

# Ecological expected utility and the mythical neural code

Jerome Feldman

Received: 17 April 2009 / Revised: 12 August 2009 / Accepted: 18 August 2009 / Published online: 4 September 2009  
© The Author(s) 2009. This article is published with open access at Springerlink.com

**Abstract** Neural spikes are an evolutionarily ancient innovation that remains nature's unique mechanism for rapid, long distance information transfer. It is now known that neural spikes sub serve a wide variety of functions and essentially all of the basic questions about the communication role of spikes have been answered. Current efforts focus on the neural communication of probabilities and utility values involved in decision making. Significant progress is being made, but many framing issues remain. One basic problem is that the metaphor of a neural code suggests a communication network rather than a recurrent computational system like the real brain. We propose studying the various manifestations of neural spike signaling as adaptations that optimize a utility function called ecological expected utility.

**Keywords** Neural coding · Ecological utility · Neuroeconomics · Spikes · Grandmother cell · Population code · Binding problem

## Introduction and background

A continuing theme in computational neuroscience has been the search for “the neural code”. In this paper, I will suggest that this is not a well formed question and has given rise to a fair amount of needless confusion. Neural spikes are an evolutionary ancient innovation that remains nature's unique mechanism for rapid long distance information

transfer (Meech and Mackie 2007). Other communication mechanisms are either much slower (e.g., hormones) or extremely local (e.g., gap junctions). It is now clear that neural spikes sub serve a wide variety of functions. Rather than trying to restate well established facts about neural spikes, this paper will develop a new and broader view. For background, we will rely on a few standard books and three fairly recent surveys by de Charms and Zador (2000), by Kreiman (2004), and by Gollisch (2009).

We will continue the tradition of examining neural signaling from the information processing perspective (Feldman 2006, Chap. 2), backgrounding the underlying biochemistry. Most research is focused, as it should be, on specific systems, but there are also important regularities in neural representation and communication. Current research is extending these general studies to decision making, but is unfortunately also falling prey to the myth of a unique neural code.

One major barrier to understanding “the neural code” is that the term itself can be misleading. First of all, using “the” presupposes that there is just one mode of neural signaling, which is known to be false. In addition, one standard meaning of a *code* is a fixed representation of information that is independent of the sender, receiver, and mechanisms of transmission (Kreiman 2004). The traditional example was Morse code, but perhaps the best known current example is the ASCII code used in computing. In ASCII, the lowercase “a” is always 1100001 and uppercase “A” is always 1000001, etc. Of course, genetic DNA is a code in this sense, although much more complex than ASCII. Neural spike signals are used in several ways in living systems, but this kind of context-free code is not, and could not be, one of them, as will be shown in the section “[Spikes in single neuron communication](#)”. We will use the phrase *neural signaling* to refer to the rapid, long

---

J. Feldman (✉)  
UC Berkeley and International Computer Science Institute,  
Berkeley, CA, USA  
e-mail: feldman@ICSI.Berkeley.EDU

distance communication mediated by spikes that is the subject of this article.

Another, quite distinct, use of the term “neural code” relates neural signaling to Shannon information theory, often called coding theory. The standard book “Spikes: Exploring the Neural Code” by (Reike et al. 1997) makes no mention at all of ASCII like codes. Codes in information theory have no symbolic meaning and this is a much better model for neural signals. The Rieke et al. book remains the most thorough treatment of neural signaling from the coding theory perspective, containing theoretical and experimental treatments of rate coding, decoding, quantity of information, and reliability. For the more recent, utility-based studies, we will rely mainly on the collection of articles in (Glimcher 2009).

Spikes are an evolutionarily ancient mechanism that is largely preserved. Quoting from John Allman (2000), who knows vastly more than I do about the brain:

Action potentials and voltage-gated sodium channels are present in jellyfish, which are the simplest organisms to possess nervous systems. The communication among neurons via action potentials and its underlying mechanism, the voltage-gated sodium channel, were essential for the development of nervous systems and without nervous systems complex animals could not exist.

Much of the mechanism behind neural spikes goes back even earlier in time (Meech and Mackie 2007; Katz 2007). Unsurprisingly, the earliest function of spiking neurons is to provide a signal for coordinated muscle action as in the swimming of the jellyfish. This kind of one-shot direct action remains one of the principal functions of neural spikes. Nature is quite conservative—once a winning design evolves it is reused and adapted. There is no reason to expect the range of variation in function for neural spikes to be less diverse than that of forelimbs, which have become arms, legs, wings, or fins. Spiking neurons are evolutionarily much older than forelimbs.

Because of the underlying chemistry, all neural spikes are of the same size and duration. The basic method of neural information transfer is, and needs to be, by *labeled lines*. Most of the information conveyed by a sensory neural spike train comes from the origin of the signal. Every patch of your skin contains a variety of specialized sensing neurons, each conveying a specific message. Similarly, the result of motor control signaling is largely determined by which muscle fibers are targeted. The other available degree of freedom is *timing*; there is a wide range of variation in the axonal conduction time of neural spikes. Nature has evolved a rich variety of mechanisms for exploiting absolute and relative time. Of course, all of the variants of neural signaling were selected for their

evolutionary fitness and this becomes important later in the article.

An additional definitional problem was that the idea of a code suggests some deterministic representation. As recently as 2000, a survey of research on the Neural Code in the Annual Review of Neuroscience (de Charms and Zador 2000) had essentially no mention of probability and utility. The zeitgeist has changed radically and Kreiman’s (2004) review and subsequent articles are largely framed in probabilistic terms. The volume by Glimcher (2009) remains the best introduction to these developments, often called “neuroeconomics”. We will look at this in some detail in later sections.

From our information processing perspective, the crucial issue is *effective signals*. We are interested in when a neuron or neural system evokes an *action* or makes a *decision*. From our purposes these are essentially the same and will be referred to as action/decision. Obviously enough, an individual neural spike is an action and can also be viewed as a (metaphorical) decision by the neuron to fire. For larger neural circuits, we still want to focus on information (coded as spikes) that (eventually) leads to a decision to do one thing rather than another. If there is no choice, there is no effective information. This becomes important when we try to analyze what function is supported by some spiking pattern.

