symmetric channel, codes perform poorly precisely when their minimum Hamming distance is small. Even though nearest neighbor decoding with respect to Hamming distance does not coincide with ML decoding on the BAC when $p \neq q$, decoding errors are still more likely to occur if codewords are close together in Hamming distance. Indeed, the poor performance of RF codes can be attributed to the very small distance between a codeword and its nearest neighbors. Since codewords correspond to regions defined by overlapping receptive fields, the Hamming distance between a codeword and its nearest neighbor is typically 1 in an RF code, which is the worst-case scenario.¹⁰ In contrast, codewords in the random comparison codes are distributed much more evenly throughout the ambient space $\{0, 1\}^n$. While there is no guarantee that the minimum

¹⁰Note that this situation would be equally problematic if we considered the full firing rate information instead of a combinatorial code. This is because small changes in firing rates would tend to produce equally valid codewords, making error detection and correction just as difficult.

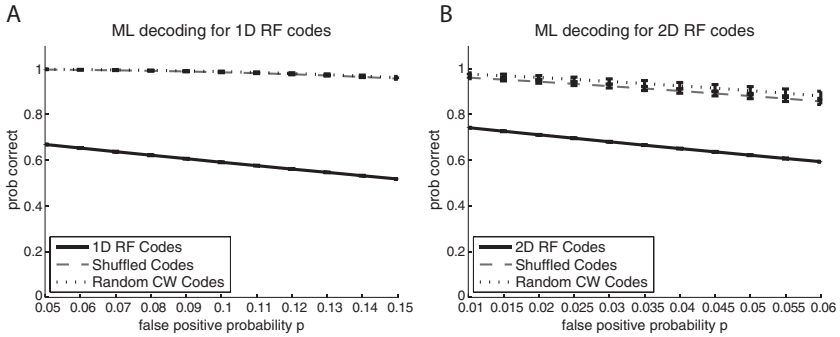


Figure 4: RF codes perform poorly under standard ML decoding. (A) With the false-negative probability $q = 0.2$ fixed, the false-positive probability p was varied in increments of 0.01 from 0.05 to 0.15 to produce different channel conditions for the BAC. On each channel, the performance of 100 1D RF codes of length 75, with mean sparsity $s = 0.165$, was compared to the performance of 100 shuffled codes and 100 random constant-weight codes of matched parameters. For each BAC condition and each code, 10,000 codewords selected uniformly at random were transmitted across the BAC and then decoded with ML decoding. The trajectories show the average performance of each code type (across the 100 sample codes) in terms of the proportion of received words that were correctly decoded. Error bars show the magnitude of one standard deviation from the average performance and are very small. While the shuffled and random constant-weight codes had similar, near-optimal performance, the 1D RF codes performed quite poorly in comparison. (B) Same as in panel A, but for 2D RF codes of length 75 and mean sparsity $s = 0.069$. Here q was fixed at 0.1, while p varied from 0.01 to 0.06 in increments of 0.005. Again, RF codes performed significantly worse than the shuffled and random codes with matched parameters.

distance on these codes is high, the *typical* distance between a codeword and its nearest neighbor is high, leading to near-optimal performance.

4.2 RF Code Redundancy Reflects the Geometry of the Stimulus Space. Given the poor error-correcting performance of RF codes, it seems unlikely that the primary function of RF code redundancy is to enable effective error correction. As outlined in the previous section, the poor performance of RF codes is the result of the very small Hamming distances between a codeword and its nearest neighbors. While these small Hamming distances are problematic for error correction, they may prove valuable in reflecting the distance relationships between stimuli, as determined by a natural metric on the stimulus space.

To further investigate this possibility, we first define a new metric on the code that assigns distances to pairs of codewords according to the

distances between the stimuli that they represent. If $\mathbf{c}, \mathbf{c}' \in \mathcal{C}$ are codewords and $\varphi : \check{X} \rightarrow \mathcal{C}$ is the (injective) encoding map, then we define the *induced stimulus space metric* $d_{\text{stim}} : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}$ by

$$d_{\text{stim}}(\mathbf{c}, \mathbf{c}') = d(\varphi^{-1}(\mathbf{c}), \varphi^{-1}(\mathbf{c}')),$$

where d is the natural metric on the (discretized) stimulus space \check{X} . For example, in the case of 2D RF codes, the stimulus space is the two-dimensional environment, and the natural metric is the Euclidean metric; in the case of 1D RF codes, the stimulus space is $[0, \pi)$, and the natural metric is the difference between angles, where 0 and π have been identified so that, for example, $d(\pi/6, 5\pi/6) = \pi/3$.

To characterize the relationship between d_{stim} and d_H on RF codes, we performed correlation analyses between these metrics on 2D RF codes and corresponding random comparison codes. For each code, we computed d_{stim} and d_H for all pairs of codewords and then computed the correlation coefficient between their values. Figure 5A shows a scatter plot of d_{stim} versus d_H values for a single 2D RF code; the high correlation is easily seen by eye. In contrast, the same analysis for a corresponding shuffled code (see Figure 5B) and a random constant-weight code (see Figure 5C) revealed no significant correlation between d_{stim} and d_H . Repeating this analysis for the receptive field and comparison codes used in Figure 4 resulted in very similar results (see Figure 5D). Thus, the codewords in RF codes appear to be distributed across $\{0, 1\}^n$ in a way that captures the geometry of the underlying stimulus space rather than in a manner that guarantees high distance between neighboring codewords.

Previous work has shown that the structure of a place field code (i.e., a 2D RF code) can be used to extract topological and geometric features of the represented environment (Curto & Itskov, 2008). We hypothesize that the primary role of RF code redundancy may be to reflect the geometry of the underlying stimulus space and that the poor error-correcting performance of RF codes may be a necessary price to pay for this feature. This poor error correction may be mitigated, however, when we reexamine the role that stimulus space geometry plays in the brain's perception of parametric stimuli.

5 Decoding with Error Tolerance in RF Codes

5.1 Error Tolerance Based on the Geometry of Stimulus Space. The brain often makes errors in estimating stimuli (van der Heijden, van der Geest, deLeeuw, Krikke, & Müsseler, 1999; Prinzmetal, Shimura, & Mikolinski, 2001; Huttenlocher, Hedges, Lourenco, Crawford, & Corrigan, 2007); these errors are considered tolerable if they result in the perception of nearby stimuli. For example, an angle of 32 degrees might be perceived as a 30 degree angle, or a precise position (x, y) in the plane might be perceived as $(x + \varepsilon_x, y + \varepsilon_y)$. If the errors are relatively small, as measured by

