

Cortical Prediction Markets

David Balduzzi
ETH Zürich, Switzerland
david.balduzzi@inf.ethz.ch

ABSTRACT

We investigate cortical learning from the perspective of mechanism design. First, we show that discretizing standard models of neurons and synaptic plasticity leads to rational agents maximizing simple scoring rules. Second, our main result is that the scoring rules are *proper*, implying that neurons faithfully encode expected utilities in their synaptic weights and encode high-scoring outcomes in their spikes. Third, with this foundation in hand, we propose a biologically plausible mechanism whereby neurons backpropagate incentives which allows them to optimize their usefulness to the rest of cortex. Finally, experiments show that networks that backpropagate incentives can learn simple tasks.

Categories and Subject Descriptors

[Theory of computation]: Multi-agent learning; [Theory of computation]: Algorithmic mechanism design

Keywords

incentives for cooperation; multiagent learning; biologically-inspired approaches; prediction markets; neural networks

1. INTRODUCTION

How does the brain encode information about the environment into its structure [27]? Inspired by recent work in prediction markets, we apply mechanism design to investigate cortical learning and the neural code [1–3, 12, 15, 19]. To the best of our knowledge this is the first paper to do so.

We start in §2 by modeling neurons as rational agents: that is, agents whose sole aim is to maximize the expected value of an objective function. To do so, we draw on a recent paper showing that *discretizing* standard models of neuronal dynamics [14] and learning [26] yields a threshold neuron with an online update rule that optimizes a simple objective [5]. By maximizing their objective function, neurons seek to optimally trade off rewards, depending on neuromodulatory signals such as dopamine, with costs, depending on resources expended on synaptic connections [6, 7].

However, it is not enough that neurons optimize locally. They should collectively converge on useful outcomes. The

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*
Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

problem of how a global (cortical) optimization procedure can be implemented at a local (neuronal) level remains open.

To tackle the problem we turn to mechanism design: How to incentivize populations of rational agents to produce desirable outcomes? An inspiring successful application of mechanism design is prediction markets, which aggregate the behavior of self-interested traders into accurate predictions of diverse real-world events [8, 21]. This has motivated research on payment schemes that encourage agents to trade in markets if the price distribution differs from their beliefs [15]. Of particular interest are *proper* scoring rules: payment schemes that incentivize rational agents to truthfully report their beliefs [19].

Our next step, §3, is therefore to analyze neuronal objective functions as *payment schemes*. This has implications in two directions. First, since the neuronal objective function decomposes as a sum over synapses, we model synapses as rational agents trading in a neuronal market, §3.1. Second, we model neurons as rational agents trading in a cortical prediction market, §4.3.

Our main result, Theorem 5, establishes a striking connection between prediction markets and cortical learning: neuronal objective functions are proper scoring rules. The remainder of the paper applies two corollaries of Theorem 5 to show that well-functioning neuronal markets form a foundation for a well-functioning cortical market – thereby gluing together the two perspectives.

Corollary 6 shows that synaptic weights encode the utility expected after pre- and post-synaptic spikes. This partially answers the question posed earlier: “How does the brain encode information about the environment into its structure?”

More importantly, the corollary provides a foundation for cooperative learning. Consider the following basic schema to incentivize rational agents to collaborate:

- (i) each agent estimates its usefulness to other agents,
- (ii) incorporates the estimate into its reward function and
- (iii) thus maximizes its usefulness to the collective.

To implement the schema, neurons must estimate their usefulness. Corollary 6 implies that synaptic weight \mathbf{w}_{ij} quantifies how useful spikes from n_i are to n_j , when n_j spikes. More generally, the sum of outgoing synaptic connections quantifies how useful a neuron’s outputs are to the rest of the system. We therefore define the *usefulness* of a neuron as, roughly, the sum of its downstream weights, §4.

In line with the schema we then show, Corollary 7, that incorporating feedback into reward functions causes neurons to (i) estimate their usefulness and (ii) maximize the estimate. This provides a new interpretation of a spike-based

backpropagation scheme [23] that is closely related to error-backpropagation [24]. In short, well-functioning neuronal markets, with synapses faithfully reporting expected utilities, can be used to build well-functioning cortical markets.

Finally, experiments in §5 confirm our theoretical results.

Scope and related work.

A well-studied framework in neuroscience is based on the idea that neurons infer the probabilities of external events, which are encoded into probabilistic population codes, see e.g. [9]. By contrast, we emphasize decisions over inferences. We are concerned with how neurons act, rather than what they infer. The two perspectives are related and it may turn out, as in prediction markets where prices can encode probabilities, that the population coding and mechanism design approaches lead to the same destination.

Our goal is to show methods from mechanism design can be fruitfully applied to fundamental questions in neuroscience. We do not advocate *specifically* the scoring rules below. These were derived from standard, but simple, neurophysiological models. It should be possible to extend our results to more realistic models.

This paper is inspired by recent work showing that carefully designed markets can be used to aggregate hypotheses generated by populations of learning algorithms [2, 20, 28].

2. A MINIMAL MODEL

At first glance, the models developed by neuroscientists are quite different from the rational agents studied in game theory. To build a bridge we utilize recent work *discretizing* a standard model from the neuroscience literature [5].

2.1 Discretized neurons

Consider a system of N binary neurons $\{n_j\}_{j=1}^N$. Let $\mathcal{O} = \{0, 1\}^N$ denote the set of possible states. Each neuron is connected to a subset of the system. Suppose neuron n_j has $K_j \ll N$ synapses. We model the restriction of the total system state to the subset received by neuron n_j with a mask projecting from $\{0, 1\}^N$ to $\{0, 1\}^{K_j}$

$$\varphi_j : \mathcal{O} \rightarrow \{0, 1\}^{K_j} : \mathbf{x} = (x_1, \dots, x_N) \mapsto (x_i)_{\{i \mid i \rightarrow j\}}. \quad (1)$$

Neuron n_j is equipped with a K_j -vector of synaptic weights, $\mathbf{w}_j \in \mathcal{H}_j = \mathbb{R}^{K_j}$. Given input $\mathbf{x} \in \mathcal{O}$, the neuron outputs a 0 or 1 according to

$$f_{\mathbf{w}_j}(\mathbf{x}) := \begin{cases} 1 & \text{if } \langle \mathbf{w}_j, \varphi_j(\mathbf{x}) \rangle - \vartheta > 0 \\ 0 & \text{else} \end{cases} \quad (2)$$

for some fixed ϑ constant across all neurons.

To simplify the exposition, we drop φ_j from the notation and let $\mathcal{H} := \mathbb{R}^N$ denote the space of synaptic weights – where synapses that do not physically exist are implicitly clamped to zero. Thus, we treat entire system states as inputs to a neuron – when in fact the mask projects out most inputs.

DEFINITION 1. *Suppose we have utility function $\mu : \mathcal{O} \rightarrow \mathbb{R}$. Following [5], define **reward function***

$$R(\mathbf{x}, \mathbf{w}_j, \mu_j) = \underbrace{\mu_j(\mathbf{x})}_{\text{utility}} \cdot \underbrace{(\langle \mathbf{w}_j, \mathbf{x} \rangle - \vartheta)}_{\text{margin}} \cdot \underbrace{f_{\mathbf{w}_j}(\mathbf{x})}_{\text{selectivity}} \quad (3)$$

Examples of utility functions are provided in §2.2 and §4.1.