Most research is focused, as it should be, on the specific systems, but there are some useful insights about neural representation and communication in general. The goal here is to consider the mechanisms of rapid, long distance communication in nervous systems. If we try to include the *content* of all the messages, we would need to explain everything about the brain. If we focus on the *form* of messages conveyed by neural spikes, a relatively clear and coherent picture emerges. We will focus on the information conveyed directly by spikes and not review the significant modulatory effects of chemical signals delivered by neural activity. Other important related topics that will not be covered in detail include spike generation, development, and learning.

In addition to information processing, the other organizing principle for this study is *resource limitations*. The most obvious resource limitation for neural action/decision is *time*. Many actions need to be fast even if that means sacrificing some accuracy. Some neural systems evolved to meet remarkable relative timing constraints, much shorter than spike intervals. A second crucial resource is *energy*; neural firing is metabolically expensive (Atwell and Laughlin 2001; Lennie 2003) and brains evolved to conserve energy while meeting performance requirements. The three factors of accuracy, timing, and resources are the core of a utility function that constrains neural computation. For advanced social animals like ourselves, there are a number

of additional considerations including learning, the exploit/explore tradeoff (Cohen et al. 2007), communication, and social cognition. Neural signaling evolved to serve all these functions, but the range of basic mechanisms involved is rather restricted.

### Spikes in single neuron communication

Although all behavior involves neural *circuits*, it is useful to first consider the role of neural spikes in the communication from a single neuron to another or to an effector cell. As always, our discussion will elide the biophysical and chemical details and focus on the information processing perspective. As discussed above, one ancient and important use of spikes is when a single spike evokes an action/decision. Single spike activity is surprisingly important in complex brains, including ours. As (Reike et al. 1997, p. 17) point out, spiking is temporally sparse—often about one spike per neuron for a salient event.

In some cases, the temporally leading spikes can be shown to directly determine human behavior. This is called “spike wave theory” (Rousselet et al. 2007) or sometimes a “latency code”. In a path breaking series of experiments the Thorpe group (Kirchner and Thorpe 2006) has shown that complex visual decisions can be made in about the time that it takes for the first spikes from a visual input to reach the brain area involved. The standard task is to push a left button if a complex scene contains any picture of an animal and a right button if not. The decision is detectable as right or left motor cortex activity in 150 ms, which is very close to the cumulative signaling and transmission delays involved. The Rousselet paper suggests (p. 1255) that the system can make the binary decision without actual recognition—using top down priming to condition the network to choose between two competing collections of criteria features.

A related finding is computationally modeled by (Serre et al. 2007). In these studies, people are shown a very wide range of visual scenes at a rate of seven images per second. Subjects have brief experiences of recognition for each scene although there is no time for eye movements. The authors are able to model this behavior with a simple biologically motivated hierarchical feed-forward connectionist network, trained on half the sample images.

The spike wave story also suggests an important point about firing patterns in neural populations. At every step in a perceptual task there are many neurons firing, but most of this is the encoding of competing potential action/decisions—the code is basically *disjunctive* not conjunctive (Jazayeri and Movshon 2006). This will be discussed in the section “Population codes”.

We should also discuss why it is not feasible for one neuron to send an abstract symbol (as in Morse or ASCII code) to another as a spike pattern. There are several related considerations from neural computation. We know experimentally that the firing of sensory (e.g., visual) neurons is a function of several stimulus variables, often intensity, position, velocity, orientation, color, etc. It would take a rather long message to convey all this as an ASCII like code and the firing rates are much too slow for this, even ignoring the stochastic nature of neural spikes. Even if such a message were somehow encoded and sent to the next level, it would require an elaborate computation to decode it and combine it with the symbolic messages of neighboring cells and then build a new symbolic message for the subsequent levels. This is what we do with language, but nothing at all like this occurs at the individual neuron level.

The most common use of neural signaling is to indicate the *strength* of some event, where the strength can code a combination of the intensity of a multidimensional event and some function of probability. Since all spikes are the same size, the strength of the signal must be encoded in the number or frequency of spikes. There are, of course, synaptic weights but these encode the strength of *connection*, not the signal. The conventional story is that the strength of a neural spike signal is conveyed by spike discharge frequency and everyone agrees that this is often the case. But we also now know that there is more time dependence than just spike frequency involved.

There is another terminological problem involving the role of spike timing in neural communication. There are at least three different versions of what *synchrony* might mean for neural signaling. There is no question that the relative timing of spikes arriving at a receiving site (synapse, dendrite, cell, etc.) can have a profound effect on the response (Kara and Reid 2003) and therefore timing is certainly relevant. Obviously enough, spikes that are temporally distant do not sum. There are also much more delicate timing interactions involved in the echolocation systems of bats, owls, etc. and also in LTP (long term potentiation) and STDP (spike timing dependent plasticity), all of which are discussed briefly in the next section.

A second timing question is whether the inter-spike interval of a single neural spike train conveys more useful information than just the firing rate. Because the base firing rate is often low, a single spike can be a rare event and therefore convey more than one bit of information. But the main area of contention has been whether there is significant additional information conveyed by the detailed timing of a spike train, beyond its average frequency. Spike trains are known to be stochastic, but if we assume that the distribution is known, two or three spikes can provide a fairly good estimate of the underlying parameter and thus

the firing frequency. The (Reike et al. 1997) book discusses this controversy in some detail; the conclusion is that very little if any useful information is conveyed by the phase structure of a single pulse train. This is not important in any case, because essentially all neural computation involves inputs from multiple sources.

A third distinct timing issue involves hypothesized periodic synchronous firing patterns as an organizing principle for conceptual binding; this will be discussed at the end of the following section.

### Spikes in neural circuits

Neurons never work in isolation; all behavior is mediated by specific circuits, from the contraction of the hydra to human speech. In higher animals, there are several levels of redundancy in these circuits and this is deeply connected to spike-based signaling. As is well known, the firing of an individual neuron is inherently probabilistic and so reliable communication requires several parallel channels. In some cases there are quite delicate timing constraints between signals on coordinated channels. But we will first discuss circuits where the timing requirement is just that the spikes are sufficient close in space and time to have their effects combine chemically.