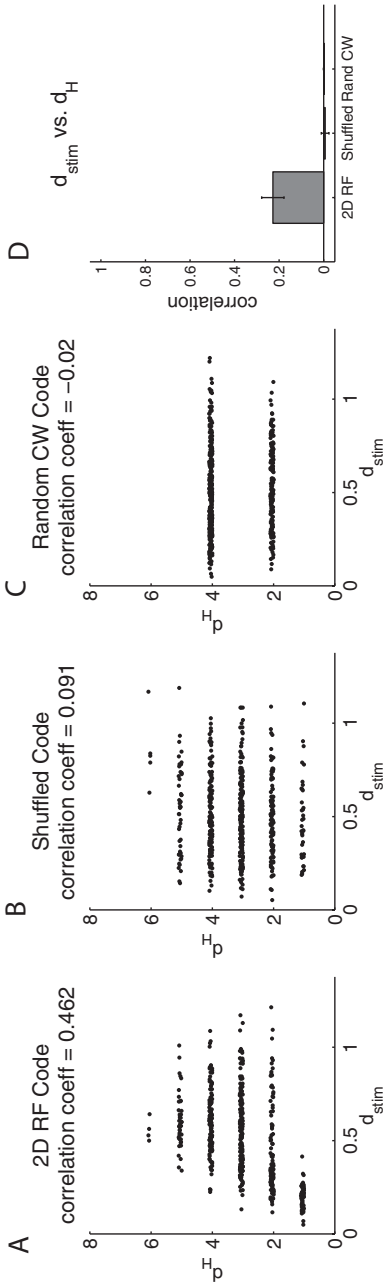


Figure 5: RF codes reflect the geometry of the stimulus space. (A) The scatter plot shows the high level of correlation (corr. coeff. = 0.462) between d_{stim} and d_H for an example 2D RF code of length 10 and sparsity $s = 0.188$. A single code of length 10 was chosen from those generated for Figure 7 to make the correlation more visually apparent. Each point in the scatter plot corresponds to a pair of distinct codewords. Random noise of size at most 0.001 in each dimension was added to each data point in order to separate points having exactly the same values of d_H and d_{stim} . (B) Same as panel A, but for a shuffled code with matched parameters. (C) Same as in panel A, but for a random constant-weight code with matched parameters. (D) Average correlation coefficient between d_{stim} and d_H for 100 2D RF codes of length 75 and mean sparsity $s = 0.069$, and for 100 shuffled and random constant-weight code of matched parameters (the same codes were used in Figure 4). While the Hamming distance d_H correlates significantly with stimulus space distance d_{stim} in the case of RF codes, there is no correlation in the case of the random or shuffled codes.

a natural metric on the stimulus space, it is reasonable to declare the signal transmission to have been successful rather than incorrect. To do this, we introduce the notion of error tolerance into our stimulus encoding and decoding paradigm using the induced stimulus space metric d_{stim} . Specifically, we can decode with an error tolerance of δ by declaring decoding to be “correct” if the decoded word $\hat{\mathbf{c}}$ is within δ of the original sent word \mathbf{c} :

$$d_{\text{stim}}(\hat{\mathbf{c}}, \mathbf{c}) < \delta.$$

This corresponds to the perceived stimulus being within a distance δ of the actual stimulus.

5.2 RF Codes “Catch Up” to Comparison Codes When Decoding with Error Tolerance. We next investigated whether the performance of RF codes improved, as compared to the comparison codes with matching parameters, when an error tolerance was introduced. For each 1D RF code and each 2D RF code used in Figure 4, we repeated the analysis using fixed channel parameters and varying instead the error tolerance with respect to the induced stimulus space metric d_{stim} . We found that RF codes quickly catch up to the random comparison codes when a small tolerance to error is introduced (see Figures 6A and 6B). In some cases, the performance of the RF codes even surpasses that of the random comparison codes.

In order to verify that the catch-up effect is not merely an artifact resulting from the assignment of random encoding maps to the comparison codes, we repeated the above analysis using Hamming distance d_H instead of d_{stim} , thus completely eliminating the influence of the encoding maps. The Hamming distance between codewords in a sparse code typically ranges from 0 to about twice the average weight, which corresponds to $d_H = 25$ for the 1D RF codes and $d_H = 10$ for the 2D RF codes considered here. The probability of correct decoding using an error tolerance measured by Hamming distance yielded similar results, with RF codes catching up to the random comparison codes for relatively small error tolerances (see Figures 6C and 6D). This suggests that errors in transmission and decoding for RF codes result in codewords that are close to the correct word not only in the induced stimulus space metric but also in Hamming distance.

The question that remains is now: Why do RF codes catch up?

5.3 ML Similarity and ML Distance. In order to gain a better understanding of why the performance of RF codes catches up to that of the random comparison codes when we allow for error tolerance, we introduce the notions of *ML similarity* and *ML distance*.¹¹ Roughly speaking, the ML

¹¹Another distance measure on neural codes was recently introduced in Tkačik, Granot-Atedgi, Segev, and Schneidman (2012).

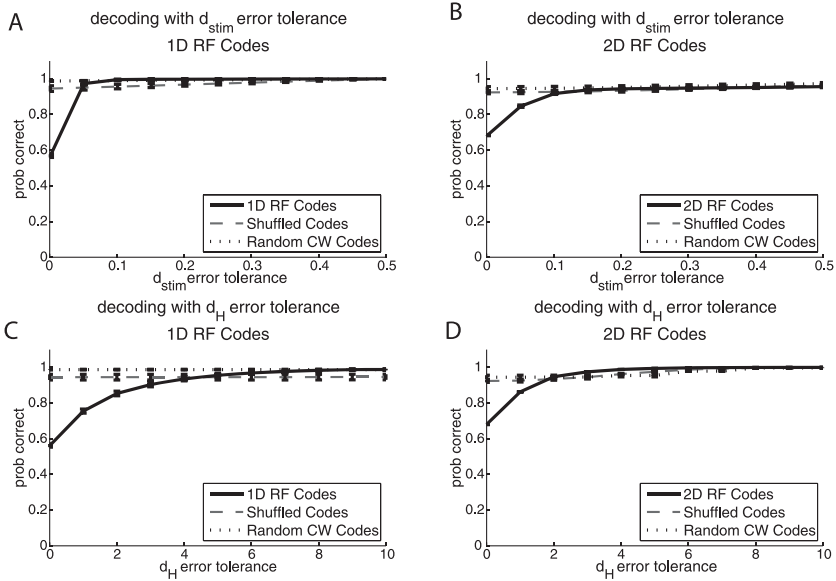


Figure 6: RF codes catch up in error-correcting performance when an error tolerance is introduced. (A) For a fixed BAC condition ($p = 0.1$, $q = 0.2$), the performance of 100 1D RF codes of length 75 was compared to the performance of 100 shuffled codes and 100 random constant-weight codes of matched parameters (the same codes were used in Figure 4). For each code, 10,000 codewords selected uniformly at random were transmitted across the BAC and then decoded with ML decoding. For each level of error tolerance, decoding was considered to be “correct” if the estimated word was within the given stimulus space distance, d_{stim} , of the correct word. Error bars show the magnitude of one standard deviation from the average performance. (B) Same as in panel A, but for 100 2D RF codes of length 75 and corresponding shuffled and random constant-weight codes of matched parameters (again, same codes as in Figure 4). Here the channel condition was fixed at $p = 0.03$ and $q = 0.1$. (C–D) These plots are analogous to panels A–B, but with the error tolerance measured using Hamming distance, d_H , rather than d_{stim} . In each case, the RF codes catch up in error-correcting performance for a small tolerance to error, at times even outperforming the shuffled and random constant-weight codes. Note: The maximum possible value for d_{stim} is normalized to be 1 in the 1D case, while in the 2D case, it is $\sqrt{2}$, corresponding to the maximum distance between two points in a 1×1 square box environment.

similarity between two codewords \mathbf{a} and \mathbf{b} is the probability that \mathbf{a} and \mathbf{b} will be confused in the process of transmission across the BAC and then running the channel output through an ML decoder. More precisely, let \mathbf{r}_a and \mathbf{r}_b be the outputs of the channel when \mathbf{a} and \mathbf{b} are input, respectively.