REMARK 1 (NOTATION FOR SPIKES). *Note that $f_{\mathbf{w}_j}(\mathbf{x})$, \mathbf{x}_j , $\mathbb{1}_{\mathbf{w}_j}$, and $\mathbb{1}_j$ all denote the output of neuron n_j ; emphasizing the function producing the output, that the output is also an input (one of many forming a vector) to other neurons, or the indicator-function aspect of the output respectively. We use $\mathbb{1}_{ij} := \mathbb{1}_i \cdot \mathbb{1}_j$ to indicate the cospiking of neurons n_i and n_j .*

Ignoring costs for a moment, suppose neurons maximize $\mathbb{E}_{(\mathbf{x}, \mu) \sim P} [R(\mathbf{x}, \mathbf{w}, \mu)]$, where $P(\mathbf{x}, \mu)$ is the joint distribution on spiking inputs and neuromodulators.

The reward function is continuously differentiable (in fact, linear) as a function of \mathbf{w} everywhere except at the kink $\langle \mathbf{w}, \mathbf{x} \rangle = \vartheta$ where it is continuous but not differentiable. We can therefore perform gradient ascent to obtain synaptic updates

$$\Delta \mathbf{w}_{ij} \propto \mu_j(\mathbf{x}) \cdot \mathbf{x}_i \cdot f_{\mathbf{w}_j}(\mathbf{x}) = \mu_j(\mathbf{x}) \cdot \mathbb{1}_{ij}. \quad (4)$$

In short, if n_j receives input $\mathbb{1}_i$ and subsequently spikes $\mathbb{1}_j$, then synapse $i \rightarrow j$ is modified proportionally to $\mu_j(\mathbf{x})$. The main theorem in [5] derives the above equations by discretizing standard models of neuronal dynamics and learning:

THEOREM 1 (DISCRETIZED NEURONS, [5]). *The fast time constant limit of Gerstner’s Spike Response Model [14] is (2). Taking the fast time constant limit of STDP [26] yields (4) with $\mu_j(\mathbf{x}) = 1$. Finally, STDP is gradient ascent on a reward function whose limit is (3).*

Spike-timing dependent plasticity is prone to overpotentiation [26], leading to epileptic seizures. In the neuroscience literature, weights are typically controlled with a depotentiation bias. We take an alternative approach, by introducing a *regularizer* $A_\bullet(\mathbf{w})$, which quantifies the *resource costs* incurred by high synaptic weights [5, 6, 18].

The optimal weights are then computed according to

$$\mathbf{w}_j^* := \operatorname{argmax}_{\mathbf{w} \in \mathcal{H}} \mathbb{E}_P [\mathcal{S}_\bullet(\mathbf{x}; \mathbf{w})] \quad (5)$$

$$:= \operatorname{argmax}_{\mathbf{w} \in \mathcal{H}} \mathbb{E}_P [R(\mathbf{x}, \mathbf{w}, \mu) - A_\bullet(\mathbf{w})] \quad (6)$$

where *scoring rule* $\mathcal{S}_\bullet(\mathbf{x}, \mathbf{w})$ balances rewards $R(\mathbf{x}, \mathbf{w}, \mu)$ against costs $A_\bullet(\mathbf{w})$. We consider two standard regularizers taken from machine learning [25] and a third, more biologically plausible, taken from [5]:

$$\begin{cases} A_2(\mathbf{w}_j) = \frac{1}{2\eta} \|\mathbf{w}_j\|_2^2 & \ell_2 \\ A_H(\mathbf{w}_j) = \frac{1}{\eta} \sum_i \mathbf{w}_{ij} \log \mathbf{w}_{ij} & \ell_H \\ A_1(\mathbf{w}_j) = \frac{1}{\eta} \|\mathbf{w}_j\|_1, \text{ where } 0 \leq \mathbf{w}_{ij} \leq 1 \text{ for all } i. & \ell_1 \end{cases}$$

Clearly ℓ_H is not a norm – we find the notation convenient.

Computing gradient ascent yields online updates

$$\Delta \mathbf{w}_{ij} \propto \mu_j(\mathbf{x}) \cdot \mathbb{1}_{ij} - \frac{1}{\eta} \cdot \begin{cases} \mathbf{w}_{ij} & \ell_2 \\ \log \mathbf{w}_{ij} + 1 & \ell_H \\ 1 & \ell_1 \end{cases} \quad (7)$$

REMARK 2 (REGULARIZERS). *Each regularizer has points in its favor. The ℓ_1 regularizer provides a simple interpretation of the saturated synaptic weights observed in some neurophysiological models [13]. The ℓ_2 regularizer allows negative synaptic weights, corresponding to inhibitory synapses. Finally, ℓ_H results in weights that can be interpreted as a probability distribution and is closely related to Hanson’s log-arithmic market scoring rule [11].*

2.2 Utility functions

Three biologically inspired utility functions are:

Example U1. (Feedforward, frequency). Utility function $\mu(\mathbf{x}) = 1$ encourages neurons to spike for inputs that are frequent and contain many spikes.

Example U2. (Feedforward, invariance). A more interesting utility function takes inputs over consecutive time steps $\mathbf{x} = (\mathbf{x}^{(t-1)}, \mathbf{x}^{(t)})$ as input and sets $\mu(\mathbf{x}) = f_{\mathbf{w}}(\mathbf{x}^{(t-1)})$. This encourages neurons to learn *stable* patterns containing many spikes, i.e. those that cause it to spike *twice* consecutively. The utility function can be extended across multiple time steps, possibly with a temporal discount factor.

Example U3. (Neuromodulators). Neuromodulatory systems signaling global rewards can be modeled via $P(\nu|\mathbf{x})$ where ν is a real-valued random variable: positive outcomes are reinforced and conversely. The utility is then $\mu(\mathbf{x}) := \mathbb{E}_{\nu \sim P(\nu|\mathbf{x})}[\nu|\mathbf{x}]$, where the expectation is with respect to the distribution on neuromodulators.

A fourth utility function is discussed in §4.1.

3. NEURONAL PREDICTION MARKETS

Scoring rules are schemes for paying agents based on their reports. Proper scoring rules, which incentivize agents to report truthfully, have proven useful in a wide range of settings including weather forecasts [10], prediction markets [15, 19] and crowdsourced learning mechanisms [2, 3].

Our main result, Theorem 5, is that the scoring rules \mathcal{S}_{\bullet} in (5) are proper for all three regularizers. The upshot is that a neuron’s synaptic weights faithfully encode¹ expectations about rewards after pre- and post- synaptic spiking activity. The form of the encoding depends on the regularizer.

3.1 Synapses as rational agents

This subsection argues that synapses are analogous to traders, operating within a neuronal market, that attempt to maximize their payout relative to their expenditures.

