Toward an analog neural substrate for production systems

Patrick Simen (psimen@princeton.edu), Marieke Van Vugt (mkvan@princeton.edu)
Fuat Balci (fbalci@princeton.edu)
Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544 USA

David Freestone (David_Freestone@brown.edu)
Department of Psychology, Brown University, Providence, RI 02912 USA

Thad Polk (tpolk@umich.edu)
Department of Psychology, University of Michigan, Ann Arbor, MI 48432 USA

Abstract

Symbolic, rule-based systems seem essential for modeling high-level cognition. Subsymbolic dynamical systems, in contrast, seem essential for modeling low-level perception and action, and can be mapped more easily onto the brain. Here we review existing work showing that critical features of symbolic production systems can be implemented in a subsymbolic, dynamical systems substrate, and that optimal tuning of connections between that substrate's analog circuit elements accounts for fundamental laws of behavior in psychology. We then show that emergent properties of these elements are reflected in behavioral and electrophysiological data, lending support to a theory about the physical substructure of productions. The theory states that: 1) productions are defined by connection strengths between circuit elements; 2) conflict resolution among competing productions is equivalent to optimal hypothesis testing; 3) sequential process timing is parallel and distributed; 4) memory allocation and representational binding are controlled by competing relaxation oscillators.

Keywords: Production system; neural network; diffusion model; random walk; reinforcement learning.

A subatomic structure for productions

Production systems underlie the most successful theories of high-level cognition, exemplified by such capabilities as planning, problem-solving, reasoning and language. Productions — if-then rules that test the contents of a working memory and trigger actions or changes to working memory as a result — have accordingly been characterized as the 'atomic components of thought' (Anderson & Lebiere, 1998). The implication is that the complex chemistry of mental life arises from, and can more easily be understood in terms of, the interactions of these simple atoms. To make the most of this analogy, however, requires a biologically plausible theory about the subatomic structure that defines these interactions. Here we propose a subatomic theory in which productions arise from the behavior of 'elementary particles' — leaky integrators, or classic neural network units — whose interactions with each other are defined by connection strengths and structured network topologies.

Any computational theory of cognition faces several challenges: How well does it conform to known laws of behavior and classic patterns of brain activity? How well does data conform to new predictions entailed by it? And how much functionality does it give you (e.g., is it computationally complete)? Here we progressively build up a design for a neural network structure that emulates the most important features

of production systems. We start with a critical core for individual productions, and then add on control mechanisms that adapt the core's behavior in order to maximize a reward function. We will attempt to show how each addition accounts for known laws, entails new (in some cases, successfully tested) predictions, and moves the resulting architecture toward full, production-system functionality. The result falls short of enabling the automatic translation of arbitrary production system programs into equivalent neural networks, but it suggests that such translations will be possible for a constrained set of such programs (and that the constraints thus identified may be of theoretical importance).

For the core, we review a specific, structured neural circuit with heuristically reward-maximizing connections that has previously been proposed as an implementation of productions (Polk, Simen, Lewis, & Freedman, 2002; Simen & Polk, 2010). After outlining the remaining mechanisms underlying key features of a neural production system architecture, we review separately published results showing the conformance of its behavioral predictions to the matching law of operant conditioning, to the logistic/softmax choice function used in reinforcement learning, and to recent, tested theories of optimal perceptual decision making. We also review new evidence supporting its predictions regarding the lateralized readiness potential (LRP) that is observed in human electroencephalography (EEG).

To the core production implementation, we add a simple timing mechanism (allowing controlled sequential processing), and we outline a proof that it conforms to the law of scalar invariance in interval timing (Gibbon, 1977). We show that this mechanism predicts behavior observed in the differential reinforcement of low rates of responding (DRL) task.

We conclude the addition of mechanisms by outlining a potential solution to two major challenges facing a connectionist production system architecture: one is the need for a flexible memory management system; the other is the variable binding problem. This problem afflicts any system in which the semantics of a representation depend only on what is connected to what, so that the components of different representations must be shared. Our proposed solution involves relaxation oscillators with tunable frequencies and duty cycles. These enable the recruitment of memory resources through fast Hebbian learning by tagging and reserving allocated net-

work units. The result is the sort of oscillatory activity that is invariably observed in invasive electrode recording and scalp EEG.