Note that \mathbf{r}_a (resp. \mathbf{r}_b) is randomly chosen from $\{0, 1\}^n$, with probability distribution determined by the channel parameters p and q and by the sent word \mathbf{a} (resp. \mathbf{b}). By definition, any ML decoder will return an ML codeword given by $\arg \max_{\mathbf{c} \in \mathcal{C}} P(\text{rec} = \mathbf{r} \mid \text{sent} = \mathbf{c})$ when \mathbf{r} is received from the channel, but this ML codeword need not be unique. To account for this, let $\lambda(\mathbf{r})$ be the set of ML codewords corresponding to the received word \mathbf{r} . We then define the *ML similarity* between codewords \mathbf{a} and \mathbf{b} to be the probability that the same word will be chosen (uniformly at random) from each of the sets $\lambda(\mathbf{r}_a)$ and $\lambda(\mathbf{r}_b)$:¹²

$$\mu_{\text{ML}}(\mathbf{a}, \mathbf{b}) \stackrel{\text{def}}{=} \sum_{\mathbf{r}_a} \sum_{\mathbf{r}_b} P(\text{rec} = \mathbf{r}_a \mid \text{sent} = \mathbf{a}) \cdot P(\text{rec} = \mathbf{r}_b \mid \text{sent} = \mathbf{b}) \frac{|\lambda(\mathbf{r}_a) \cap \lambda(\mathbf{r}_b)|}{|\lambda(\mathbf{r}_a)| |\lambda(\mathbf{r}_b)|}.$$

In other words, $\mu_{\text{ML}}(\mathbf{a}, \mathbf{b})$ is the probability that if \mathbf{a} and \mathbf{b} are each sent across the channel, then the same codeword will be returned in each case by the decoder. In particular, $\mu_{\text{ML}}(\mathbf{a}, \mathbf{a})$ is the probability that the same word will be returned after sending \mathbf{a} twice across the channel and decoding. Note that typically, $\mu_{\text{ML}}(\mathbf{a}, \mathbf{a}) < 1$ and $\mu_{\text{ML}}(\mathbf{a}, \mathbf{a}) \neq \mu_{\text{ML}}(\mathbf{b}, \mathbf{b})$ for $\mathbf{a} \neq \mathbf{b}$.

In order to compare μ_{ML} to distance measures such as d_{stim} and d_H , we can use the usual trick of taking the negative of the logarithm in order to convert similarity to distance:

$$\tilde{d}_{\text{ML}}(\mathbf{a}, \mathbf{b}) \stackrel{\text{def}}{=} -\ln \mu_{\text{ML}}(\mathbf{a}, \mathbf{b}).$$

It is clear, however, that \tilde{d}_{ML} is not a metric because $\tilde{d}_{\text{ML}}(\mathbf{a}, \mathbf{a}) \neq 0$ in general. We can fix this problem by first normalizing,

$$d_{\text{ML}}(\mathbf{a}, \mathbf{b}) \stackrel{\text{def}}{=} -\ln \left(\frac{\mu_{\text{ML}}(\mathbf{a}, \mathbf{b})}{\sqrt{\mu_{\text{ML}}(\mathbf{a}, \mathbf{a}) \mu_{\text{ML}}(\mathbf{b}, \mathbf{b})}} \right),$$

so that $d_{\text{ML}}(\mathbf{a}, \mathbf{a}) = 0$ for all words in $\{0, 1\}^n$. We call d_{ML} the *ML distance*. Unfortunately, d_{ML} still fails to be a metric on $\{0, 1\}^n$, as the triangle inequality is not generally satisfied (see appendix A.3), although it may be a metric when restricted to a particular code.

Despite not being a metric on $\{0, 1\}^n$, d_{ML} is useful as an indicator of how close the ML decoder comes to outputting the correct idealized codeword. By definition, ML decoding errors will have large ML similarity to the

¹²This definition does not explicitly depend on the channel parameters, although details of the channel are implicitly used in the computation of $P(\text{rec}|\text{sent})$.

correct codeword. In other words, even if $\hat{\mathbf{c}}_{\text{ML}} \neq \mathbf{c}$, the value of $d_{\text{ML}}(\hat{\mathbf{c}}_{\text{ML}}, \mathbf{c})$ will be relatively small. Unlike Hamming distance, d_{ML} naturally captures the notion that two codewords are close if they are likely to be confused after having been sent through the BAC channel and decoded with the ML decoder.¹³ In practice, however, d_{ML} is much more difficult to compute than Hamming distance. Fortunately, as we will see in the next section, there is a high correlation between d_{ML} and d_H , so that d_H may be used as a proxy for d_{ML} when using d_{ML} becomes computationally intractable.

5.4 Explanation of the Catch-Up Phenomenon. The ML distance d_{ML} is defined so that ML decoding errors have small ML distance to the correct codeword, regardless of the code. However, tolerating small errors makes sense only if errors are quantified by distances between stimuli, given by the induced stimulus space metric d_{stim} . The fact that RF codes catch up in error correction when an error tolerance with respect to d_{stim} is introduced suggests that on these codes, d_{stim} and d_{ML} correlate well, whereas on the comparison codes they do not. In other words, although the codewords in RF codes are not well separated inside $\{0, 1\}^n$, decoding errors tend to return codewords that represent very similar stimuli, and are hence largely tolerable.

To verify this intuition, we performed correlation analyses between d_{stim} and d_{ML} on 2D RF codes and corresponding random comparison codes. For each code, we computed d_{stim} and d_{ML} for all pairs of codewords and then computed the correlation coefficient between these two measures. Because finding d_{ML} among all pairs of codewords in a code with many neurons was computationally intractable, we performed this analysis on short codes having only 10 neurons, or length 10. Figure 7A shows a scatter plot of d_{stim} versus d_{ML} values for a single 2D RF code; the high correlation is easily seen by eye. In contrast, the same analysis for a corresponding shuffled code (see Figure 7B) and a random constant-weight code (see Figure 7C) revealed no significant correlation between d_{stim} and d_{ML} . Repeating this analysis for 10 matched sets of codes, each consisting of a 2D RF code, a corresponding shuffled code, and a corresponding random constant-weight code resulted in very similar results (see Figure 7D).

In order to test if the correlation between d_{stim} and d_{ML} might continue to hold for our longer codes with $n = 75$ neurons, we first investigated whether Hamming distance d_H could be used as a proxy for d_{ML} , as the latter can be computationally intractable. Indeed, on all of our length 10 codes, we found near-perfect correlation between d_H and d_{ML} (see Figure 7E). We then computed correlation coefficients using d_H instead of d_{ML} for the length

¹³On the binary symmetric channel (BSC), Hamming distance *does* measure the likelihood of two codewords being confused after ML decoding of errors introduced by the channel.