Prediction market traders buy and sell contingent securities. The simplest case is an *Arrow-Debreu security*, which pays out \$1 if an outcome belongs to a particular set, and \$0 otherwise [4]. For example an Arrow-Debreu security could pay \$1 if and only if candidate X wins an election. The price a trader will pay depends on her expectations about whether X will win. It turns out that the prices of securities in well-designed, liquid markets reliably aggregate traders’ diverse, private information into public estimates of the probabilities of outcomes [16].

n_j	neuron	market
$i \rightarrow j$	synapse	trader
$\mathbb{1}_i$	spike	security
$A_{\bullet}(\mathbf{w}_{ij})$	regularizer at i	cost to $i \rightarrow j$
$\mathbf{w}_{ij} \mathbb{1}_i$	weight \times spike	$\mathbb{1}_i$ s bought by $i \rightarrow j$
$\langle \mathbf{w}_j, \mathbf{x} \rangle$	total current	bundle of securities
$\langle \mathbf{w}_j, \mathbf{x} \rangle \mathbb{1}_j$	current \times spike	collective bid
$\mu_j(\mathbf{x}) \mathbf{w}_{ij} \mathbb{1}_i$	reward of $i \rightarrow j$	payout to $i \rightarrow j$

Since the neuronal scoring rule decomposes into sum $\mathcal{S}_{\bullet} = \sum_i \mathcal{S}_{\bullet}^i$, we can model not only neurons, but also synapses, as rational score-maximizing agents. Synapse $i \rightarrow j$ receives payment

$$\mathcal{S}_{\bullet}^i := \mathbb{E}_P[\langle \mathbf{w}_{ij} \mathbb{1}_i - \vartheta \rangle \cdot \mu_j(\mathbf{x}) \mathbb{1}_j - A_{\bullet}(\mathbf{w}_{ij})], \quad (8)$$

¹“Truthful reporting” is not appropriate when referring to neurons. We use the phrase “faithful encoding” instead.

where $\mathbb{1}_j$ depends on vector \mathbf{w}_j and couples the synapses.

Synapse $i \rightarrow j$ invests amount $A_{\bullet}(\mathbf{w}_{ij})$ to set its weight to \mathbf{w}_{ij} . In return, it receives quantity $\mu_j(\mathbf{x}) \mathbb{1}_i$ of security $\mathbb{1}_i$.

Like paper money, the securities $\mathbb{1}_i$ have no intrinsic worth. Instead, they are bundled into total current $\langle \mathbf{w}_j, \mathbf{x} \rangle$. If the bundle exceeds threshold ϑ then n_j spikes. That is, n_j uses the bundle to bid on an extrinsic event: the utility $\mu_j(\mathbf{x})$.

After bidding, neuron n_j receives payout $\mu_j(\mathbf{x}) \langle \mathbf{w}_j, \mathbf{x} \rangle \mathbb{1}_j$, of which it distributes an amount $\mu_j(\mathbf{x}) \mathbf{w}_{ij} \mathbb{1}_i$ to each synapse proportional to its contribution $\mathbf{w}_{ij} \mathbb{1}_i$ to the bundle. Synapses only receive payouts when they spike. Payouts can be positive or negative.

Summarizing, synapses optimize the payout, $\mu_j(\mathbf{x}) \mathbf{w}_{ij} \mathbb{1}_i$ resulting from their contribution $\mathbf{w}_{ij} \mathbb{1}_i$ to the collective bid, against their cost $A_{\bullet}(\mathbf{w}_{ij})$. The neuron’s bid $\langle \mathbf{w}_j, \mathbf{x} \rangle \mathbb{1}_j$ is thus a collective prediction of high utility by its synapses.

3.2 Proper scoring rules

The remainder of this section uses *properness* to precisely quantify how synaptic weights relate to utility expectations.

DEFINITION 2. Let $\mathcal{P}_{\mathcal{O}}$ be a set of probability distributions on states \mathcal{O} and define a **property** as a function $\Gamma : \mathcal{P}_{\mathcal{O}} \rightarrow \mathcal{H}$. Scoring rule $S : \mathcal{O} \times \mathcal{H} \rightarrow \mathbb{R}$ is **proper** [19] for property $\Gamma : \mathcal{P}_{\mathcal{O}} \rightarrow \mathcal{H}$ if for all $P \in \mathcal{P}_{\mathcal{O}}$

$$\Gamma(P) \in \operatorname{argmax}_{\mathbf{w} \in \operatorname{range}(\Gamma)} \mathbb{E}_P[S(\mathbf{x}; \mathbf{w})]. \quad (9)$$

Properness is the common-sensical requirement that the true value, $\mathbf{w} = \Gamma(P)$, is a score maximizer, $\mathbf{w} \in \operatorname{argmax} \mathbb{E}_P[S]$. In short: “you get what you think you are paying for”.

Proper scoring rules can be constructed as follows [3]. Given functions $\rho : \mathcal{O} \rightarrow \mathcal{H}$ and $F : \mathcal{H} \rightarrow \mathbb{R}$, define

$$S_F : \mathcal{O} \times \mathcal{H} \rightarrow \mathbb{R} : (\mathbf{x}; \mathbf{w}) \mapsto -D_F(\rho(\mathbf{x}), \mathbf{w}) - F(\rho(\mathbf{x}))$$

where $D_F(\mathbf{x}, \mathbf{y}) := F(\mathbf{x}) - F(\mathbf{y}) - \langle \nabla F(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle$ is the Bregman divergence. It is shown in [3] that:

PROPOSITION 2 (LINEAR PROPER SCORING RULES). *If F is convex then S_F is a proper scoring rule for linear property $\Gamma : P \mapsto \mathbb{E}_P[\rho(\mathbf{x})]$.*

3.3 Proper scoring for discretized neurons

This section adapts Proposition 2 to discretized neurons. As a warmup, we show that dropping the selectivity term from (3) yields proper scoring rules.

LEMMA 3. Let $\mathcal{T}_{\bullet}(\mathbf{x}; \mathbf{w}_j) := \mu_j(\mathbf{x}) \cdot (\langle \mathbf{w}_j, \mathbf{x} \rangle - \vartheta) - A_{\bullet}(\mathbf{w}_j)$ be scoring rules. These are proper for $\Gamma_{\bullet}^T : \mathcal{P}_{\mathcal{O}} \rightarrow \mathcal{H} = \mathbb{R}^N$,

$$\Gamma_{\bullet}^T : P \mapsto G_{\bullet} \left(\mathbb{E}_P[\mu_j(\mathbf{x}) \cdot \mathbf{x}] \right) \text{ for } \begin{cases} G_2(\mathbf{v}) = \eta \cdot \mathbf{v} \\ G_H(\mathbf{v}) = e^{\eta \cdot \mathbf{v} - 1} \\ G_1(\mathbf{v}) = \mathbb{1}_{\eta \cdot \mathbf{v} > 1}, \end{cases}$$

where $\mathbb{1}_{\eta \cdot \bullet > 1}$ an N -vector of indicator functions returning 1 when $\eta \cdot \bullet > 1$ and 0 otherwise.

PROOF. We drop ϑ since it is independent of \mathbf{w}_j . Define hypothesis space $\mathcal{H} = \mathbb{R}^N$ and map

$$\rho_{\mu} : \mathcal{O} \rightarrow \mathcal{H} : \mathbf{x} \mapsto \mu(\mathbf{x}) \cdot \mathbf{x}.$$

We consider the three cases in turn.