Although some local circuits (e.g., in the retina) use gap junctions, most are mediated by neural spikes. These local computations are almost never considered part of the “neural code” and this is another flaw in the standard formulation. The Dayan and Abbott book (2001) and the Gerstner and Kistler book (2002) have good treatments. The most basic circuits involve competition, cooperation or the combination of both. The paradigmatic example of competition is mutual inhibition; this has two major effects on theories of neural signaling. Mutual inhibition is what allows a neural population with competing activation patterns to come to a specific action/decision; there does need to be a homunculus decider. Competing circuits play an important role in neuroeconomics and will be discussed in the section “[Probability, utility, and fitness](#)”. Inhibition also allows for redundant back-up circuits to be in place in case of damage to the primary circuit for some function, this is known technically as “release from inhibition” (Snyder and Sinex 2002).

Cooperation between input signal streams also has many realizations. The most basic is the use of multiple pathways for greater reliability and dynamic range. Another important use of coordinated signaling is in hyper-acuity. As was mentioned earlier, the spike signal from sensory (e.g., visual) neurons encodes the strength of match to a broad, multi-dimensional, receptive field. Cooperating signals allow for the sensing of differences beyond the

discrimination of the input receptors. As a toy example, the coincident firing of a cell that detected values in the range [1–10] with one sensitive to values [8–16] signals the much tighter range [8–10]. This coarse coding circuit mechanism can be seen as a way of extending the representational power of neural firing. This well established neural mechanism of coding the strength of a multi-dimensional signal is another barrier to postulating a simple readable neural code.

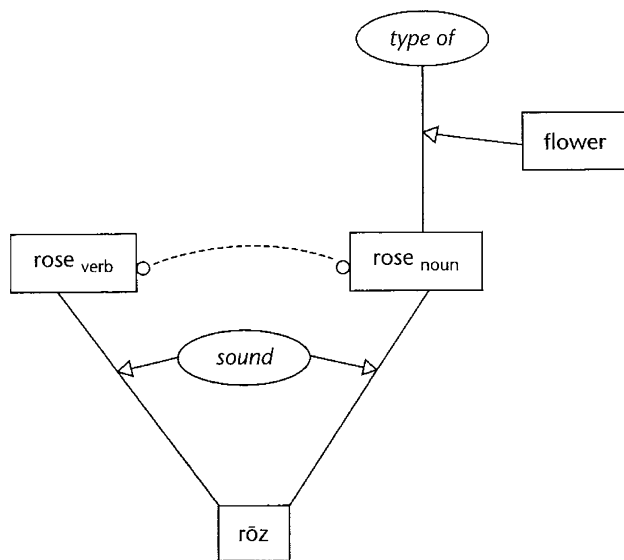
So far, we have focused on the immediate use of neural spikes for action/decision. However, there are also important indirect effects, studied under the names “spreading activation” and “priming”. Intuitively, spreading activation lies behind the fact that your thoughts will often shift among loosely related ideas without any conscious action. The brain is very richly connected and activation of one thought provides collateral activation to ones that are conceptually linked.

It is widely believed that a general *best-fit* process involving both competition and cooperation is the fundamental process underlying recognition in vision, language, etc. Given two competing analyses of an image (e.g., Necker Cube) or a sentence, we normally settle on one and do not notice the other. It is difficult to exactly prove this theory, but it is consistent with a great deal of evidence. Several of the ideas just described can be seen in the classical experiment of (Tanenhaus et al. 1979).

Subjects were asked to decide quickly whether letters flashed on a screen formed an English word. It was already well known that responses could be improved by “priming”, presenting a clue slightly before the test—either visually or auditorially. For example, hearing the word “rose” makes people faster to indicate that the text “flower” is an English word. In this experiment, subjects heard sentences with a misleading word sense like “They all rose” where the target word is not semantically related to the sound clue. The results depended entirely on the relative timing of the sound and image stimuli.

If the sound “rose” in the sentence above was timed to be only slightly (<200 ms) before the test image, it still had a priming effect on a word (flower) that was related to a unintended meaning of the sound that was never even noticed by the subject. This is generally agreed to indicate that there is pre-attentive parallel activation of all the words consistent with a heard sound. With a somewhat longer interval between the sound clue and the target image, the priming effect disappears and there is even a slight slowing from the neutral case. The hypothesized underlying neural circuitry is modeled in Fig. 1.

The priming link between the noun “rose” and “flower” is modeled as arising from the semantic relation that a rose is a type of flower. The triangular node models a local 2/3 circuit (Carpenter and Grossberg 1987) hypothesized to



**Fig. 1** Model circuit for cross-modal priming (from Feldman 2006, p. 90)

capture such semantic relations. The mutual inhibition between the two senses of rose is diagrammed as a dashed line with circular tips. In general, this kind of cooperative/competitive architecture underlies a wide range of theoretical and experimental work on neural representation. Again, the communication and settling of such networks is not usually considered as part of “the neural code”. More generally, the spiking behavior of receiving neurons is also a function of their internal state, which is yet another reason why there is not a fixed neural code.

We have discussed the circuit of Fig. 1 as local, but it actually involves brain areas for sound, image, and conceptual information. This is not unusual; almost all behavior uses multiple brain areas in complex interactions. For example, there are many more feedback connections from primary visual cortex to the lateral geniculate than the feed-forward links that are usually studied. In general, a major problem with the metaphor of the neural code is that it suggests a communication channel rather than a recurrent computational system.

### Time sensitive computations

Now we consider the time sensitive interactions. All neural communication is constrained by the time constants of the underlying chemistry, but there are some mechanisms that have tighter bounds and are generally called coincidence detection. For example, the circuits that process visual motion rely on the coincidence of a current input with a signal from cells that responded to a similar input earlier in time at a nearby location. This encodes evidence for specific motion and it is why we see the discrete frames of film

or TV as continuous motion. At the extreme end, the conduction time of spikes along axons is used with coincidence to support very fine distinctions in the arrival time of sound at the two ears, mostly notably in owls. Owls and bats make distinctions that correspond to timing differences at the ten microsecond level—much faster than neural switching times. Similar fine timing distinctions occur in dolphins, electric fish, etc. (Carr 1993); this is another example of the use of spike signaling that cannot be called a code.