The elementary particles

The basic building block we will use is a stochastic neural network unit. We begin its description by considering it as a deterministic system. At each moment, it computes a weighted sum of its current inputs, then computes an exponentially decaying average of recent weighted sums, and finally amplifies the result by a gain function that is approximately linear (but which saturates at very low and very high input levels). This quantity is broadcast to other units, over connections whose strengths determine their relative contribution in those units' weighted sum computations. Formally, the output of the ith unit is V_i, the leaky integral of summed input is x_i, and the dynamics are defined as follows:

Figure 1: **a, b:** A neural network unit's rate of activation change (dV/dt) as a function of input I and output V for units with fixed I and balanced (**a**) or strong (**b**) excitatory, recurrent connections. Equilibrium curves are solid; velocities dV/dt are indicated by arrows and shading (light > 0, dark < 0). **c:** 'Catastrophe manifold' formed by the equilibrium curves of Eq. 3 as the self-excitatory, recurrent weight strength w_{ii} ranges from 0 to 2. Three network symbols are also illustrated. Sigmoid: weak self-excitation, leaky integration ($w_{ii} < 1$ for a unit with $\lambda = 4$). Rounded step-function: balanced self-excitation, perfect integration ($w_{ii} = 1$). S symbol: strong self-excitation, hysteresis and bistable switching (or 'latch') behavior ($w_{ii} > 1$). **d:** A latch based on hysteresis. States above the dashed curve converge to the upper solid curve; states below converge to the lower solid curve. This latch can store a 1 (upper gray region) or a 0 (lower gray region) as long as input is held between A and B. Bit-flipping during constant I is least likely when $I = (A+B)/2$.

$$I_i = \sum_{j=1}^{n} w_{ij} \cdot V_j, \tag{1}$$

$$\tau \cdot \frac{dx_i}{dt} = -x_i + I_i, \tag{2}$$

$$\text{and } V_i(t) = f(x_i(t)) = [1 + \exp(-\lambda \cdot (x_i - \beta))]^{-1}. \tag{3}$$

Parameters λ and β determine the steepness and position of the sigmoidal activation function f, and τ determines the decay rate of exponential averaging (large τ gives slow decay).

In addition to deterministic dynamics, we assume that noise enters the system from units that have direct sensory inputs, and also from the connections between units themselves. To model these assumptions, we use stochastic differential equations, in which we represent white noise with a useful abuse of notation as $\eta \equiv dW/dt$ (multiplication by dt then gives the standard notation dW in our equations; cf. Gardiner, 2004). This quantity represents the time-derivative of a Brownian motion, or Wiener process, $W(t)$.[1] The standard deviation of η is 1, but can be changed to any value c by multiplying by c. Here, we multiply η by the square root of the weighted input, an assumption which is consistent with an even more microscopic level of neural modeling: we assume that spiking neurons are Poisson processes, and that leaky integrators model their population-level behavior. The variance of sums of these independent processes is the sum of their variances. Thus, if we consider increases in a given weight w_{ij} to be equivalent to the addition of independent Poisson processes (because of the addition of noisy synaptic connections), we get a noise standard deviation equal to the square root of net input. Formally, then, the full, stochastic unit description is as follows:

$$\tau \cdot \frac{dx_i}{dt} = -x_i + \left(I_i + c_i\sqrt{I_i} \cdot \eta\right)$$
$$\Rightarrow \tau \cdot dx_i = (-x_i + I_i)\, dt + c_i\sqrt{I_i}\, dW_i$$
$$\Rightarrow \tau \cdot dV_i \approx (-x_i + f(I_i))\, dt + c_i\sqrt{I_i}\, dW_i \tag{4}$$

(See Simen and Polk (2010) for justification of the last approximation, which moves the noise term outside the nonlinear function f.)

This system can be numerically simulated on a computer (and perhaps be more easily understood) as a discrete-time difference equation (Gardiner, 2004):

$$\tau \cdot V_i(t+\Delta t) \approx V_i(t) + (-x_i + f(I_i))\, \Delta t + c_i\sqrt{I_i}. \tag{5}$$

It is now critical for our purposes to consider the effects of recurrent excitation of a unit by itself ($w_{ii} > 0$). The strength of this self-excitation determines which of three, qualitatively distinct types of behavior a unit exhibits (Simen & Polk, 2010). For $w_{ii} < 1$, the system acts like a leaky integrator;