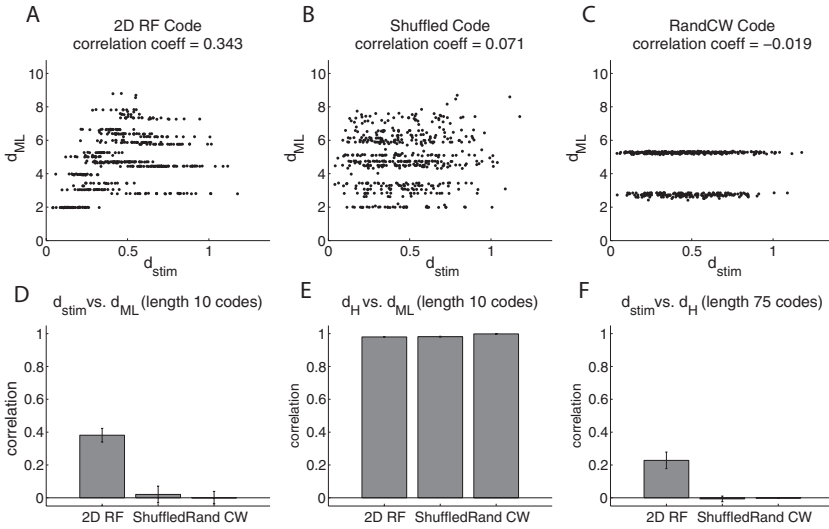


Figure 7: Correlations between stimulus space distance d_{stim} , Hamming distance d_H , and ML distance d_{ML} . (A) The scatter plot shows the correlation between d_{stim} and d_{ML} for a single 2D RF code of length 10 and sparsity $s = 0.188$. (A code of length 10 was chosen because d_{ML} is computationally intractable for longer codes.) Each point in the scatter plot corresponds to a pair of distinct codewords; random noise of size at most 0.001 in each dimension was added to each data point in order to separate points having exactly the same values of d_{stim} and d_{ML} . The values for d_{ML} were computed for channel parameters $p = 0.03$ and $q = 0.1$. (B) Same as in panel A, but for a shuffled code with matched parameters to the 2D RF code. (C) Same as in panel A, but for a random constant-weight code with matched parameters. (D) Average correlation coefficient between d_{stim} and d_{ML} for 10 2D RF codes of length 10 and mean sparsity $s = 0.191$ and 10 shuffled and random constant-weight codes of matched parameters. All d_{ML} values were computed for channel parameters $p = 0.03$ and $q = 0.1$. Error bars denote standard deviations. (E) Average correlation coefficient between d_H and d_{ML} for the same codes and channel condition used in panel D. The high correlation across codes suggests that d_H may be used as a proxy for d_{ML} in cases where d_{ML} is computationally intractable. (F) Average correlation coefficient between d_{stim} and d_H for 100 2D RF codes of length 75 and mean sparsity $s = 0.069$, and for 100 shuffled and random constant-weight codes of matched parameters (same codes as in Figure 4; note that this panel reproduces Figure 5D). Here we think of d_H as a proxy for d_{ML} . Because d_{ML} was not computed for this plot, calculations involving the larger codes were computationally feasible. Given the correlation patterns in panels A–F, it is likely that d_{stim} and d_{ML} are significantly correlated for large RF codes, but not for the shuffled or random constant-weight comparison codes.

75 2D RF codes and corresponding comparison codes that were analyzed in Figures 4, 5, and 6. As expected, there was a significant correlation between d_{stim} and d_H for RF codes but not for the random comparison codes (see Figure 7F). It is thus likely that d_{stim} and d_{ML} are well correlated for the large RF codes that displayed the catch-up phenomenon (see Figure 6) but not for the comparison codes.

6 Discussion

We have seen that although RF codes are highly redundant, they do not have particularly good error-correcting capability, performing far worse than random comparison codes of matching size, length, sparsity, and redundancy. This poor performance is perhaps not surprising when we consider the close proximity between RF codewords inside $\{0, 1\}^n$, a feature that limits the number of errors that can be corrected. On the other hand, RF code redundancy seems well-suited for preserving relationships between encoded stimuli, allowing these codes to reflect the geometry of the represented stimulus space. Interestingly, RF codes quickly catch up to the random comparison codes in error-correcting capability when a small tolerance to error is introduced. The reason for this catch-up is that errors in RF codes tend to result in nearby codewords that represent similar stimuli, a property that is not characteristic of the random comparison codes. Our analysis suggests that in the context of neural codes, there may be a natural trade-off between a code's efficiency and error-correcting capability and its ability to reflect relationships between stimuli. It would be interesting to investigate whether RF codes are somehow optimal in this regard, though this is beyond the scope of this letter. Likewise, it would be interesting to test the biological plausibility of the error tolerance values that are required for RF codes to catch up. For visual orientation discrimination in human psychophysics experiments, the perceptual errors range from about 4 degrees to 12 degrees (out of 180 degrees) (Mareschal & Shapley, 2004; Li, Thier, & Wehrhahm, 2000); this is roughly consistent with a 5% error tolerance, a level that resulted in complete catch-up for the 1D RF codes (see Figure 6A).

Throughout this work, we have assumed that neurons are independent. This assumption arose as a consequence of using the BAC as a channel model for noise, which operates on each neuron independently (see section 3.3). While somewhat controversial (Schneidman, Bialek, & Berry, 2003), some experimental evidence supports the independence assumption (Gawne & Richmond, 1993; Nirenberg, Carcieri, Jacobs, & Latham, 2001), in addition to a significant body of theoretical work suggesting that ignoring noise correlations does not have a significant impact on the decoding of neural population responses (Abbott & Dayan, 1999; Averbeck & Lee, 2004; Latham & Nirenberg, 2005). Nevertheless, it is quite possible

that the error-correcting capabilities of RF codes may increase (or decrease) if this assumption is relaxed (Averbeck, Latham, & Pouget, 2006). It would thus be interesting to explore a similar analysis for channel models that produce correlated noise, though this is beyond the scope of this letter.

We have also assumed a perfect understanding of the encoding map; however, it is possible that error-correcting capabilities vary significantly according to what aspect of the stimulus is being represented, similar to what has been found in information-theoretic analyses (Nemenman, Lewen, Bialek, & de Ruyter van Steveninck, 2008). Furthermore, in assessing the error-correcting properties of RF codes as compared to random comparison codes, we used a decoder that was optimal for all codes. If instead we used a biologically motivated decoder, such as those suggested in Deneve et al. (1999) and Beck et al. (2008), the performance of the random comparison codes may be significantly compromised, leading to a relative improvement in error correction for RF codes.

Mathematical coding theory has been very successful in devising codes that are optimal or nearly optimal for correcting noisy transmission errors in a variety of engineering applications (MacWilliams & Sloane, 1983; Wicker, 1994; Huffman & Pless, 2003). We believe this perspective will also become increasingly fruitful in neuroscience, as it provides novel and rigorous methods for analyzing neural codes in cases where the encoding map is relatively well understood. In particular, mathematical coding theory can help to clarify apparent paradoxes in neural coding, such as the prevalence of redundancy when it is assumed that neural circuits should maximize information. Finally, we believe the coding theory perspective will eventually provide the right framework for analyzing the trade-offs that are inherent in codes that are specialized for information transfer and processing in the brain.