There is now a great deal known about the chemical details on the timing sensitivity of post-synaptic events. Timing fit is crucial in a variety of ways, including LTP (Long Term Potentiation, Gerstner et al. 1997) and STDP (spike timing dependent plasticity; Bender et al. 2006; Izhikevich 2007), but development and learning are beyond the scope of this review.

### The variable binding problem

Another area of considerable research and controversy explores the possibility that separate coordinated phases of neural spiking play a central role in the “binding problem”. The basic and easier binding problem concerns how we can coherently see a bouncing red ball and a blue book given that these properties are computed in separate brain areas. The harder, *variable binding*, problem involves how we remember and draw inferences from complex relations (Barrett et al. 2008). For example, the sentence “John gave the book to Dick and the red thing to Jane” automatically leads to inferences about who has what (Shastri 2001). Computational models of these processes are easy in conventional programming, but no one has a convincing story of how the brain achieves this.

The most popular and thoroughly studied neural model of variable binding is based on the idea of synchronous neural firing patterns. Suppose that (somehow) the pattern of activity for storing facts about objects was divided into some small number ( $\sim 8$ ) of phase periods. Then all the properties of the ball could be active in phase 3 and those of the book in phase 5. When our sample sentence was heard, the system would just add the new ownership facts to the appropriate phase. Such a system is computationally feasible and has been extensively discussed (Barrett et al. 2008). But despite some suggestive early experimental findings, there is good evidence (Shadlen and Movshon 1999) that this is not the mechanism that the brain employs.

### Population codes

Another terminological confusion arises in the use of “population codes” or “distributed representations”. In the

past, there were heated debates about whether neural representations were basically punctuate with a “grandmother cell” (Gross 2002; Bowers 2009) for each element of interest or basically holographic (with each item represented by a pattern involving all the units in a large population). It has been known for decades (Feldman 1988) that neither extreme could be computationally feasible for the neural systems of nature.

Having just one neuron coding an element of interest (concept) is impossible for several reasons. The most obvious is that the known death of cells would cause concepts to disappear. Also, the firing of individual cells is stochastic and would not be a reliable representation. Computationally, it is easy to see that there are not nearly enough neurons in the brain to capture all the possible combinations of shapes, sizes, colors, etc. that we recognize, let alone all the non-visual concepts. In fact, the pure grandmother cell story has always been a straw man—using a small number ( $\sim 10$ ) cells per concept would overcome all these difficulties.

The holographic alternative is more attractive because it is studied with the techniques of statistical mechanics. But it is equally implausible. This is easy to see informally and was established technically at least as early as (Willshaw et al. 1969). Suppose that we want to represent some set of concepts (e.g., English words) as a pattern of activity over some number  $N$  (say 10,000,000) neurons. The key problem is cross-talk: if multiple words are simultaneously active, how can we avoid interference among their respective patterns. Willshaw showed that the best answer is to have each concept represented by the activity of only about  $\log N$  units, which would be about 24 neurons in our example. There are many other computational problems with holographic models (Feldman 1988); for example if a concept required a pattern over all  $N$  units, how would that concept combine with other concepts or be transmitted to other brain regions.

There is now a wide range of converging experimental evidence (Bowers 2009; Olshausen and Field 1996; Vinje and Gallant 2000; Purushothaman and Bradley 2005; Quiroga et al. 2008a, b) showing that neural coding relies on the behavior of a modest number (tens to hundreds) of units. There is also overlap—the same neuron is often involved in the representation of different items. For various reasons, not all of them technical, some people continue to refer to these sparse representations as “population codes”. An equivalent, and much more appropriate, characterization would be *redundant circuits*.

One reason for the early suggestions proposing holographic codes is that, in many brain areas, a large fraction of the neurons fire in response to a relevant event. What could they all be doing if not jointly coding that event? We have already encountered the basic answer—there are

many possible interpretations of an isolated stimulus (whether sensory or deeper). For the required rapid response, it is optimal to consider the possibilities in parallel. In other words, the population firing pattern is basically *disjunctive*, not conjunctive. Of course, not all interpretations are equally likely and the population firing pattern can be viewed as encoding a probability distribution over the possible causes of the input (Barlow 2001; Jazayeri and Movshon 2006).

A related terminological problem involves two distinct uses of the term “sparse”. As described above, the more common usage refers to the fact that the neural representation of some item (e.g., a sound or an image) is carried by a small fraction of the population of neurons in the relevant neural area. But the term is also used (e.g., in the Kreiman (2004) survey) to refer to the fact that the firing of neural spikes is sparse in time. Rates greater than 100 spikes/second are unusual and there are systems with much slower base rates. Neural communication is sparse in both time and unit count and there are the usual metabolic pressures that require this (Lennie 2003).

In summary, there is now a broad consensus on neural spike signaling. There are a number of specialized structures involving delicate timing and the relative time of spike arrival is important for plasticity. But the main mechanism for spike signaling is frequency coding in specific circuits of moderate redundancy. Current research on “the neural code” is focused on how neural systems deal with information and decisions under uncertainty.

### Probability, utility, and fitness

Although people are not perfect utility maximizers, there is no question that much of human behavior is describable in terms of probabilities and utilities. In recent years, this has given rise to a renewed interest in “the neural code” for decision making. The hypothesis is that there are *general* mechanisms of representing probabilities and utilities and associated decision rules. Unfortunately, there are indications of a new round of confused reasoning based on the metaphor of the neural code and the communication channel model of neural computation.

Noise is inherent both in our perception of the world and in the chemistry of neural firing. So any full explanation of neural spikes will need to be probabilistic in some way. This is generally accepted and, in addition, almost all treatments include explicit prior probability estimates and are therefore Bayesian. Further, some information and decisions are more important than others so utility theory must play an important role. There is a great deal of elegant and informative work on theoretical and experimental Bayesian modeling of neural communication and decision

and much is being learned this way. A related set of developments is part of the Neuroeconomics effort. The (Glimcher 2009) collection has introductory articles on all aspects of neurally based utility research.