Observe that convex function $F_2(\mathbf{x}) = \frac{1}{2\eta} \|\mathbf{x}\|_2^2$ yields scoring rule $S_2(\mathbf{x}, \mathbf{w}_j) = \langle \mu(\mathbf{x}) \cdot \mathbf{x}, \mathbf{w}_j \rangle - \frac{1}{2\eta} \|\mathbf{w}_j\|_2^2$, which implies \mathcal{T}_2 is proper by Proposition 2.

For \mathcal{T}_H , restrict \mathcal{H} to the subset of \mathbb{R}^N where $\sum_i \exp(\frac{1}{\eta} \mathbf{w}_{ij}) = 1$ and define $\psi : \mathcal{H} \rightarrow \mathcal{H} = \mathbb{R}^N$ taking $\mathbf{w}_{ij} \mapsto \frac{\exp(\frac{1}{\eta} \mathbf{w}_{ij})}{\sum_i \exp(\frac{1}{\eta} \mathbf{w}_{ij})} = \exp(\frac{1}{\eta} \mathbf{w}_{ij})$. Convex function $F_H(\mathbf{x}) = \eta \log(\sum_{i=1}^n \exp(\frac{1}{\eta} \mathbf{x}_i))$ yields

$$\begin{aligned} S_H(\mathbf{x}, \mathbf{w}_j) &= F_H(\mathbf{w}_j) + \langle \psi(\mathbf{w}_j), \mu_j(\mathbf{x}) \cdot \mathbf{x} - \mathbf{w}_j \rangle \\ &= \langle \psi(\mathbf{w}_j), \mu_j(\mathbf{x}) \cdot \mathbf{x} \rangle - \frac{1}{\eta} \langle \psi(\mathbf{w}_j), \log \psi(\mathbf{w}_j) \rangle, \end{aligned}$$

since $F_H(\mathbf{w}_j) = 0$. By Proposition 2 it follows that \mathcal{T}_H is proper for linear property $\mathbb{E}_P[\eta \cdot \mathbf{v}]$. The result follows for $\Gamma_H^\mathcal{T}$ since $e^{\bullet-1}$ is monotonic. We use the nonlinear ‘‘exp’’ representation since it directly corresponds to synaptic weights which will be useful in Theorem 5.

Proposition 2 does not apply to \mathcal{T}_1 , so we derive properness by other means. Computing gradients obtains

$$\Delta \mathbf{w}_j \propto \mathbb{E}_P[\mu_j(\mathbf{x}) \cdot \mathbf{x}] - \frac{1}{\eta}$$

which has a stationary point when all synaptic weights are equal to the scalar η . The stationary point is unstable – a local minima rather than maxima. Synapses with $\mathbb{E}_P[\eta \cdot \mu_j(\mathbf{x}) \cdot \mathbf{1}_i] > 1$ are forced to boundary condition $\mathbf{w}_{ij} = 1$; others are forced to 0 (for simplicity we assume no expectation is precisely 1).

The range of $\Gamma_1^\mathcal{T}$ is the set of N -vectors of 0s and 1s. Any $\mathbf{w} \in \text{range}(\Gamma_1^\mathcal{T})$ differing from $\Gamma_1^\mathcal{T}(P)$ has non-zero gradient and hence a lower score, implying $\Gamma_1^\mathcal{T}$ is proper. \square

The selectivity term in (3) introduces a complication into the scoring rule: potentiating a synaptic weight may cause a neuron to stumble over a sharp change in its utility function that is hidden by the selectivity term. Although the reward function is continuous in \mathbf{w} its derivative is not: there is a kink. We bound the jump after crossing a kink via

ASSUMPTION 1 (NO NASTY SURPRISES). *If $\Delta_{ij} = \mathbb{E}_P[\mu_j(\mathbf{x}) \mathbf{1}_i \mathbf{1}_{\mathbf{w}_j}] - \partial_i A_\bullet(\mathbf{w}_j) > 0$ then there exists $\epsilon > 0$ such that*

$$\mathbb{E}_P[\mu_j(\mathbf{x}) \mathbf{1}_{\epsilon \cdot \Delta_{ij}}] > -\Delta_{ij},$$

where $\mathbf{1}_{\Delta_{ij}} := \mathbf{1}_{\mathbf{w}_j + \epsilon \cdot \Delta_{ij}} - \mathbf{1}_{\mathbf{w}_j}$.

Assumption 1 implies that sufficiently small synaptic updates, Eq. (7), always increase a neuron’s score:

LEMMA 4 (SMOOTH ASCENT). *Under Assumption 1, if $\Delta_{ij} > 0$ then there exists $\epsilon > 0$ such that*

$$\mathbb{E}_P[\mathcal{S}_\bullet(\mathbf{x}, \mathbf{w}_j + \epsilon \cdot \Delta_{ij})] > \mathbb{E}_P[\mathcal{S}_\bullet(\mathbf{x}, \mathbf{w}_j)]$$

and similarly for $\Delta_{ij} < 0$.

PROOF. Straightforward computation. \square

Informally, if high utility follows n_i and n_j co-spiking, then Assumption 1 says that the utility of *new* inputs, causing n_j to spike when synapse i increases by Δ_{ij} , is not *too* negative. If the assumption fails then the neuron will continuously potentiate and depotentiate synapse i as the gradient jumps from positive to negative. This is analogous to the behavior of a perceptron when confronted with classes that are not linearly separable.

Nasty surprises can be avoided in at least two ways. First, by designing the utility function so that it behaves well with

respect to the distribution the neuron encounters. Second, by allowing neuron n_j to modify its regularization parameter η_j . Going further, one could introduce additional degrees of freedom by associating an η_{ij} with each synapse (note the regularizers are sums over synapses) that is tweaked when a neuron detects that one of its synapses jumps back and forth. We do not pursue these ideas here.

Before proving our main result, we introduce some notation. Given \mathbf{w} , let $\mathbf{1}_{\mathbf{w}} := f_{\mathbf{w}}^{-1}(1) \subset \mathcal{O}$ and let $\mathbb{P}_{\mathcal{O}}$ denote the powerset of \mathcal{O} . Enlarge the hypothesis space to $\mathcal{H}' := \mathbb{P}_{\mathcal{O}} \times \mathcal{H}$ with embedding $\psi : \mathcal{H} \rightarrow \mathcal{H}' : \mathbf{w} \mapsto (\mathbf{1}_{\mathbf{w}}, \mathbf{w})$. Let $\mathbf{w}^* = \text{argmax}_{\mathbf{w} \in \mathcal{H}} \mathbb{E}_P[\mathcal{S}_\bullet]$.

THEOREM 5 (NEURONAL SCORING RULES ARE PROPER). *Under Assumption 1, scoring rules \mathcal{S}_\bullet are proper for property $\Gamma_\bullet : \mathcal{P}_{\mathcal{O}} \rightarrow \mathcal{H}' = \mathbb{P}_{\mathcal{O}} \times \mathcal{H}$,*

$$\Gamma_\bullet : P \mapsto \left(\mathbf{1}_{\mathbf{w}^*}, G_\bullet(\mathbb{E}_P[\mu(\mathbf{x}) \cdot \mathbf{x} \cdot \mathbf{1}_{\mathbf{w}^*}]) \right), \quad (10)$$

where $G_\bullet(\mathbf{v}) \in \{\eta \cdot \mathbf{v}, e^{\eta \cdot \mathbf{v} - 1}, \mathbf{1}_{\eta \cdot \mathbf{v} > 1}\}$ depending on the choice of regularizer.