---

[1] W in fact is non-differentiable, but it is the limit of a sequence of slightly smoother, differentiable noise processes, so it can be used without danger.
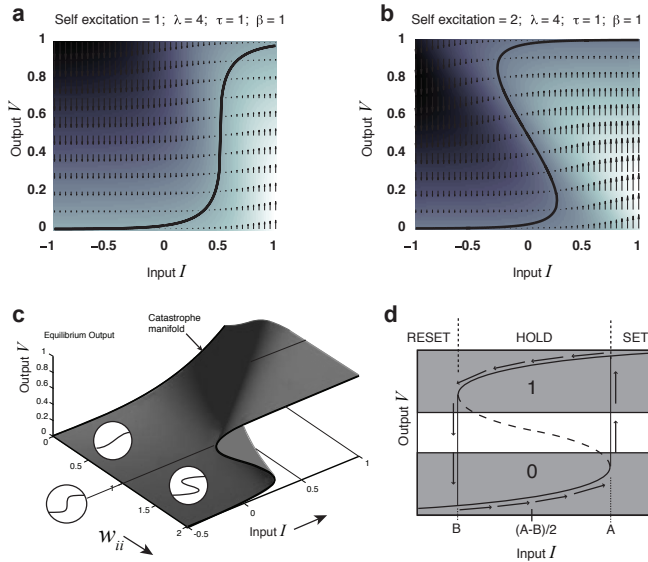
as w_{ii} grows, the leak is reduced. When the self-excitation exactly balances the leak ($w_{ii} = 1$), the unit acts like a perfect integrator (until it saturates). For $w_{ii} > 1$, the system is unstable and is forced upward against the upper ceiling on its activation (1), or downward toward its lower floor (0); thus it acts like a binary switch. Furthermore, such a unit displays hysteresis, so that it can both trigger abrupt changes and also store a bit. Fig. 1 shows the dynamics of such a unit.

In general, leaky integration (weak self-excitation) is useful because it low-pass filters its input, thereby removing much of the high frequency noise contributed by connections and by the environment. Perfect integration (balanced self-excitation) is needed for optimal hypothesis testing (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006). Bistability (strong self-excitation) is needed for triggering subsequent steps of sequential processes and for maintaining the current state of working memory. The behavioral and electrophysiological data we consider bears on the predictions made by these bistable units and the integrators that feed into them.

Neural productions, timers and oscillators

Fig. 2 shows the basic building blocks of the proposed architecture; in the remainder of the paper, we explain how each block functions, and assess how well each accords with known laws and new empirical data. The left column shows the 3 unit types (**a**,**b**,**c**). Simen and Polk (2010) detail how a complete set of logic operations (AND, OR, NOT) can be built from the bistable units in **c** by parameterizing their input strengths. Panel **d** shows a simple if-then rule structure: the leaky integrator filters noise from its inputs, and if the sum exceeds a critical level, the bistable unit switches from (approximately) 0 to (approximately) 1. This is analogous to the process of 'matching' the contents of working memory (which can be made to depend on arbitrarily many symbolic preconditions using a cascade of logic gates). The degree of match may be an analog quantity, and whether this is sufficient to cause a bit flip in the output unit determines whether the production will 'fire'. Furthermore, the weights on inputs to the if-stage may also encode preferences between productions that have an equal degree of supporting evidence.

If more than one production matches, however, there may be conflict between them. At least at the motor output stage (e.g., SOAR's 'operators'), such conflict must be resolved. Here we consider conflict resolution as a process of competition between matching productions (Fig. 2 **e**), with the outcome biased toward selection of the production with the strongest amount of preference-weighted evidence. Since noise is everywhere, this reduces to a well-defined hypothesis testing problem, for which simple, near-optimal algorithms exist. These algorithms — sequential probability ratio tests (SPRTs) — can be parameterized to maximize expected utility in the case of two-alternative choices (Bogacz et al., 2006), and can approximately maximize utility for a greater number of competing alternatives (McMillen & Holmes, 2006). For a difficult decision, the process of deciding via lateral inhi-
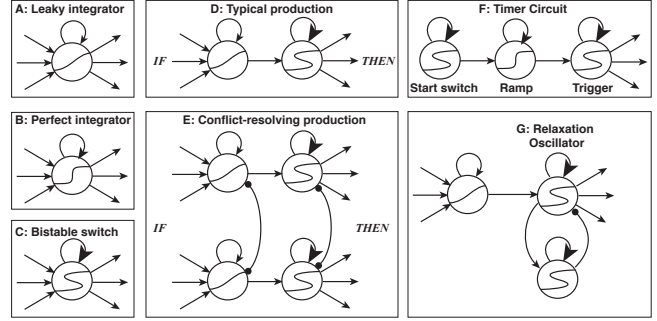