Much of the work in neuroeconomics focuses on the relation between animal (usually human) decision making and possible neural substrates. This can be seen as an extension of the earlier search for “the neural code” discussed in the sections on “[Introduction and background](#)”, “[Spikes in single neuron communication](#)”, “[Spikes in neural circuits](#)”, “[Population codes](#)” of this article. The underlying belief is that there is some universal encoding of the phenomenon. As with the earlier efforts, this entails the risk of oversimplifying an operation, here decision making, which has many manifestations at multiple levels in all animals. For example, Paul Glimcher (2009, p. 508), an unquestioned leader in Neuroeconomics, presents what seems to be a *dictum ex cathedra* that “utility is ordinal”. Ordinal utility is preference without quantitative norms and this is indeed much easier to assess in people. But animals (including people) are constantly making multidimensional choices and it would be impossible to do this without some calibration of the relative strengths of all competing drives. Experimentally, there are already results suggesting that both ordinal and cardinal utilities are neurally encoded (Pine et al. 2009; Platt and Padoa-Schioppa 2009, p. 448). More importantly, the whole idea of claiming universality from general reasoning or from one constrained experiment is misguided.

But the Neuroeconomics effort is producing some valuable insights into the possible neural processing of probabilities and utilities. Bayesian posterior probabilities have a simple relation to numerical (cardinal) utilities in a single trial decision based on passive observation. The expected utility of each choice is just the sum of the probability, under that choice, of each possible outcome weighted by the utility of that outcome. A standard example is an animal choosing between two possible unknown food sources, which it believes have different probabilities of having two kinds of food. The expected value of each potential source is the just the sum of the values of the two food types, weighted by the probability of finding each at that source. It turns out that even simple animals come fairly close to optimal foraging strategies in rather complex situations (Kamil et al. 1987). No one believes that, e.g., insects, have explicit neural representation of expected utilities—effective foraging is an evolutionary requirement (Parker 2006).

This is our first direct encounter with the topic of “[EEU and the functions of neural spike signaling](#)” of this article: *Ecological Expected Utility* (EEU). The central insight of Neuroeconomics is that the notion of “maximizing expected utility” from economics is a powerful tool for

helping to understand neural computation. I suggest that expected utility (EU) does provide an appropriate criterion for modeling neural communication and computation, but that a much richer and more subtle notion of utility, which could be called “ecological utility”, is needed.

In its most general form, this can be seen as a formalization of the core biological idea of evolutionary fitness. As the great biologist Theodosius Dobzhansky (1900–1975) famously stated “Nothing in biology makes sense except in the light of evolution”. Animals are effective foragers, because lineages that were not efficient lost out to competitors (McDermott et al. 2008). Much of the current work on “the neural code” for utility and decisions attempts to internalize representation and computation to specific dedicated mechanisms. As in the earlier work, seeking a “code” masks the inherently recurrent nature of neural computing. A prototypical experiment studies monkeys on one highly constrained task and makes a mathematical model of the neural signals in one pathway. Essentially all of these studies are strictly feed forward. The survey by Knill and Pouget (2004) discusses several computational neural models of uncertainty.

As soon as we include any active information gathering strategy, the simple link between posterior probability and expected utility breaks down. The main source of uncertainty for animal decisions is not noise, but *limited observability*. We usually do not have access to all of the information that might be helpful in making an action/decision. It is obviously better to gain information about something that is more important to you. Again, this has been the subject of extensive research in decision theory as the “value of information” (Feldman and Sproul 1977; Behrens et al. 2007). One important application of these ideas to neural systems is in the study of information gathering saccades (voluntary eye movements) and covert attention. This will be discussed in the next section; the main point here is that the utility function that must be maximized in animals is much more complex than that of a visual discrimination task.

### **EEU and the functions of neural spike signaling**

We are now in a position to characterize the various known functions of neural spikes and the roles they play in behavior. From the EEU perspective, the basic question concerns what kind of action/decision is supported by the various spiking disciplines. Let’s start from the two extremes. At the low end, there are reflexes and also a number of behaviors that are triggered by a single spike, either in isolation or as the first spike (or wave of spikes) from some stimulus. At the other extreme consider a major life decision, like whether to accept a job offer. This kind

of binary action/decision usually requires an extended period of active information gathering, often involving considerable travel and time. These two extremes are often described as Type 1 and Type 2 decisions (Kahneman 2003). As we will now show, there is actually a continuum of action/decision processes involving varying amounts of processing and information gathering.

One of the most productive techniques of cognitive psychology involves studying people's eye movements as they execute behaviors. These eye-movement studies in psychology pre-suppose that saccade planning is central to optimal behavior. Obviously enough, where you look has a much stronger effect on the information gained than any details of subsequent processing (Yarbus 1967).

We are constantly making large and small decisions. Every act of perception involves decisions, for example disambiguating words as in the “rose” example of Fig. 1. There is an enormous range of neural decision making, with a vast array of different information gathering and evaluating strategies. For concreteness, we will focus on overt (saccades) and covert visual attention. As is well known, people make three or four saccades per second and each of these is goal driven action/decision. The recent survey by Gollish (Gollisch 2009) is largely concerned with neural signaling in early vision, taking overt and covert attention seriously. There are a number of well known effects of saccades on the visual signals transmitted from the retina. During saccades there is signal suppression. Also “efference copy” provides the visual system with a prediction of the saccade target and this is used to prime that area. There is some shift in receptive fields towards the target and this helps maintain coherence. More surprisingly some (rabbit) retinal cells seem to switch polarity (from ON to OFF) shortly after a saccade.

We do know a fair amount about the brain circuits that plan saccades (and covert attention choices), but taking this seriously requires explicitly modeling active perception. It is perfectly possible to technically treat the choice of where to next saccade as an “expected value of information” computation in utility theory (Torralba et al. 2006; Brodersen et al. 2008). Torralba et al. model the choice of saccade targets as Bayesian optimization combining local feature information with global measurements suggesting the general scene type. Interestingly, the model (and people) does not need to know the general scene type to take advantage of it. The Torralba paper includes detailed comparison of scan patterns of the subjects and the model. Again, the metaphor of a neural code just does not fit.