PROOF. Property $P \mapsto \mathbf{1}_{\mathbf{w}^*}$ is proper by construction; we therefore focus on the synaptic term $G_\bullet(\cdot)$ in Eq. (10).

Computing gradients for \mathcal{S}_2 and \mathcal{S}_H yields stationary points

$$\begin{aligned} \mathbf{w}^* &= \mathbb{E}_P[\eta \cdot \mu(\mathbf{x}) \cdot \mathbf{x} \cdot \mathbf{1}_{\mathbf{w}^*}] \quad \text{and} \\ \mathbf{w}^* &\propto \exp\left(\mathbb{E}_P[\eta \cdot \mu(\mathbf{x}) \cdot \mathbf{x} \cdot \mathbf{1}_{\mathbf{w}^*}] - 1\right) \end{aligned}$$

respectively which are stable maxima under Assumption 1 by Lemma 4. As argued before, a weight vector \mathbf{w} that does not have zero-gradient cannot be a maxima, and the argument follows from Lemma 3.

Similar reasoning applies to \mathcal{S}_1 . \square

REMARK 3 (INDIRECT ELICITATION). *Eliciting properties from distributions was studied in [19], which drew a distinction between elicitable and directly elicitable properties. For example, the variance can only be elicited by a scoring rule if the mean is elicited as well. Similarly, $G_\bullet(\mathbb{E}_P[\mu(\mathbf{x}) \cdot \mathbf{1}_i \mathbf{1}_{\mathbf{w}^*}])$ cannot be elicited directly, but only in conjunction with $\mathbf{1}_{\mathbf{w}^*}$.*

Neurons only modify their synapses to incorporate rewards when spiking, Eq. (4). This encourages specialization, but also implies that individual neurons may never discover that spiking for certain inputs results in very high utility. More formally, the kink makes \mathcal{S}_\bullet non-convex, so gradient ascent is not guaranteed to find the global optimum.

Nevertheless, the relationship between synaptic weights and expected utilities in Theorem 5 still holds:

COROLLARY 6 (SYNAPTIC CODE). *Let $\tilde{\mathbf{w}}$ be the (in general local) maximum of \mathcal{S}_\bullet obtained by gradient ascent with Eq. (7). If Assumption 1 holds then $\tilde{\mathbf{w}}$ satisfies*

$$\tilde{\mathbf{w}} = G_\bullet\left(\mathbb{E}_P[\mu(\mathbf{x}) \cdot \mathbf{x} \cdot \mathbf{1}_{\tilde{\mathbf{w}}}] \right). \quad (11)$$

Note that Eq. (11) is not in closed form since $\tilde{\mathbf{w}}$ appears on both the left- and right-hand sides.

PROOF. Since local maxima are stationary points, the proof follows the same argument as Theorem 5. \square

A discretized neuron n_j thus faithfully encodes two properties of its input distribution. First, its spikes encode a set of inputs for which spiking is locally optimal. Second, its synaptic weights encode the *expected utility per synapse* when n_i and n_j co-spike.

REMARK 4 (NEURAL CODE). *Corollary 6 provides an interesting interpretation of the meaning of spikes. A neuron spikes if the dot product $\langle \mathbf{w}_j, \mathbf{x} \rangle$ is above threshold ϑ . That is, neuron n_j 's spike means that the current system state \mathbf{x} is significant (above threshold) when evaluated against the utility expectations \mathbf{w}_j that were previously encoded into n_j 's structure.*

Similarly to how stock price movements encode information about which sectors of an economy are expected to yield high profits in the near future; spikes and synaptic weights encode expectations about future rewards.

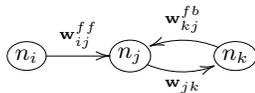
4. CORTICAL PREDICTION MARKETS

This section investigates how neurons can estimate their usefulness to downstream neurons, and so allocate their resources such that the benefit to *other* neurons is maximized. In short, we introduce a utility function that incentivizes neurons to optimize their usefulness to other neurons.

4.1 Backpropagation: errors or incentives?

To provide context, we recall related work on incorporating spikes into a reward signal. Neuromodulators provide a primary reward system. However, neurons whose actions do not directly result in pleasure or pain may require more indirect incentives. In machine learning, multilayer networks are often trained by backpropagating errors [24]. However, backpropagation (BP) is biologically implausible – it requires pathways for backpropagating errors which have not been observed in cortex [23].

As an alternative, [23] proposed attention-gated reinforcement learning (AGREL), which uses feedback spikes as attention signals to modulate learning. AGREL abstracts two features of feedback (NMDA) connections in cortex: (i) they prolong, but do not initiate, spiking activity and (ii) they have a multiplicative effect on synaptic updates.



AGREL updates feedforward weights according to

$$\Delta \mathbf{w}_{ij} \propto \mathbf{w}_{kj}^{fb} \mathbf{x}_k \cdot \mathbf{x}_i \mathbf{x}_j \cdot (1 - \mathbf{x}_j) \cdot f(\delta), \quad (12)$$

where $f(\delta)$ is a global reward signal. Here, neurons have real-valued outputs and $(1 - \mathbf{x}_j)$ is a regularizer that prevents n_j from overactivating. The main result of [23] is that average weight changes $\mathbb{E}[\Delta \mathbf{w}_{ij}]$ under (12) coincide with BP. AGREL thus provides a biologically plausible substitute for BP.

Inspired by AGREL, we introduce a 4th utility function:

Example U4. (Feedback). Identify disjoint upstream and downstream populations, \mathbf{x}^{ff} and \mathbf{x}^{fb} respectively, and define \mathcal{H}^{ff} and \mathcal{H}^{fb} by clamping weights not in the respective populations to zero using Eq (1). Define the feedback utility as $\mu_j^{fb}(\mathbf{x}) := \langle \mathbf{w}_j^{fb}, \mathbf{x} \rangle$ for $\mathbf{w}_j^{fb} \in \mathcal{H}^{fb}$.

A neuron with feedback utility maximizes

$$\mathbb{E}_P \left[\langle \mathbf{w}_j^{fb}, \mathbf{x} \rangle (\langle \mathbf{w}_j^{ff}, \mathbf{x} \rangle - \vartheta) \mathbb{1}_j - A_\bullet(\mathbf{w}_j) \right] \quad (13)$$

and so aligns its feedforward $\langle \mathbf{w}_j^{ff}, \mathbf{x} \rangle$ and feedback $\langle \mathbf{w}_j^{fb}, \mathbf{x} \rangle$ current whenever the neuron itself spikes.

Computing gradient ascent on scoring rule $\mathcal{S}_\bullet^{fb}(\mathbf{w}; \mathbf{x}) = \langle \mathbf{w}^{fb}, \mathbf{x} \rangle \cdot (\langle \mathbf{w}^{ff}, \mathbf{x} \rangle - \vartheta) \cdot \mathbb{1}_j - A_\bullet(\mathbf{w})$ obtains

$$\Delta \mathbf{w}_{ij}^{ff} \propto \langle \mathbf{w}_j^{fb}, \mathbf{x} \rangle \cdot \mathbb{1}_{ij} - \partial_i A_\bullet(\mathbf{w}), \quad (14)$$

which differs from AGREL (12) by using $\partial_i A_\bullet$ as regularizer instead of $(1 - \mathbf{x}_j)$ and extending feedback from a single neuron, $\mathbf{w}_{kj}^{fb} \mathbf{x}_k$, to many neurons, $\langle \mathbf{w}_j^{fb}, \mathbf{x} \rangle$. We also drop the global reward signal $f(\delta)$ since we are interested in the pure backpropagation case; it can easily be reinstated.