Appendix A: ML and MAP Decoding

A.1 ML Decoding on the BAC. Here we derive a simple expression for the ML decoder on the binary asymmetric channel with false-positive probability p and false-negative probability q , as in Figure 2B. Recall that the ML decoder is given by

$$\hat{\mathbf{c}}_{\text{ML}} = \arg \max_{\mathbf{c} \in \mathcal{C}} P(\text{rec} = \mathbf{r} \mid \text{sent} = \mathbf{c}),$$

where $\mathbf{r} \in \{0, 1\}^n$ is the received word, or “actual response,” of the population of n neurons, and \mathcal{C} is the neural code. Because the channel is assumed to act on each neuron independently, $P(\text{rec} = \mathbf{r} \mid \text{sent} = \mathbf{c})$ will depend on

only the following quantities:

$$t_{00}(\mathbf{c}, \mathbf{r}) = \# \text{ of 0s that match between } \mathbf{c} \text{ and } \mathbf{r}.$$

$$t_{11}(\mathbf{c}, \mathbf{r}) = \# \text{ of 1s that match between } \mathbf{c} \text{ and } \mathbf{r}.$$

$$t_{01}(\mathbf{c}, \mathbf{r}) = \# \text{ of 0s in } \mathbf{c} \text{ that correspond to 1s in } \mathbf{r}.$$

$$t_{10}(\mathbf{c}, \mathbf{r}) = \# \text{ of 1s in } \mathbf{c} \text{ that correspond to 0s in } \mathbf{r}.$$

With this, it is straightforward to compute

$$P(\text{rec} = \mathbf{r} \mid \text{sent} = \mathbf{c}) = (1 - p)^{t_{00}(\mathbf{c}, \mathbf{r})} p^{t_{01}(\mathbf{c}, \mathbf{r})} (1 - q)^{t_{11}(\mathbf{c}, \mathbf{r})} q^{t_{10}(\mathbf{c}, \mathbf{r})}.$$

Using the obvious identities,

$$t_{01}(\mathbf{c}, \mathbf{r}) + t_{11}(\mathbf{c}, \mathbf{r}) = w_H(\mathbf{r})$$

$$t_{10}(\mathbf{c}, \mathbf{r}) + t_{00}(\mathbf{c}, \mathbf{r}) = n - w_H(\mathbf{r}),$$

we find

$$\begin{aligned} P(\text{rec} = \mathbf{r} \mid \text{sent} = \mathbf{c}) &= (1 - p)^{t_{00}(\mathbf{c}, \mathbf{r})} p^{w_H(\mathbf{r}) - t_{11}(\mathbf{c}, \mathbf{r})} (1 - q)^{t_{11}(\mathbf{c}, \mathbf{r})} q^{n - w_H(\mathbf{r}) - t_{00}(\mathbf{c}, \mathbf{r})}. \end{aligned}$$

When we do the maximization over $\mathbf{c} \in \mathcal{C}$, we can ignore terms that are independent of \mathbf{c} , and we obtain

$$\begin{aligned} \hat{\mathbf{c}}_{\text{ML}} &= \arg \max_{\mathbf{c} \in \mathcal{C}} \left[\left(\frac{1 - p}{q} \right)^{t_{00}(\mathbf{c}, \mathbf{r})} \left(\frac{1 - q}{p} \right)^{t_{11}(\mathbf{c}, \mathbf{r})} \right] \\ &= \arg \max_{\mathbf{c} \in \mathcal{C}} \left[t_{00}(\mathbf{c}, \mathbf{r}) \ln \left(\frac{1 - p}{q} \right) + t_{11}(\mathbf{c}, \mathbf{r}) \ln \left(\frac{1 - q}{p} \right) \right] \quad (*) \end{aligned}$$

If we further observe that

$$t_{11}(\mathbf{c}, \mathbf{r}) = \mathbf{c} \cdot \mathbf{r},$$

$$t_{00}(\mathbf{c}, \mathbf{r}) = (\mathbf{1} - \mathbf{c}) \cdot (\mathbf{1} - \mathbf{r}) = n - w_H(\mathbf{c}) - w_H(\mathbf{r}) + \mathbf{c} \cdot \mathbf{r},$$

where $\mathbf{1} \in \{0, 1\}^n$ is the all-ones word, and again ignore terms that are independent of \mathbf{c} , we obtain

$$\hat{\mathbf{c}}_{\text{ML}} = \arg \max_{\mathbf{c} \in \mathcal{C}} \left[\mathbf{c} \cdot \mathbf{r} \ln \left(\frac{(1 - p)(1 - q)}{pq} \right) - w_H(\mathbf{c}) \ln \left(\frac{1 - p}{q} \right) \right]. \quad (\text{A.1})$$

Since we assume $p, q < 1/2$, the decoder maximizes the number $\mathbf{c} \cdot \mathbf{r}$ of matching 1s between the sent and received words, subject to a penalty term that is proportional to the weight (i.e., the number of active neurons) of the sent word.

Note that for $p = q < 1/2$, as on the BSC, equation (*) becomes

$$\begin{aligned}\hat{\mathbf{c}}_{\text{ML}} &= \arg \max_{\mathbf{c} \in \mathcal{C}} [t_{00}(\mathbf{c}, \mathbf{r}) + t_{11}(\mathbf{c}, \mathbf{r})] \\ &= \arg \max_{\mathbf{c} \in \mathcal{C}} [n - d_H(\mathbf{c}, \mathbf{r})] = \arg \min_{\mathbf{c} \in \mathcal{C}} [d_H(\mathbf{c}, \mathbf{r})],\end{aligned}$$

where

$$d_H(\mathbf{c}, \mathbf{r}) = t_{01}(\mathbf{c}, \mathbf{r}) + t_{10}(\mathbf{c}, \mathbf{r}) = |\{i \in [n] \mid c_i \neq r_i\}|,$$

is the Hamming distance between two words in $\{0, 1\}^n$. In other words, ML decoding is equivalent to nearest neighbor decoding, with respect to Hamming distance, on the BSC.

A.2 Comparison of ML and MAP Decoding Using Bayes' Rule. Given two events A and B such that the probability of B is nonzero, Bayes' rule states,

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}.$$

We can use this theorem to relate the ML and MAP decoders:

$$\begin{aligned}\hat{\mathbf{c}}_{\text{MAP}} &= \arg \max_{\mathbf{c} \in \mathcal{C}} P(\text{sent} = \mathbf{c} \mid \text{rec} = \mathbf{r}) \\ &= \arg \max_{\mathbf{c} \in \mathcal{C}} \frac{P(\text{rec} = \mathbf{r} \mid \text{sent} = \mathbf{c})P(\text{sent} = \mathbf{c})}{P(\text{rec} = \mathbf{r})} \\ &= \arg \max_{\mathbf{c} \in \mathcal{C}} P(\text{rec} = \mathbf{r} \mid \text{sent} = \mathbf{c})P(\text{sent} = \mathbf{c}).\end{aligned}$$

In the case that all words are sent with equal probability, $P(\text{sent} = \mathbf{c})$ is constant over all codewords $\mathbf{c} \in \mathcal{C}$, we have

$$\hat{\mathbf{c}}_{\text{MAP}} = \arg \max_{\mathbf{c} \in \mathcal{C}} P(\text{rec} = \mathbf{r} \mid \text{sent} = \mathbf{c}) = \hat{\mathbf{c}}_{\text{ML}}.$$

Thus, the two decoders coincide when all stimuli are equally likely. In the case where some codewords are more likely to be transmitted than others, however, MAP and ML decoding need not coincide.