This is not to dismiss the search for common mechanisms. Even within the restricted feed-forward paradigm, it is possible to suggest some possible general methods of neural decision making. The (Gold and Shadlen 2007) review article presents an extremely clear description of the

economic and neural background and many of the central issues within the framework of simple sensory-motor tasks. This is the best current introduction to the field and totally avoids the notion of a “neural code” although the idea of decision variables plays a crucial role in their treatment.

Another interesting hypothesis comes from the Ganguli et al. (2008) model of a study of attention and distractors in the Lateral IntraParietal visual area LIP. Monkeys were trained to make a saccade after a delay to the position of the ring target that had a gap. On half the trials, a distractor ring was flashed during the delay. Attention, as measured by improved contrast sensitivity, would be split between the remembered saccade target and the distractor location. The main finding is that the time for the sensitivity enhancement to equalize between the two targets was almost constant for each animal, independent of location and of the base firing rate of the neurons involved. The paper includes an elaborate sparse coding model, but the main proposal does not depend strongly on the details. They suggest that the time to recover from a distractor should be constant and needs to be a system (rather than a local) property of the LIP network, because units differ in their local dynamics. There could be a general architecture that supports this functionality.

More generally, there are a number of subtleties involved in defining an ecologically appropriate notion of utility. Evolutionary selection for fitness guarantees that an animal will not be too incompetent. Of course, selection does not operate only on individuals so species survival may well depend on actions that are not, by any direct measure, optimal for the individual.

Even within one individual, the expected utility to be maximized should be amortized over life experience, including adaptation and learning. Further, even within one behavioral episode, the optimal behavior is often rather more complex than the current industry standard “Bayesian Brain” story would suggest. As we saw above, looking for general optimality in a single feed-forward perceptual discrimination task has the character of looking for the lost ring under the streetlight.

## Summary

Neural spikes are an evolutionary ancient innovation that remains nature's exclusive mechanism for rapid long distance information transfer. There is no unique “neural code”, but the number of spike-based communication techniques is limited. From our perspective there are not so many qualitatively different functional roles played by neural spikes and each of these is best understood as optimizing some EEU requirement. Each of these



mechanisms can be seen as a highly effective strategy for solving different information processing problems.

Ecological Expected Utility provides a unifying theme for the communication functions sub-served by neural spikes. EEU is determined by evolutionary fitness of the organism's genome. It cannot be computed directly and is not static over time, but it is what nature optimizes. By keeping this explicitly in mind, we can come to a better understanding of neural signaling and brain function.

Again obviously, animals often need to act before taking time to fully consider all their options and seek additional input. This can also be formalized within a utility theory framework and some work along these lines has been done in AI (at least) under the title “anytime planning” (Likhachev et al. 2008). The need for rapid, approximately optimal, actions appears to have everything to do with the neural coding. In fact, we can do fairly well by assuming that neural spike signaling evolved to do as little as possible for each task.

The minimal neural action/decision is the *reflex*, which can be monosynaptic. Even reflexes are conditioned by top down activation; neural systems are never just passive. The next level of complexity is multistep spike wave signaling, as discussed in the section “[Spikes in single neuron communication](#)”; some pre-specified action/decisions can be chosen by the temporally first of competing spike signals. The third level of complexity involves feedback loops. These can be local as in lateral inhibition or across brain regions as evidenced by the ubiquitous bidirectional connections between levels. Neural loops inherently involve settling time and therefore slower decisions. Loops also entail that the signals on a given pathway change over time.

All of the mechanisms discussed above are passive; they do not take into account the animals own actions to gain information. In the previous section we reviewed some of the most basic information gathering strategies, overt and covert visual attention. But of course there are many others involving bodily movement, active exploration, language, etc. To be meaningful, any neural signals from perception must be interpreted with respect to the context in which they are received, again eliminating the possibility of a fixed neural code.

Neural spikes are a restricted mechanism, but it's the only one we've got for rapid long distance signaling. Spikes are limited in speed, accuracy, reliability, and bandwidth. There is no unique neural code, but instead a wonderful collection of mechanisms that exploit neural signaling for a remarkable set of functions. The current focus on studying the general semantic content of neural signals is an essential component of understanding the brain, but labeling this effort as a search for “the neural code” remains a profoundly bad idea.

**Acknowledgments** I would like to thank Jose Carmena, Joachim Diederich, Sridhar Narayanan, David Zipser, and the referees for helpful comments and discussion.

**Open Access** This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

## References

- Allman JM (1999) *Evolving brains*. Scientific American Press, New York, p 16
- Atwell D, Laughlin SB (2001) An energy budget for signaling in the grey matter of the brain. *J Cereb Blood Flow Metab* 21:1133–1145
- Barlow H (2001) Redundancy reduction revisited. *Netw Comput Neural Syst* 12:241–253
- Barrett L, Feldman JA, MacDermed L (2008) A somewhat new solution to the variable binding problem. *Neural Comput* 20(9):2361–2378
- Behrens TE, Woolrich MW, Walton ME et al (2007) Learning the value of information in an uncertain world. *Nat Neurosci* 10(9):1214–1221
- Bender V, Bender K, Brasier DJ, Feldman DE (2006) Two coincidence detectors for spike timing-dependent plasticity in somatosensory cortex. *J Neuro* 26:4166–4177
- Bowers JS (2009) On the biological plausibility of grandmother cells: implications for neural network theories in psychology and neuroscience. *Psychol Rev* 116(1):252–282
- Brodersen KH, Penny WD, Harrison LM et al (2008) Integrated Bayesian models of learning and decision making for saccadic eye movements. *Neural Netw* 21(9):1247–1260
- Carpenter GA, Grossberg S (1987) A massively parallel architecture for a self-organizing neural pattern recognition machine. *Comput Vis Graph Image Process* 7:54–115
- Carr CE (1993) Processing of temporal information in the brain. *Annu Rev Neurosci* 16:223–243
- Cohen JD, McClure SM, Yu AJ (2007) Should I stay or should I go? Exploration versus exploitation. *Philos Trans R Soc B Biol Sci* 362:933–942
- Dayan P, Abbott LF (2001) *Theoretical neuroscience: computational and mathematical modeling of neural systems*. MIT Press, Cambridge
- de Charms RC, Zador A (2000) Neural representation and cortical code. *Annu Rev Neurosci* 23:613–647
- Feldman JA (1988) Computational constraints on higher neural representations. In: Schwartz E (ed) *Proceedings of the system development foundation symposium on computational neuroscience*. Bradford Books/MIT Press, Cambridge, April 1988
- Feldman JA (2006) *From molecule to metaphor: a neural theory of language*. Bradford Books/MIT Press, Cambridge
- Feldman JA, Sproul RF (1977) Decision theory and AI II: the hungry monkey. *Cogn Sci* 2:158–192
- Ganguli S, Bisley JW, Roitman JD et al (2008) One-dimensional dynamics of attention and decision making in LIP. *Neuron* 58:15–25
- Gerstner W, Kreiter AK, Markham H, Herz AVM (1997) Neural codes: firing rates and beyond. *PNAS* 94:12740–12741
- Gerstner W, Kistler WM (2002) *Spiking neuron models: single neurons, populations, plasticity*. Cambridge University Press, Cambridge