Note that the utility function μ^{fb} is itself plastic. Neuron n_j not only modifies feedforward weights to maximize its score, it also modifies feedback weights to increase the maximum achievable score:

$$\Delta \mathbf{w}_{kj}^{fb} \propto (\langle \mathbf{w}_j^{ff}, \mathbf{x} \rangle - \vartheta) \cdot \mathbb{1}_{jk} - \partial_k A_\bullet(\mathbf{w}). \quad (15)$$

4.2 Estimating usefulness with feedback

As suggested in the introduction, one way to encourage collaboration is for each neuron to estimate its usefulness to the rest of the system and optimize that estimate. By Corollary 6, a faithful measure of the usefulness of n_j 's output to the rest of cortex is the sum of active downstream synaptic weights:

DEFINITION 3. *The usefulness $\mathcal{V}_j(\mathbf{x})$ of a spike by n_j is the sum of the synaptic weights of downstream neurons that co-spike with n_j :*

$$\mathcal{V}_j(\mathbf{x}) := \sum_{\{k \mid j \rightarrow k\}} \mathbf{w}_{jk} \mathbb{1}_{jk} = \langle \mathbf{w}_{j\bullet}, \mathbf{x} \rangle \mathbb{1}_j. \quad (16)$$

Intuitively, $\mathcal{V}_j(\mathbf{x})$ is the total utility that spiking downstream neurons expect after n_j spikes.

Neurons cannot compute their usefulness directly, since the utilities of downstream neurons are private. They must therefore make do with publicly available data: *spikes by other neurons*. We therefore propose that neurons use feedback, which they can actually compute, as a proxy for usefulness, which would be ideal.

As a consequence of Corollary 6, we quantify how closely feedback-utility approximates usefulness (16):

COROLLARY 7 (ESTIMATING USEFULNESS WITH FEEDBACK). *Neuron n_j equipped with utility function $\mu^{fb}(\mathbf{x})$ approximately maximizes its usefulness $\mathcal{V}(n_j)$ to the rest of cortex, where the failure of the approximation is*

$$\sum_k \mathbb{1}_{jk} \left[\underbrace{G_\bullet \left(\mathbb{E}[\mu_k(\mathbf{x}) \mathbb{1}_{jk}] \right)}_{\mathbf{w}_{jk}} - \underbrace{G_\bullet \left(\mathbb{E}[\langle \mathbf{w}_j^{fb}, \mathbf{x} \rangle \mathbb{1}_{jk}] \right)}_{\mathbf{w}_{kj}^{fb}} \right].$$

Thus, the quality of $\langle \mathbf{w}_j^{fb}, \mathbf{x} \rangle$ as an estimate of $\mathcal{V}_j(\mathbf{x})$ depends on how closely n_j 's feedforward inputs $\langle \mathbf{w}_j^{ff}, \mathbf{x} \rangle$ approximate the sum of the downstream utilities $\mu_k(\mathbf{x})$.

PROOF. The usefulness and utility of n_j are

$$\mathcal{V}_j(\mathbf{x}) = \langle \mathbf{w}_{j\bullet}, \mathbf{x} \rangle \mathbb{1}_j \quad \text{and} \quad \mu_j^{fb}(\mathbf{x}) = \langle \mathbf{w}_{j\bullet}^{fb}, \mathbf{x} \rangle \quad \text{respectively.}$$

The utility μ_j^{fb} is multiplied by $\mathbb{1}_j$ when it is used in scoring rules, so the difference comes down to the weights. Corollary 6 implies the optimal feedforward weights are

$$\mathbf{w}_{jk} = G_\bullet \left(\mathbb{E}_P [\mu_k(\mathbf{x}) \cdot \mathbb{1}_{jk}] \right)$$

so that the usefulness of n_j is

$$\mathcal{V}_j(\mathbf{x}) = \sum_{\{k \mid j \rightarrow k\}} \mathbb{1}_{jk} \cdot G_\bullet \left(\mathbb{E}_P [\mu_k(\mathbf{x}) \mathbb{1}_{jk}] \right).$$

Again by Corollary 6, the properness of the scoring function implies the optimal weights for $k \rightarrow j$ satisfy

$$\mathbf{w}_{kj}^{fb} = G_\bullet \left(\mathbb{E}_P [\langle \mathbf{w}^{ff}, \mathbf{x} \rangle \mathbb{1}_{jk}] \right)$$

and we are done. \square

Experiments in §5 demonstrate that \mathbf{w}_{kj}^{fb} is a good proxy for \mathbf{w}_{jk} in some interesting cases.

A detailed analysis of the relationship between approximations $\mathbf{w}_{jk} \approx \mathbf{w}_{kj}^{fb}$, distribution $P(\mathbf{x})$, and utility functions $\mu_k(\mathbf{x})$ is beyond the scope of this paper.

4.3 Neurons as rational agents

Section §3.1 suggested neurons are analogous to markets in which synapses trade. This section presents a second analogy, where cortex forms a market in which neurons trade.

Recall that neurons are rational agents that optimize their expected reward balanced against a cost term, Theorem 1:

$$\mathbf{w}_j^* := \operatorname{argmax}_{\mathbf{w}_j \in \mathcal{H}} \mathbb{E}_P \left[(\langle \mathbf{w}_j, \mathbf{x} \rangle - \vartheta) \cdot \mu_j(\mathbf{x}) \mathbb{1}_j - A_\bullet(\mathbf{w}_j) \right].$$

n_j	neuron	trader
$\langle \mathbf{w}_{\bullet j}^{ff}, \mathbf{x} \rangle \mathbb{1}_j$	ff current \times spike	purchases by n_j
$\langle \mathbf{w}_{j\bullet}, \mathbf{x} \rangle \mathbb{1}_j$	usefulness $\mathcal{V}_j(\mathbf{x})$ of n_j	use made of n_j
$\langle \mathbf{w}_{\bullet j}^{fb}, \mathbf{x} \rangle \mathbb{1}_j$	fb current \times spike	recorded usage = payment to n_j

The key idea is that each neuron should optimize its usefulness to the rest of the brain. Building on Corollary 6, usefulness is defined as $\mathcal{V}_j(\mathbf{x}) = \sum_{j \rightarrow k} \mathbf{w}_{jk} \mathbb{1}_{jk}$. That is, the quantity of $\mathbb{1}_j$ used by downstream neurons in their internal markets. Unfortunately, n_j does not have access to this number. Similarly to how musicians are paid for actual sales rather than downloads of their music, neurons need to record when their outputs are used. They therefore use feedback to compute $\langle \mathbf{w}_{\bullet j}^{fb}, \mathbf{x} \rangle \mathbb{1}_j$, which acts as a proxy² for $\langle \mathbf{w}_{\bullet j}^{fb}, \mathbf{x} \rangle \mathbb{1}_j$.