Suppose that the probability of a codeword being sent can be approximated by assuming individual 0s and 1s are transmitted with independent probabilities consistent with the sparsity s of the code:

$$P(\text{sent} = 1) \approx s,$$

$$P(\text{sent} = 0) \approx 1 - s.$$

Under these assumptions, we approximate

$$P(\text{sent} = \mathbf{c}) \approx s^{w_H(\mathbf{c})} (1 - s)^{n - w_H(\mathbf{c})}.$$

Using Bayes' rule and equation A.1, this gives an approximation for the MAP decoder as

$$\begin{aligned} \hat{\mathbf{c}}_{\text{MAP}} &\approx \arg \max_{\mathbf{c} \in \mathcal{C}} [P(\text{rec} = \mathbf{r} \mid \text{sent} = \mathbf{c}) s^{w_H(\mathbf{c})} (1 - s)^{n - w_H(\mathbf{c})}] \\ &= \arg \max_{\mathbf{c} \in \mathcal{C}} [\ln P(\text{rec} = \mathbf{r} \mid \text{sent} = \mathbf{c}) + w_H(\mathbf{c}) \ln s \\ &\quad + (n - w_H(\mathbf{c})) \ln(1 - s)] \\ &= \arg \max_{\mathbf{c} \in \mathcal{C}} \left[(\mathbf{c} \cdot \mathbf{r}) \ln \left(\frac{(1 - p)(1 - q)}{pq} \right) - w_H(\mathbf{c}) \ln \left(\frac{1 - p}{q} \right) \right. \\ &\quad \left. - w_H(\mathbf{c}) \ln \left(\frac{1 - s}{s} \right) \right] \\ &= \arg \max_{\mathbf{c} \in \mathcal{C}} \left[(\mathbf{c} \cdot \mathbf{r}) \ln \left(\frac{(1 - p)(1 - q)}{pq} \right) - w_H(\mathbf{c}) \ln \left(\frac{(1 - p)(1 - s)}{qs} \right) \right]. \end{aligned}$$

Comparing this approximation to equation A.1, we see that the difference between MAP and ML for sparse codes ($s < 1/2$) is that the approximate MAP decoder has a larger penalty term associated with the weight $w_H(\mathbf{c})$. This means that the approximate MAP decoder will sometimes return lower-weight codewords than the ML decoder. Unlike MAP, the ML decoder is completely indifferent to the code sparsity parameter s .

In our simulations with 2D RF codes, we have found that the above MAP approximation outperforms ML decoding when codewords in the distribution of transmitted words are weighted in a manner dictated by the sparsity of the code (see Figure 3). What if the codeword distribution is instead weighted by the sizes of the stimulus space regions corresponding to each codeword? In this case, Figure 8 shows that ML decoding outperforms the MAP approximation, further justifying our use of ML decoding in our analysis of the error-correcting properties of RF codes.

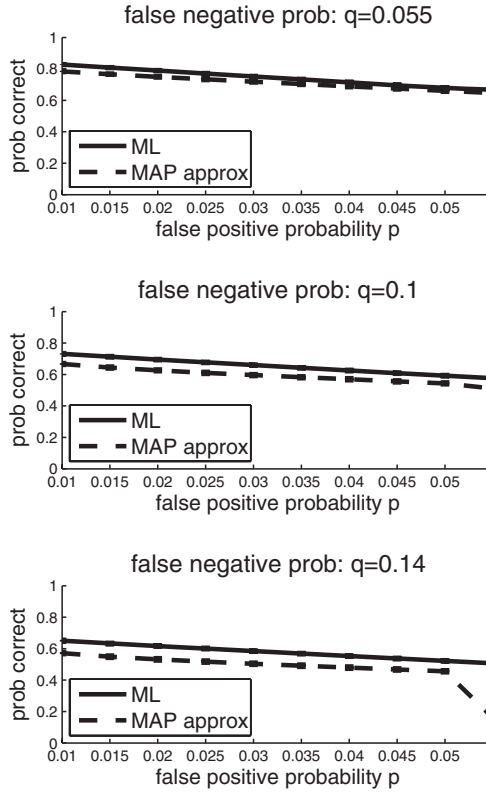


Figure 8: ML decoding outperforms approximate MAP decoding for a distribution of codewords weighted by the size of the codeword region. With the false-negative probability q fixed at $q = 0.055$ (top), $q = 0.1$ (middle) and $q = 0.14$ (bottom), the false-positive probability p was varied in increments of 0.005 from 0.01 to 0.055 to produce different channel conditions for the BAC. On each channel, the performance of 100 2D RF codes of length 75 and mean sparsity $s = 0.596$ was assessed using the standard ML decoder and our approximation to the MAP decoder. For each BAC condition and each code, 10,000 codewords were selected according to a weighted probability distribution, where the probability of sending codeword c was proportional to the area of $\varphi^{-1}(c)$, which was approximated by the fraction of points from the 300×300 fine grid that fell within the region $\varphi^{-1}(c)$ (if no grid point lay inside a region, we counted it as 1 grid point to ensure the probability was nonzero). Each codeword was transmitted across the BAC and decoded using both the ML decoder and the approximate MAP decoder. The fraction of correctly decoded words was then averaged across the 100 codes, with error bars denoting standard deviations. ML decoding consistently outperformed approximate MAP decoding for each channel condition, even though the opposite was true when codewords were weighted according to the sparsity of the code (see Figure 3A).

A.3 Failure of the Triangle Inequality for d_{ML} . Recall that the ML distance d_{ML} is defined by

$$d_{\text{ML}}(\mathbf{a}, \mathbf{b}) \stackrel{\text{def}}{=} -\ln\left(\frac{\mu_{\text{ML}}(\mathbf{a}, \mathbf{b})}{\sqrt{\mu_{\text{ML}}(\mathbf{a}, \mathbf{a})\mu_{\text{ML}}(\mathbf{b}, \mathbf{b})}}\right),$$

where the ML similarity μ_{ML} is the probability that \mathbf{a} and \mathbf{b} will be confused in the process of transmission across the BAC and then running the channel output through an ML decoder.