Figure 2: Basic building blocks. Arrowheads indicate excitation, circleheads inhibition. **a, b, c:** Elementary particles; arrows: excitatory inputs. **d:** Production topology. **e:** Conflict resolution via lateral inhibition (circles: inhibition); inhibition between switches is optional. **f:** Interval timer. **g:** Relaxation oscillator added to production output unit.

bition (a form of attractor dynamics) can be parameterized to implement an SPRT. This requires only that the lateral inhibitory strengths between input units equal -1. An example of these dynamics is shown in Fig. 3. Thus, the firing of a single production is equivalent to a statistical hypothesis test.
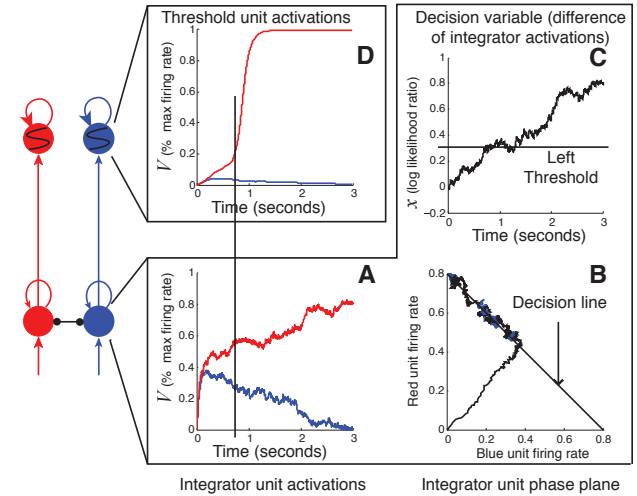

Figure 3: Hypothesis testing via lateral inhibition. The 2D system in the bottom layer reduces to a single dimension, along which a random walk to threshold occurs (implemented by attractor dynamics in the top layer).

A critical question facing the proposed architecture, however, is whether the timing of these firings can be coordinated and sequentialized without reference to a central system clock. Our problem is the same as that facing digital circuit designers, who have long relied on a central clock and synchronous updating to preclude critical race conditions and other signal timing hazards. Our solution is to use these pro-

duction implementations to form processing bottlenecks, and to use handshake completion signals between computing elements for asynchronous, distributed timing control (Simen & Polk, 2010). The most difficult question is whether we can implement productions of the form: If A, Then B and Not A. Naively wiring up a system to implement such a production can cause critical race conditions or metastability.

Our solution derives from the hysteresis properties of our bistable units. Fig. 4 shows that a sequence of such units can be wired up so that an input unit stays active long enough to trigger an output unit, which in turn inhibits the input. If the input unit did not resist this inhibition, it could fail to latch the output before shutting off. Elsewhere we have detailed the specific conditions that ensure proper sequential latching. To ensure that timing issues can be handled, we use the timer circuit in Fig. 2 **f** to implement an analogue of the delay gates used in digital logic. This mechanism activates a 'start' switch unit on the left, then integrates that signal in a 'ramp' unit, weighted by the start-to-ramp weight, until it triggers the 'trigger' unit to flip from 0 to 1. The delay duration is equal to this threshold value divided by the start-to-ramp weight. These dynamics are very similar to those implementing hypothesis-testing in Fig. 3, but now the only evidence is the passing of time.
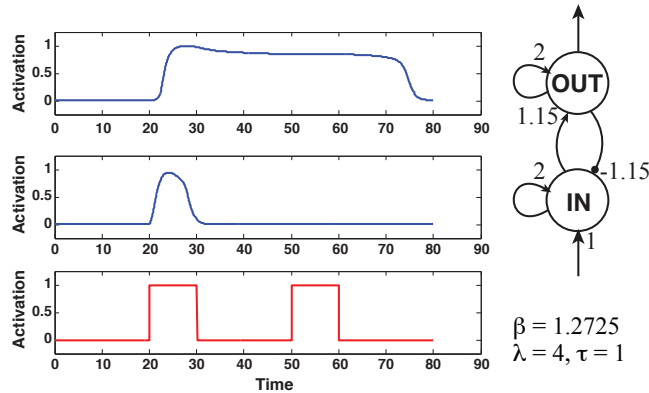

Figure 4: A production that negates its own if-condition. Bottom layer: input signal (red). Middle layer: IN unit activation. Top layer: OUT unit activation.