Intuitively, n_j simultaneously sets its feedback connections on the downstream traders that most frequently purchase its spikes, and sets its feedforward connections on the upstream traders that sell the most useful spikes.

The result is a mesh of intertwining neuronal chains – optimized for usefulness at every link by the invisible hand of the cortical market – that connects sensory inputs to motor actions.

5. EXPERIMENTS

We investigate the empirical performance of discretized neurons. The experiments are designed to show that: (i) the ideas above can be implemented with minimal modifications; (ii) synaptic weights encode environmental statistics and rewards; (iii) feedback improves performance; and (iv) feedback reliably estimates a neuron’s usefulness.

We have therefore constructed networks, inspired by [22], that learn tasks designed so that the embedding of expected utilities into synaptic weights is easy to visualize.

²Recorded usage could over- or under- estimate true usage. Section §5 shows that it is a good guide in practice.

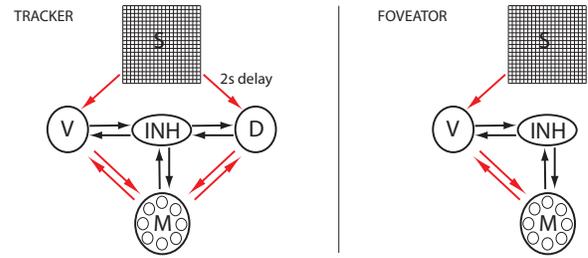


Figure 2: Foveator and Tracker architectures. Arrows are initialized randomly. Red arrows are plastic; black are fixed.

Our goal is not to compete with the state of the art. Rather, our aim is to introduce mechanism design techniques into the analysis and construction of networks. A pressing open question is whether more sophisticated networks, such as those developed by the deep learning community, can be understood or improved via mechanism design.

Network architectures.

The **tracker network**, Fig. 2 left, has a sensory grid S of 20×20 neurons, intermediate layers V and D with 100 neurons each, motor layer M , and 100 randomly connected inhibitory neurons INH . Signals from S to D are *delayed* so that V and D receive different temporal snapshots of S . Synapses are plastic except those to or from inhibitory neurons. M is divided into 8 areas of 10 neurons each. Actuators engage when they receive more than 10 spikes. The network is initialized randomly.

The tracker network tracks targets traveling along an edge of the visual field. Motor areas are rewarded ($\mu_j = +1$) or punished ($\mu_j = -1$) according to whether or not the action correctly anticipates where the target is headed and from which direction (4×2 possibilities). Note the motor layer receives neuromodulatory signals whereas the intermediary layers do not and learn from feedback.

The **foveator network**, Fig. 2 right, drops D and has fewer inhibitory neurons.

The foveation task is to move the fovea (center of the retina) onto an object appearing on the edge of the visual field. Each motor area controls an actuator that moves the fovea in a compass direction (N, NW, W, etc). After a movement, the corresponding area is rewarded ($\mu_j = 1$) if the object is closer to the center and punished (-1) otherwise.

We tweak the discretized neuron in §2 to make the dynamics closer to continuous time models of cortical neurons.

First, we introduce a voltage term V , which provides neurons with a steadily decaying “memory” of previously received spikes. Neurons spike when $V > \vartheta$, after which $V \leftarrow 0$. When the neuron does not spike, V is updated according to

$$V \leftarrow V + \langle \mathbf{w}, \mathbf{x} \rangle - \delta.$$

Neurons maintain an exponentially decaying trace reflecting recent output spikes:

$$\text{trace}_j \leftarrow 0.95 \cdot (\text{trace}_j + 0.4 \cdot \mathbb{1}_j)$$

Neurons in subsystems V and D update their feedforward synapses according to

$$\Delta \mathbf{w}_{ij} \propto \langle \mathbf{w}^{fb}, \mathbf{x} \rangle \cdot \mathbb{1}_i \cdot \text{trace}_j$$

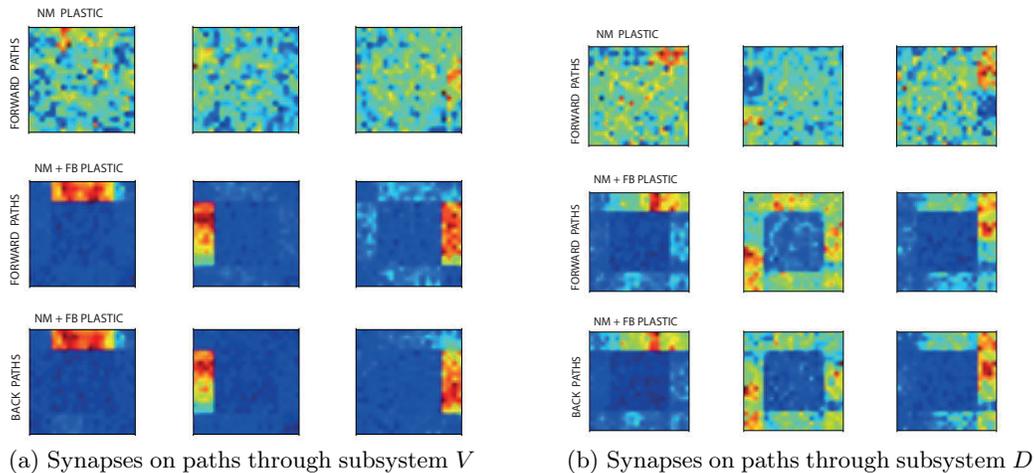


Figure 1: Average synaptic weights in Tracker network.

and similarly for feedback. Thus, trace_j is substituted for $\mathbb{1}_j$ to temporally smooth out learning. Neurons in subsystem M update their synapses according to $\Delta \mathbf{w}_{ij} \propto \mu_j(\mathbf{x}) \cdot \sum_{t'=t-m}^t \mathbb{1}_{ij}$ where the sum is over tics since the last neuromodulatory signal, similar to the trace implemented in V and D .

Finally, we tweak the regularization. Instead of continually regularizing by $A_1(\mathbf{w})$, we regularize at discrete intervals, analogous to a hypothesized role of sleep [5]. Regularization consists of setting the K strongest synapses to 1, and pruning the rest (i.e. setting their weights to 0). The number K is fixed within each layer, but varies across layers.

Visualizing synaptic strengths.

By Corollary 7, the quality of a neuron’s usefulness-estimate can be computed and visualized by comparing average feedforward and feedback weights.

Weights are visualized in Fig 1 and 3 as follows. For each area in M , we average over all *feedforward* paths $S \rightarrow V \rightarrow M$ and *feedback* paths $S \rightarrow V \leftarrow M$ respectively, and similarly for D . To save space, 3 out of 8 areas in M are plotted. Plots are averaged over 20 runs. Blue denotes low values; red denotes high.

Results.

(i) The tasks are easy and the networks rapidly (within a few thousand tics) achieve 98% and 95% accuracy. The tracker outperforms the foveator, possibly because the foveator modifies its environment by actively moving the center of the retina, whereas the tracker does not.

The delay line is essential to tracker performance: if the delay is set to zero then the network performs little better than chance, and the structure of the environment is not learned at all.