- Glimcher PW (2009) Choice: towards a standard back-pocket model. In: Glimcher PW, Camerer C, Poldrack RA et al (eds) *Neuroeconomics: decision making and the brain*, 1st edn. Academic Press, San Diego
- Gold JJ, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574
- Gollisch T (2009) Throwing a glance at the neural code: rapid information transmission in the visual system. *HFSP J* 3:36–46
- Gross CG (2002) Genealogy of the “Grandmother Cell”. *Neuroscientist* 8(5):512–518
- Izhikevich EM (2007) Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cereb Cortex* 17:2443–2452
- Jazayeri M, Movshon JA (2006) Optimal representation of sensory information by neural populations. *Nat Neurosci* 9(5):690–696
- Kahneman D (2003) Maps of bounded rationality: psychology for behavior economics. *Am Econ Rev* 93(5):1449–1475
- Kamil AC, Krebs JR, Pulliam HR (1987) *Foraging behavior*. Plenum Press, New York and London
- Kara P, Reid RC (2003) Efficacy of retinal spokes in driving cortical responses. *J Neurosci* 23(24):8547–8557
- Katz PS (2007) Evolution and development of neural circuits in invertebrates. *Curr Opin Neurobiol* 17(1):59–64
- Kirchner H, Thorpe SJ (2006) Ultra-rapid object detection with saccadic eye movements: visual processing speed revisited. *Vis Res* 46(11):1762–1776
- Knill DC, Pouget A (2004) The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci* 27(12):712–719
- Kreiman G (2004) Neural coding: computational and biophysical perspectives. *Phys Life Rev* 1(2):71–102
- Lennie P (2003) The cost of cortical computation. *Curr Biol* 13(6):493–497
- Likhachev M, Ferguson D, Gordon G et al (2008) Anytime search in dynamic graphs. *Artif Intell* 172(14):1613–1643
- McDermott R, Fowler JH, Smirnov O (2008) On the evolutionary origin of prospect theory preferences. *J Polit* 70(2):335–350
- Meech RW, Mackie GO (2007) Evolution of excitability in lower metazoans. In: North G, Greenspan RJ (eds) *Invertebrate neurobiology*. Cold Spring Harbor Laboratory Press, New York
- Olshausen BA, Field DJ (1996) Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381(6583):607–609
- Parker G (2006) Behavioral ecology: natural history as science. In: Lucas JR, Simmons LW (eds) *Essays in animal behavior*. Academic Press, New York
- Pine A, Seymour B, Roiser JP, Bossaerts P, Friston K, Curran HV, Dolan RJ (2009) Encoding of marginal utility across time in the human brain. *J Neurosci* 29(30):9575–9581
- Platt M, Padoa-Schioppa C (2009) Neuronal representations of value. In: Glimcher PW, Camerer C, Poldrack RA et al (eds) *Neuroeconomics: decision making and the brain*, 1st edn. Academic Press, San Diego
- Purushothaman G, Bradley DC (2005) Neural population code for fine perceptual decisions in area MIT. *Nat Neurosci* 8(1):99–106
- Quiroga RQ, Kreiman G, Koch C, Fried I (2008a) Sparse but not “Grandmother-cell” coding in the medial temporal lobe. *Trends Cogn Sci* 12(3):87–91
- Quiroga RQ, Mukamel R, Isham EA et al (2008b) Human single-neuron responses at the threshold of conscious recognition. *Proc Natl Acad Sci USA* 105(9):3599–3604
- Reike F, Warland D, de Ruytter R, Bialek W (1997) *Spikes: exploring the neural code*. MIT Press, Cambridge
- Rousselet GA, Mace MJ, Thorpe SJ, Fabre-Thorpe M (2007) Limits of event-related potential differences in tracking object processing speed. *J Cogn Neurosci* 19(8):1241–1258
- Serre T, Oliva A, Poggio T (2007) A feedforward architecture accounts for rapid categorization. *Proc Natl Acad Sci USA* 104(15):6424–6429
- Shadlen M, Movshon J (1999) Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* 24:67–77
- Shastri L (2001) A computational model of episodic memory formation in the hippocampal system. *Neurocomputing* 38:889–897
- Snyder RL, Sinex DG (2002) Immediate changes in tuning of inferior colliculus neurons following acute lesions of cat spiral ganglion. *J Neurophysiol* 87(1):434–452
- Tanenhaus MK, Leiman JM, Seidenberg MS (1979) Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *J Verbal Learn Verbal Behav* 18(4):427–440
- Torralba A, Oliva A, Castelhano MS, Henderson JM (2006) Contextual guidance of eye movements and attention in real-world scenes: the role of global features in object search. *Psychol Rev* 113(4):766–786
- Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287(5456):1273–1276
- Willshaw DJ, Buneman OP, Longuet-Higgins HC (1969) Non-holographic associative memory. *Nature* 222:960–962
- Yarbus AL (1967) *Eye movements and vision*. Plenum, New York (Originally published in Russian 1962)