With these building blocks in hand, we can build arbitrarily complex circuits that implement logic gates and finite state machines, and thus special-purpose production systems. However, we still face the same critical problems facing all connectionist systems: if the semantics of a representation depend on what is connected to what, then how do separate representations share subcomponents? Or if their subcomponents conflict, then how are the proper subcomponents bound with the proper parent representation? Temporal synchrony has been widely considered to be a potential solution. The architectural assumptions are made that whatever is simultaneously active refers to the same entity, and distinct entities share different oscillation phases. We implement these as-

sumptions using the same machinery that underlies productions which cancel their own if-conditions.

Fig. 2 **g** shows that for each production trigger, we can assign an inhibitor. If a production fires, its output unit activates and triggers its own cancellation after a controllable delay (depending on connection strengths). However, the firing of a production can trigger a stored, hidden variable in a third bistable unit, which forces reactivation of the production after the inhibitor falls silent. This process repeats, triggering oscillations. When productions compete with each other, they push their active periods out of phase with each other, as shown in Fig. 5. When they do not, excitation causes them to entrain to the same phase. Thus conflicting representations locally decide which gets to broadcast information globally. If we allow for a plasticity signal that globally increases the learning rate of Hebbian connection plasticity between units, and if we activate this signal only at critical times, then we can burn in connections (possibly temporary connections) between units simply by activating them.
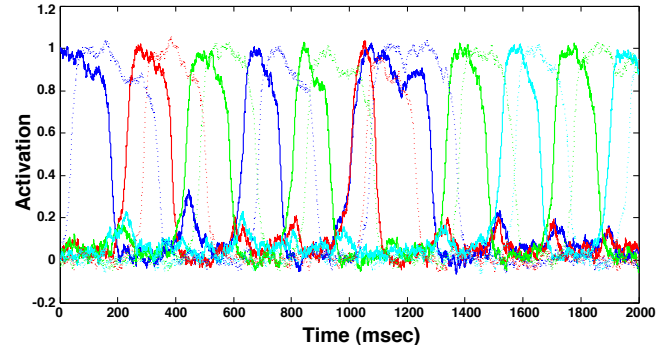

Figure 5: Relaxation oscillations among competing representations, allowing sharing of a single broadcast channel. Each solid color corresponds to one representation's bistable output unit; dashed curves correspond to the output's inhibitor.

More work will be needed to determine the scope of this approach to dynamic symbol and rule creation, to the implementation of a data type system such as exists in ACT-R, and to the binding problem more generally. With enough additional assumptions about the structure of the basic building blocks, it would be possible to translate between any given symbolic architecture and an architecture built from the components we have outlined. This must be the case in a trivial sense because our components are equivalent to circuits of resistors, capacitors and transistors. To stand as a psychologically plausible mapping, however, any subsymbolic theory will have to account for empirical data. We now focus on the kind of data for which our subsymbolic approach has something definite to say, leaving a more detailed investigation of the dynamics of connection-strength change for future work.

Laws of behavior

Two laws of behavior bear directly on the plausibility of the architectural building blocks. The first, known as the 'match-

ing law', states the following: that the ratio of rates of two (or more) different types of behavior that an animal engages in equals the ratio of the rewards earned for those behaviors. We showed in Simen and Cohen (2009) that the network in Fig. 3 reproduces this behavior. That network involves conflict between two productions that are supported by exactly the same amount of perceptual evidence. The exponentially weighted reward history of each response is encoded in the weight between the input unit and output unit of a production. This effectively changes the random walk thresholds for each response, while the walk itself is unbiased toward any response. The average result, in the case of two alternatives, is a state of exact matching of the behavior and reward ratios. When, instead, the reward history is encoded in connections from sensory inputs to the laterally inhibiting units, and input-output unit weights are held fixed, the model implements a softmax or logistic choice function defined on the difference between the reward histories. Evidence abounds for one or the other choice function in the instrumental conditioning literature since the time of Skinner. Thus the basic implementation of preferences for certain responses over others in the architecture meets a well-known psychological constraint on learning from reinforcement.