Tracker	% correct	# correct
only M plastic	95	243
all plastic	98	672
Foveator		
only M plastic	93	80
all plastic	95	207

(ii) The middle rows of Fig 1 and 3 show how rewards and

environmental statistics are incorporated into the networks’ feedforward structure.

For the tracker network, the V -area synapses learn trajectories; whereas the D -area learns the *starting points* of trajectories. The combination of instantaneous lines in V (which learn directions) and delay lines in D (which learn starting points) thus allows the network to implicitly compute derivatives and thereby determine directions of travel.

For the foveator, it is easy to read off the correspondence between the NE, N, and NW movements of the actuators and the locations of objects driving the movements.

(iii) Shutting off feedback plasticity (top rows of Fig 1 and 3) slightly worsens performance, from 98% to 95% for the tracker and from 95% to 93% for the foveator. However, it dramatically worsens the “reaction times” of the networks, quantified as the number of times the actuators correctly engage per 1000 tics.

Indeed, looking at synaptic weights *without* feedback plasticity, top rows of the figures, we find that the structure of the rewards and environment is barely visible.

(iv) Finally, when feedback plasticity is turned on, average synaptic weights over feedforward paths $T \rightarrow V/D \rightarrow M$ (middle rows) and feedback paths $T \rightarrow V/D \leftarrow M$ (bottom rows) are almost identical, demonstrating that neurons in V and D accurately estimate their usefulness to downstream M neurons using feedback.

6. CONCLUSION

This paper applied tools from mechanism design to investigate a simple model of cortical neurons. The main result is that, under a technical assumption, neurons faithfully encode expected utilities into their synaptic weights. If the result can be extended to more realistic models, then it will provide a powerful new approach to understanding the relationship between cortical structure and function.

There is reason to be optimistic: continuous time models require exponential discount factors analogous to interest rates – which are well-understood in mechanism design.

An important corollary is a novel interpretation of the role of spiking feedback in cortex: neurons can use feedback spikes to estimate their usefulness to the rest of cortex, and then learn to maximize that estimate. Going further, it may

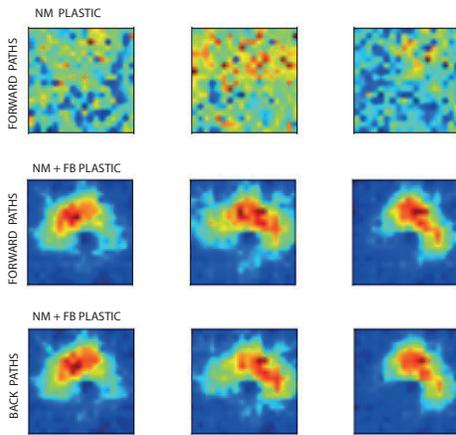


Figure 3: Average weights in Foveator network.

be that retroaxonal signals, backpropagated through axons, may play an important role in helping neurons estimate their usefulness [17].

We have used the simplest possible scoring rules, derived from standard models, to provide proof of principle. It will be interesting to explore more realistic models taken from the neuroscience literature, and also more powerful models such as those developed for deep neural networks.

Finally, although the flow of ideas in this paper is one-sided – from mechanism design to neuronal models – future work should be more symmetric. The cortex aggregates information far more effectively than the auctions and online markets studied in game theory. This suggests there are powerful design principles waiting to be uncovered.

7. ACKNOWLEDGEMENTS

I am grateful to Hastagiri Vanchinathan for encouraging me to look into mechanism design.

8. REFERENCES

- [1] J. Abernethy, Y. Chen, and J. W. Vaughan. Efficient Market Making via Convex Optimization, and a Connection to Online Learning. *ACM TEAC*, 1, 2013.
- [2] J. Abernethy and R. Frongillo. A Collaborative Mechanism for Crowdsourcing Prediction Problems. In *NIPS*, 2011.
- [3] J. Abernethy and R. Frongillo. A Characterization of Scoring Rules for Linear Properties. In *COLT*, 2012.
- [4] K. J. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.
- [5] D. Balduzzi and M. Besserve. Towards a learning-theoretic analysis of spike-timing dependent plasticity. In *NIPS*, 2012.
- [6] D. Balduzzi, P. A. Ortega, and M. Besserve. Metabolic cost as an organizing principle for cooperative learning. *Advances in Complex Systems*, 16(2/3), 2013.
- [7] D. Balduzzi and G. Tononi. What can neurons do for their brain? Communicate selectivity with spikes. *Theory in Biosciences*, 132(1):27–39, 2013.
- [8] J. Berg, R. Forsythe, F. Nelson, and T. Rietz. Results from a dozen years of election futures markets research. In C. Plott and V. Smith, editors, *Handbook of Experimental Economics Results*. 2001.
- [9] M. Boerlin and S. Denève. Spike-based population coding and working memory. *PLoS Comput Biol*, 7(2):e1001080, Feb 2011.
- [10] G. Brier. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1):1–3, 1950.
- [11] Y. Chen and D. Pennock. A utility framework for bounded-loss market makers. In *UAI*, 2007.
- [12] Y. Chen and J. Wortman Vaughn. A new understanding of prediction markets via no-regret learning. In *ACM EC*, 2010.
- [13] S. Fusi and L. Abbott. Limits on the memory storage capacity of bounded synapses. *Nature Neuroscience*, 10(4):485–493, 2007.
- [14] W. Gerstner and W. Kistler. *Spiking Neuron Models*. Cambridge University Press, 2002.
- [15] R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15, 2007.
- [16] R. Hanson, R. Oprea, and D. Porter. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior and Organization*, 4:449–459, 2006.
- [17] K. D. Harris. Stability of the fittest: organizing learning through retroaxonal signals. *Trends Neurosci*, 31:130–136, 2008.
- [18] A. Hasenstaub, S. Otte, E. Callaway, and T. J. Sejnowski. Metabolic cost as a unifying principle governing neuronal biophysics. *Proc Natl Acad Sci U S A*, 107(27):12329–34, Jul 2010.
- [19] N. Lambert, D. Pennock, and Y. Shoham. Eliciting properties of probability distributions. In *ACM EC*, 2008.
- [20] N. Lay and A. Barbu. Supervised aggregation of classifiers using artificial prediction markets. In *ICML*, 2010.
- [21] J. Ledyard, R. Hanson, and T. Ishikida. An experimental test of combinatorial information markets. *Journal of Economic Behavior and Organization*, 69:182–189, 2009.
- [22] A. Nere, U. Olcese, D. Balduzzi, and G. Tononi. A neuromorphic architecture for object recognition and motion anticipation using burst-STDP. *PLoS One*, 7(5):e36958, 2012.
- [23] P. R. Roelfsema and A. van Ooyen. Attention-gated reinforcement learning of internal representations for classification. *Neural Comput*, 17(10):2176–2214, 2005.
- [24] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *Nature*, 323:533–536, 1986.
- [25] S. Shalev-Shwartz and Y. Singer. A primal-dual perspective of online learning algorithms. *Machine Learning*, 69(2-3):115–142, 2007.
- [26] S. Song, K. D. Miller, and L. F. Abbott. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci*, 3(9), 2000.
- [27] G. B. Stanley. Reading and writing the neural code. *Nat Neurosci*, 16(3), 2013.
- [28] A. Storkey. Machine Learning Markets. In *AISTATS*, 2011.