It is clear that d_{ML} is a pseudo-semimetric on $\{0, 1\}^n$; that is, for all $\mathbf{a}, \mathbf{b} \in \{0, 1\}^n$ we have $d_{\text{ML}}(\mathbf{a}, \mathbf{b}) \geq 0$, $d_{\text{ML}}(\mathbf{a}, \mathbf{a}) = 0$, and $d_{\text{ML}}(\mathbf{a}, \mathbf{b}) = d_{\text{ML}}(\mathbf{b}, \mathbf{a})$. However, d_{ML} is not a metric or even a pseudo-metric on $\{0, 1\}^n$ because it fails to satisfy the triangle inequality. As an example, consider the code $\mathcal{C} = \{(1, 1, 0), (1, 0, 1), (0, 0, 1)\}$, and take $\mathbf{x} = (0, 0, 1)$, $\mathbf{y} = (0, 0, 0)$, and $\mathbf{z} = (0, 1, 0)$. For channel conditions $p = 0.05$ and $q = 0.07$, we obtain

$$d_{\text{ML}}(\mathbf{x}, \mathbf{y}) + d_{\text{ML}}(\mathbf{y}, \mathbf{z}) = .005 + 3.217 = 3.222 < 4.072 = d_{\text{ML}}(\mathbf{x}, \mathbf{z}).$$

It is interesting to note, however, that both the triangle inequality $d_{\text{ML}}(\mathbf{a}, \mathbf{b}) + d_{\text{ML}}(\mathbf{b}, \mathbf{c}) \geq d_{\text{ML}}(\mathbf{a}, \mathbf{c})$ and the condition that $d(\mathbf{a}, \mathbf{b}) = 0$ only if $\mathbf{a} = \mathbf{b}$ hold in all examples we have tried when \mathbf{a} , \mathbf{b} , and \mathbf{c} are chosen to be codewords in some code \mathcal{C} . In other words, it is unknown to us whether d_{ML} is a metric when restricted to a code $\mathcal{C} \subset \{0, 1\}^n$, even though it is not a metric on the entire ambient space $\{0, 1\}^n$.

Appendix B: Details of the Simulations

B.1 Generation of 1D RF Codes. To generate the 1D RF codes used in our simulations, we took the length of the stimulus space to be 1 and identified the points 0 and 1 since the stimuli represent angles in $[0, \pi)$. Each receptive field (tuning curve) was chosen to be an arc of the stimulus space. We chose our receptive fields to have a constant radius of 0.08, which corresponds to a radius of 14.4 degrees in the orientation selectivity model. This parameter matches that in Somers, Nelson, and Sur (1995), where tuning curves in the visual cortex were set to have half-width-half-amplitudes of 14.9 degrees, based on experimental data from Watkins and Berkley (1974) and Orban (1984). Each receptive field was specified by its center point. We used 75 receptive fields to cover the stimulus space, and so our codewords had length 75. The centers of the receptive fields were selected uniformly at random from the stimulus space, with the following modification: while the stimulus space remained uncovered, the centers were placed randomly in the uncovered region. This modification allowed us to guarantee that the stimulus space would be covered by the receptive

fields; we used a fine grid of 300 uniformly spaced test points to find uncovered regions in the stimulus space.

By examining all pairwise intersections of receptive fields, we found all the regions cut out by the receptive fields, and each such region defined a codeword (see Figure 1A). Note that each codeword corresponds to a convex region of the stimulus space. The *center of mass* of a codeword is the center point of the interval to which the codeword corresponds.

B.2 Generation of 2D RF Codes. To generate the 2D RF codes used in our simulations, we took the stimulus space to be a 1×1 square box environment. Each receptive field was the intersection of the stimulus space with a disk whose center lay within the stimulus space. All disks were chosen to have the same radius; this is consistent with findings that place fields in the dorsal hippocampus are generally circular and of similar sizes (Jung, Weiner, & McNaughton, 1994; Maurer, Vanhoads, Sutherland, Lipa, & McNaughton, 2005). We chose the radius of our receptive fields to be 0.15, that is, 15% of the width of the stimulus space, to produce codes having a reasonable sparsity of ~ 0.07 . As with the 1D RF codes, we generated 75 receptive fields to cover the space, with each receptive field identified by its center point. In our simulations, the center points of the receptive fields were dropped uniformly at random in the stimulus space, with the same modification as for the 1D RF codes: while the space remained uncovered, the centers of the disks were placed uniformly at random in regions of the space that had yet to be covered. We used a fine grid of 300×300 uniformly spaced test points to find uncovered regions in the stimulus space.

Again, by examining all intersections of receptive fields, we found all regions cut out by the receptive fields, and each region defined a codeword (see Figure 1B). Unlike with the 1D RF codes, however, the codeword regions in the 2D RF codes were not guaranteed to be convex or even connected subsets of the stimulus space, although the typical region was at least connected. For the purpose of defining a stimulus space distance on these codes, we defined the center of mass of a codeword to be an appropriate approximation of the center of mass of the region corresponding to the codeword, regardless of whether that center lay within the region. When the codeword region was large enough to contain points from the 300×300 fine grid, we took the center of mass of the codeword to be the center of mass of the grid points contained in the codeword region. A small number of codewords had regions that were narrow crescents or other small shapes that avoided all grid points; in these cases, the center of mass of the codeword was taken to be the center of mass of the receptive field boundary intersection points that defined the region.

For Figure 7, we generated 10 new 2D RF codes of length 10. For these smaller codes, the radius was chosen to be 0.25 to ensure reasonable coverage of the space. All other parameters were as described above.

B.3 Details of Error Correction Simulations. As a result of the chosen receptive field radii, the mean sparsity of the 1D RF codes was $s = 0.165$, while the mean sparsity of the 2D RF codes was $s = 0.069$. To test how effective each of these types of codes was compared to the random codes with matched parameters, we chose to make the error probabilities as high as possible while still abiding by our BAC channel constraints and maintaining a reasonable value for the expected number of errors in each transmission. Thus, we set $q = 0.20$ for the 1D RF codes and $q = 0.10$ for the 2D RF codes.¹⁴ To test the performance of these codes over varying degrees of channel asymmetry, the value of p was chosen to range from 0.05 to 0.15 in increments of 0.01 for the 1D RF codes, while p ranged from 0.01 to 0.06 in increments of 0.005 for the 2D RF codes.

Acknowledgments

C.C. was supported by NSF DMS 0920845 and an Alfred P. Sloan Research Fellowship. V.I. was supported by NSF DMS 0967377 and NSF DMS 1122519. K.M. was supported by NSF DMS 0903517 and NSF DMS 0838463. Z.R. was supported by Department of Education GAANN grant P200A060126. J.L.W. was supported by NSF DMS 0903517.

References

- Abbott, L., & Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Computation*, *11*, 91–101.
- Adesberger-Mangan, D., & Levy, W. (1992). Information maintenance and statistical dependence reduction in simple neural networks. *Biological Cybernetics*, *67*, 469–477.
- Attkin, J. (1992). Could information theory provide an ecological theory of sensory processing? *Network*, *3*, 213–251.
- Attneave, F. (1954). Some informational aspects of visual perception. *Psychology Reviews*, *61*, 183–193.
- Averbeck, B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, *7*, 358–366.
- Averbeck, B., & Lee, D. (2004). Coding and transmission of information by neural ensembles. *Trends in Neurosciences*, *27*(4), 225–230.
- Barlow, H. (1961). Possible principles underlying the transformation of sensory messages. In W. Rosenblith (Ed.), *Sensory communication* (pp. 217–234). Cambridge, MA: MIT Press.
- Barlow, H. (2001). Redundancy reduction revisited. *Network: Computational Neural Systems*, *12*, 241–253.