The other law regards timed behavior. A variety of different timing experiments show that the standard deviation of response times in such tasks is equal to a constant times the mean. The distribution of such responses is usually approximately Gaussian. The timing model in Fig. 2 **f**, accounts for this law. When a unit balances its self-excitation against its leak, it acts as an integrator. The model uses a simple error-correction rule to set the connection strength w from the start switch unit so as to ramp up to a level sufficient to trigger a switch from 0 to 1 in the output trigger. The integrator acts as a drift-diffusion process, since it integrates a constant drift term, w, corrupted by noise of amplitude $c\sqrt{w}$:

$$dV = w \cdot dt + c \cdot \sqrt{w} \cdot dW. \quad (6)$$

The trigger unit at the end of the chain defines a threshold on this diffusion process; call it z. Such a process produces a Wald, or inverse Gaussian, distribution of first-passage times (Luce, 1986). The mean RT of this process is z/w, and the standard deviation σ is $c\sqrt{z}/w$. Given that the ramping integrator unit cannot rise above a certain activation because of its saturation nonlinearity, then if we wish to minimize RT variability, we have the choice of minimizing z or maximizing w. The square root in the numerator indicates that increasing w will effect a larger reduction in variability than an equal increase in z. This implies that for all intervals, we should set z to a constant value that is as large as possible, without requiring the integrator to enter its highly nonlinear activation range. This in turn implies that $\sigma = \gamma z/w$, with $\gamma = c/\sqrt{z}$. That is, RT standard deviation is in constant proportion to the mean. Furthermore, as long as c is not too great — with a psychologically plausible value of 0.1 to 0.2, for example – the Wald distribution has very little skewness, and looks almost normal (and a slight positive skewness is often observed

in timing data anyway). Thus the model reproduces scalar invariance, and meets a second strong, empirical constraint.

Other behavioral and EEG predictions

We now examine two new predictions that regard the specific mechanism used to implement thresholds. In most decision making models (e.g., Bogacz et al., 2006), such thresholds are simply assumed to exist as a step function or Heaviside function, with a sharp discontinuity at the threshold. The bistable trigger mechanism described in Fig. 1 makes no such assumption, but nevertheless acts approximately as an all-or-none, digital device. Its hysteresis properties are critical for sequential processing, as we have shown, but does it make any testable predictions?

One is that if an input to a trigger unit with strong self-excitation is just below the point needed to trigger a transition from low to high activation, there will nevertheless be occasional triggerings due only to noise. This phenomenon — known as the escape from a double-well potential (Gardiner, 2004) — produces escape-time distributions defined in terms of exponential functions of the well depth (in our case, the remaining distance to the threshold).
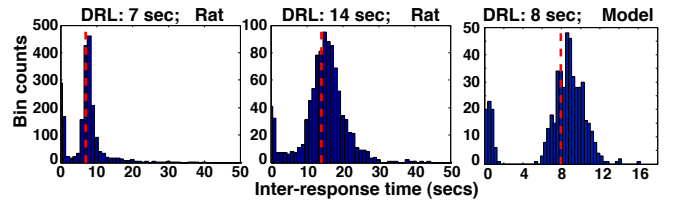

Figure 6: Two-component (exponential + scale-invariant Gaussian) response time distributions (rat, **a, b**, model, **c**).

In fact, a nearly exponential component of response times shows up in rat data collected in the differential reinforcement of low rates of responding (DRL) task. In this task, an animal must wait some minimum amount of time before making a response. Any response after this time is rewarded; any response that occurs prior to this waiting time relative to their last response resets the clock. Animals learn to wait in this task until shortly after the deadline, but they also emit a proportion of very fast responses that are apparently not controlled by a timer. Our model of this task involves using the timer circuit in Fig. 2**f** to implement the nearly Gaussian component of such RT distributions, but it also allows for direct connections between the start-switch and response trigger. This produces a proportion of fast responses that are nearly exponentially distributed. We reason that such a connection exists because of the way these tasks are acquired by animals: first, a contingency between some input stimulus and the response mechanism must be learned; second, a learned delay between responses is shaped through training.