## Related works

- Abbott LF (1994) Decoding neuronal firing and modeling neural networks. *Q Rev Biophys* 27:291–331
- Aur D, Jog MS (2007) Reading the neural code: what do spikes mean for behavior? *Nat Precedings*. doi:[10.1038/npre.2007.61.1](https://doi.org/10.1038/npre.2007.61.1)
- Barber MJ, Clark JW, Anderson CH (2003) Neural representation of probabilistic information. *Neural Comput* 15:1843–1864
- Beck JM, Ma W-J, Kiani R et al (2008) Probabilistic population codes for Bayesian decision making. *Neuron* 60(6):1142–1152
- Bisley JW, Goldberg ME (2006) Neural correlates of attention and distractibility in the lateral intraparietal area. *J Neurophysiol* 95:1696–1717
- Burr D, Tozzi A, Morrone MC (2007) Neural mechanisms for timing visual events are spatially selective in real-world coordinates. *Nat Neurosci* 10(4):423–425
- Butts DA, Weng C, Jin J (2007) Temporal precision in the neural code and the timescales of natural vision. *Nature* 449:92–95
- Corrado G, Doya K (2007) Understanding neural coding through the model-based analysis of decision making. *J Neurosci* 27(31):8178–8180
- Corrado GS, Sugrue LP, Brown JR et al (2009) The trouble with choice: studying decision variables in the brain. In: Glimcher PW, Camerer C, Poldrack RA (eds) *Neuroeconomics: decision making and the brain*, 1st edn. Academic Press, London San Diego
- Deneve S (2008) Bayesian spiking neurons I: inference. *Neural Comp* 20:91–117
- Deneve S, Latham PE, Pouget A (1999) Reading population codes: a neural implementation of ideal observers. *Nat Neurosci* 2(8):740–746
- Eggermont JJ (2001) Between sound and perception: reviewing the search for a neural code. *Hear Res* 157:1–42
- Ehinger K, Hidalgo-Sotelo B, Torralba A, Oliva A (2009) Modeling search for people in 900 scenes: a combined source model of eye guidance. *Vis Cogn* (in press)
- Faisal AA, Selen LPJ, Wolpert DM (2008) Noise in the nervous system. *Nat Rev Neurosci* 9(4):292–303

- Fox CR, Poldrack RA (2009) Prospect theory and the brain. In: Glimcher PW, Camerer C, Poldrack RA et al (eds) *Neuroeconomics: decision making and the brain*, 1st edn. Academic Press, San Diego
- Georgopoulos AP et al (1986) Neuronal population coding of movement direction. *Science* 233:1416–1419
- Gerstner W, Kistler WM (2002) *Spiking neuron models: single neurons, populations, plasticity*. Cambridge University Press, Cambridge
- Glimcher PW, Rustichini A (2004) Neuroeconomics: the conciliation of brain and decision. *Science* 306:447–454
- Globerson A, Stark E, Vaadia W et al (2009) The minimum information principle and its application to neural code analysis. *Proc Natl Acad Sci USA* 106(9):3490–3495
- Gold JI, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. *Trends Cogn Sci* 5(1):10–16
- Gutnisky DA, Dragoi V (2008) Adaptive coding of visual information in neural populations. *Nature* 452(7184):220–224
- Kiebel SJ, Daunizeau J, Friston KJ (2008) A hierarchy of time-scales and the brain. *PLoS Comput Biol*. doi:10.1371/journal.pcbi.1000209
- Lee D, Wang X-J (2009) Mechanisms for stochastic decision making in the primate frontal cortex: single-neuron recording and circuit modeling. In: Glimcher PW, Camerer C, Poldrack RA et al (eds) *Neuroeconomics: decision making and the brain*, 1st edn. Academic Press, San Diego
- Ma WJ, Beck JM, Latham PE et al (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9(11):1432–1438
- Passaglia C, Dodge F, Herzog E et al (1997) Deciphering a neural code for vision. *Proc Natl Acad Sci USA* 94(23):12649–12654
- Pillow JW, Shlens J, Paninski L et al (2008) Spatio-temporal correlations and visual signaling in a complete neuronal population. *Nature* 454(7207):995–999
- Pouget A, Dayan P, Zemel RS (2003) Inference and computation with population codes. *Ann Rev Neurosci* 26:283–410
- Ratcliff R, Smith PL (2004) A comparison of sequential sampling models for two-choice reaction time. *Psychol Rev* 111:333–367
- Rozell CJ, Johnson DH, Baraniuk RG, Olshausen BA (2008) Sparse coding via thresholding and local competition in neural circuits. *Neural Comput* 20:2526–2563
- Rustichini A (2009) Neuroeconomics: formal models of decision making and cognitive neuroscience. In: Glimcher PW, Camerer C, Poldrack RA et al (eds) *Neuroeconomics: decision making and the brain*, 1st edn. Academic Press, San Diego
- Schwartz O, Hsu A, Dayan P (2007) Space and time in visual context. *Nat Rev Neurosci* 8(7):522–535
- Shafir S, Reich T, Tsur E et al (2008) Perceptual accuracy and conflicting effects of certainty on risk-taking behaviour. *Nature* 453(7197):917–920
- Theunissen F, Miller JP (1995) Temporal encoding in nervous systems: a rigorous definition. *J Comput Neurosci* 2(7):149–162
- Tiesinga P, Fellous J-M, Sejnowski TJ (2008) Regulation of spike timing in visual cortical circuits. *Nat Rev Neurosci* 9(2):97–109
- Trepel C, Fox CR, Poldrack RA (2005) Prospect theory on the brain? Toward a cognitive neuroscience of decision under risk. *Cogn Brain Res* 23(1):34–50
- Yu AJ, Dayan P (2005) Uncertainty, neuromodulation and attention. *Neuron* 46(4):681–692
- Zemel RS, Dayan P, Pouget A (1997) Probabilistic interpretation of population codes. *Neural Comp* 10(2):403–430