¹⁴In addition to the simulations shown here with the above parameters, we also tested the code performance over a range of both larger and smaller values of q and obtained similar results.

- Beck, J., Ma, W. J., Kiani, R., Hanks, T., Churchland, A., Roitman, J., et al. (2008). Probabilistic population codes for Bayesian decision making. *Neuron*, *60*, 1142–1152.
- Ben-Yishai, R., Bar-Or, R. L., & Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci. USA*, *92*(9), 3844–3848.
- Bialek, W., Rieke, F., de Ruyter van Stevenick, R., & Warland, D. (1991). Reading a neural code. *Science*, *252*(5014), 1854–1857.
- Borst, A., & Theunissen, F. E. (1999). Information theory and neural coding. *Nature Neuroscience*, *2*(11), 947–957.
- Cover, T., & Thomas, J. (2006). *Elements of information theory*. New York: Wiley.
- Curto, C., & Itskov, V. (2008). Cell groups reveal structure of stimulus space. *PLoS Computational Biology*, *4*(10), e1000205.
- deCharms, R. C., & Zador, A. (2000). Neural representation and the cortical code. *Annual Reviews in Neuroscience*, *23*, 613–647.
- Deneve, S., Latham, P. E., & Pouget, A. (1999). Reading population codes: A neural implementation of ideal observers. *Nature Neuroscience*, *2*(8), 740–745.
- Gawne, T., & Richmond, B. (1993). How independent are the messages carried by adjacent inferior temporal cortical neurons? *Journal of Neuroscience*, *13*(7), 2758–2771.
- Hopfield, J. (2008). Searching for memories, Sudoku, implicit check bits, and the iterative use of not-always-correct rapid neural computation. *Neural Computation*, *20*(5), 1119–1164.
- Huffman, W., & Pless, V. (2003). *Fundamentals of error-correcting codes*. Cambridge, MA: Cambridge University Press.
- Huttenlocher, J., Hedges, L., Lourenco, S., Crawford, L., & Corrigan, B. (2007). Estimating stimuli from contrasting categories: Truncation due to boundaries. *Journal of Experimental Psychology*, *136*(3), 502–519.
- Jacobs, A., Fridman, G., Douglas, R., Alam, N., Latham, P., Prusky, G., et al. (2009). Ruling out and ruling in neural codes. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(14), 5936–5941.
- Jung, M., Weiner, S., & McNaughton, B. (1994). Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. *Journal of Neuroscience*, *14*, 7347–7356.
- Latham, P. E., & Nirenberg, S. (2005). Synergy, redundancy, and independence in population codes, revisited. *Journal of Neuroscience*, *25*(21), 5195–5206.
- Levy, W., & Baxter, R. (1996). Energy efficient neural codes. *Neural Computation*, *8*, 531–543.
- Li, W., Thier, P., & Wehrhahm, C. (2000). Contextual influence on orientation discrimination of humans and responses of neurons in V1 of alert monkeys. *Journal of Neurophysiology*, *83*, 941–954.
- London, M., Roth, A., Beeren, L., Häusser, M., & Latham, P. (2010, July). Sensitivity to perturbations in vivo implies high noise and suggests rate coding in the cortex. *Nature*, *466*, 123–127.
- Luczak, A., Barthó, P., & Harris, K. (2009). Spontaneous events outline the realm of possible sensory responses in the auditory cortex. *Neuron*, *62*(3), 413–425.
- Ma, W. J., Beck, J., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, *9*, 1432–1438.

- MacWilliams, F. J., & Sloane, N. J. A. (1983). *The theory of error-correcting codes*. Amsterdam: North-Holland.
- Mareschal, I., & Shapley, R. M. (2004). Effects of contrast and size on orientation discrimination. *Vision Research*, *44*, 57–67.
- Maurer, A., Vanrhoads, S., Sutherland, G., Lipa, P., & McNaughton, B. (2005). Self-motion and the origin of differential spatial scaling along the septo-temporal axis of the hippocampus. *Hippocampus*, *15*, 841–852.
- McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M. B. (2006). Path integration and the neural basis of the “cognitive map.” *Nat. Rev. Neurosci.*, *7*(8), 663–678.
- Nemenman, I., Lewen, G. D., Bialek, W., & de Ruyter van Steveninck, R. (2008). Neural coding of natural stimuli: Information at sub-millisecond resolution. *PLoS Comp. Bio.*, *4*(3).
- Nirenberg, S., Carcieri, S., Jacobs, A., & Latham, P. E. (2001). Retinal ganglion cells act largely as independent encoders. *Nature*, *411*, 698–701.
- O’Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, *34*(1), 171–175.
- Orban, G. (1984). *Neuronal operations in the visual cortex*. Berlin: Springer.
- Osborne, L., Palmer, S., Lisberger, S., & Bialek, W. (2008). The neural basis for combinatorial coding in a cortical population response. *Journal of Neuroscience*, *50*(28), 13522–13531.
- Prinzmetal, W., Shimura, A., & Mikolinski, M. (2001). The Ponzo illusion and the perception of orientation. *Perception and Psychophysics*, *63*(1), 99–114.
- Puchalla, J., Schneidman, E., Harris, R., & Berry, M. (2005). Redundancy in the population code of the retina. *Neuron*, *46*, 493–504.
- Quiroga, R. Q., & Panzeri, S. (2009). Extracting information from neuronal populations: Information theory and decoding approaches. *Nature Reviews Neuroscience*, *10*, 173–185.
- Rieke, F., Warland, D., de Ruyter van Steveninck, R., & Bialek, W. (1999). *Spikes: Exploring the neural code*. Cambridge, MA: MIT Press.
- Schneidman, E., Berry, M., Segev, R., & Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, *440*, 1007–1012.
- Schneidman, E., Bialek, W., & Berry, M. (2003). Synergy, redundancy, and independence in population codes. *Journal of Neuroscience*, *23*(37), 11539–11553.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, *27*, 379–423, 623–656.
- Somers, D. C., Nelson, S. B., & Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortical simple cells. *Journal of Neuroscience*, *15*, 5448–5465.
- Sreenivasan, S., & Fiete, I. (2011). Grid cells generate an analog error-correcting code for singularly precise neural computation. *Nature Neuroscience*, *14*(10), 1330–1337.
- Tkačik, G., Granot-Atedgi, E., Segev, R., & Schneidman, E. (2012). *Retinal metric: A stimulus distance measure derived from population neural responses*. <http://arxiv.org/abs/1205.6598v1>

- van der Heijden, A., van der Geest, J., de Leeuw, F., Krikke, K., & Müsseler, J. (1999). Sources of position-perception errors for small isolated targets. *Psychological Research*, *62*(1), 20–35.
- Watkins, D., & Berkley, M. (1974). The orientation selectivity of single neurons in cat striate cortex. *Experimental Brain Research*, *19*, 433–446.
- Wicker, S. B. (1994). *Reed-Solomon codes and their applications*. Piscataway, NJ: IEEE Press.

Received September 28, 2012; accepted January 8, 2013.