The second prediction the bistable trigger mechanism makes regards the lateralized readiness potential (LRP) observed in any task with a motor response that occurs on one
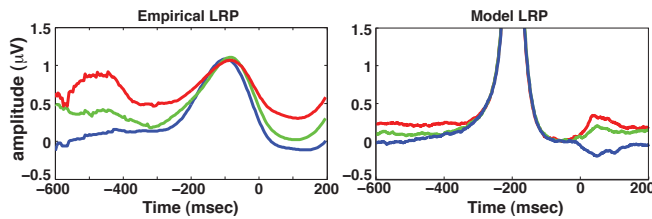
Figure 7: **a:** Average, response-locked LRP data (microvolts) from 8 human participants performing left vs. right dot-motion discrimination, with stimulus odds equal to 60:40 (blue), 75:25 (green), 90:10 (red). Behavioral responses occur at time 0. Data is baseline-corrected to align peaks. **b:** Model LRP (left threshold unit activation subtracted from right in Fig. 3, followed by a bilateral shutoff signal), with constant bias toward the right. The order of LRP differences between conditions is captured, but (as shown by the y-axis limits) capturing the smaller magnitude of the empirical peak requires additional assumptions.

side of the body. Prior to the movement, a voltage builds up over the part of motor cortex that is contralateral to the movement. A voltage also builds up on the same side as the movement, but not to the same degree. Then, just before the response is made, the LRP returns to baseline, because the voltage on both sides of the head over motor cortex becomes large and equal. We hypothesize that motor cortex houses response triggers, and we examined what would happen to a circuit in which a prior probabilities favored, say, a left button press rather than right button press. Although our bistable switches are nearly binary, they do involve slow, graded changes in activation level prior to the point at which they transition to a high activation. Because of this, and because this happens to a greater extent for the response trigger that is about to activate than for a trigger for the other, competing response, a difference in trigger activations develops, as shown in Fig. 7 **a**. We interpret this difference as a readiness potential. As a result, a consistent bias toward one response over the other should show up as an LRP both before and after the response. Such biases are expected in two-alternative perceptual decision making tasks with rewards for correct responses in which one stimulus is more frequently presented than the other (Bogacz et al., 2006). Simen et al. (2009) showed that human behavior in such tasks is consistent with the predicted bias toward the more frequent stimulus.

New LRP data from the same task shows the predicted physiological signature of such a constant bias: for a condition in which a right button-press response is always more likely to be correct than a left button-press, a persistent LRP should occur, with magnitude increasing as the prior probability increases for a given response. Data from 8 participants confirms this stimulus bias prediction (Fig. 7).

Conclusion

Here we have shown strong behavioral and electrophysiological evidence for key components of a neurally implemented production system architecture. These include bistable response units, and competitive response selection and hypothesis testing that are equivalent to random walk attractor dynamics. Thus the assumptions we made in order to achieve basic production system functionality seem to be justified. Much work remains to determine just how much production system functionality can truly be emulated by such systems. However, it is clear from working examples that simple cognitive models of problem solving can be so implemented (Polk et al., 2002; Simen & Polk, 2010), and we have outlined a mechanistic implementation of the temporal synchrony solution to the binding problem and the problem of dynamic linking among representations — problems which bedevil neural network architectures, but which are handled easily in standard production systems. We hope that future work on this topic will illustrate what constraints need to be imposed on production system programs in order for them to be 'compiled' into an equivalent neural network.

References

Anderson, J., & Lebiere, C. (1998). *The Atomic Components of Thought*. Lawrence-Erlbaum Associates.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced choice tasks. *Psychological Review*, *113*(4), 700–765.

Gardiner, C. W. (2004). *Handbook of Stochastic Methods* (Third ed.). New York, NY: Springer-Verlag.

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, *84*, 279–325.

Luce, R. D. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization*. New York: Oxford University Press.

McMillen, T., & Holmes, P. (2006). The dynamics of choice among multiple alternatives. *Journal of Mathematical Psychology*, *50*, 30–57.

Polk, T. A., Simen, P. A., Lewis, R. L., & Freedman, E. G. (2002). A computational approach to control in complex cognition. *Cognitive Brain Research*, *15*(1), 71-83.

Simen, P. A., & Cohen, J. D. (2009). Explicit melioration by a neural diffusion model. *Brain Research*, *1299*, 95–117.

Simen, P. A., Contreras, D., Buck, C., Hu, P., Holmes, P., & Cohen, J. D. (2009). Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *Journal of Experimental Psychology: Human Perception and Performance*, *35*, 1865–1897.

Simen, P. A., & Polk, T. A. (2010). A symbolic/subsymbolic interface protocol for cognitive modeling. *Logic Journal of the IGPL*, *18*, 705–761.