

Toward Network Intelligence

Alex Pentland

pentland@mit.edu

Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.

This article proposes a conceptual framework to guide research in neural computation by relating it to mathematical progress in other fields and to examples illustrative of biological networks. The goal is to provide insight into how biological networks, and possibly large artificial networks such as foundation models, transition from analog computation to an analog approximation of symbolic computation. From the mathematical perspective, I focus on the development of consistent symbolic representations and optimal policies for action selection within network settings. From the biological perspective, I give examples of human and animal social network behavior that may be described using these mathematical models.

1 Introduction ---

Progress in understanding the phenomenon of intelligence as a result of neural computation has been astounding since the time I was helping Terry Sejnowski and Steve Zucker, and later Geoffrey Hinton, run workshops on computational neuroscience at the Woods Hole Marine Biological Laboratory. Despite amazing progress, many of the fundamental problems in understanding network intelligence remain unanswered. Intelligence is defined by the *Oxford English Dictionary* as the ability to acquire and apply knowledge and skills. Today, most neural computation avoids the hardest parts of modeling intelligent behavior by framing tasks so that they avoid problems of representation, context, and intention by having humans select data that address a specific knowledge or skill task and preestablishing training or success criteria for the computation.

Extremely large artificial networks such as foundation models, which exhibit surprising capacity and flexibility, are rich descriptions of relationships within the data they are trained on because their training is not supervised. However, they do not by themselves have the intentionality or generativity that biological entities require for survival, although they provide sophisticated perception and example generation capabilities that could be a major part of an agent that could survive in the real world. An important question that I consider is whether foundation models have

somehow passed a threshold and possess the abstraction and generativity capabilities characteristic of symbolic systems.

At an even more general level, there is no “computational theory” in the sense of Marr and Poggio (1976), who describe the evolutionary fitness problem that biological networks are solving and connect the specifics of the network’s computation to solution of the fundamental problems of survival and reproduction. In other words, we do not understand why biological neural networks have the structure and behavior that we observe.

In the spirit of exploring possibilities for a computational theory of network intelligence, I focus on three basic issues. The first is how the capabilities of symbolic computation arise, since this is widely considered to be fundamental to human-level intelligence. Second, rapid, stable, and generalizable learning is still something of a mystery; most of today’s methods require huge numbers of examples that accurately sample the problem distribution. And third, there is the difficulty that evolution and Darwinian selection require learning actions that improve fitness and survivability rather than just learning accurate representations or mimicking carefully crafted training examples. I address each of these three in turn, first by describing recent mathematical progress in adjacent fields and then presenting examples of biological networks that illustrate the ideas.

2 From Signals to Symbols

Why do humans have symbolic intelligence? Perhaps the most common hypothesis is that it is to support language and thus social learning. Symbolic communication requires establishing a shared vocabulary (a set of symbols) with a clear association to the external world. Unlike the traditional biological evolution models where organisms become adapted to their local niches (and unlike traditional quantization theory, where codebooks are adapted to local source distributions), language evolution is necessarily a social phenomenon, since without social interaction, there is no need for shared vocabularies. Consequently, communication vocabularies must evolve to balance individual concerns and social exchange. In the neural computation setting, this means that states of each interacting neural assembly must have the same environmental associations as the states of connected, cooperating neurons.

Formally, symbolic representation is based on signal description (e.g., quantization), that is, the assignment of symbols to ranges of sensory signals. In our 2021 IEEE signal processing paper (Mani, Varshney, & Pentland, 2021), we prove that under surprisingly general conditions, a quantization network game with cooperative communication will evolve a symbolic representation and settle into a Nash equilibrium with local agents having a shared vocabulary.

Using game theory to characterize this type of longer-range descriptive interaction has long been proposed for understanding how neurons

cooperate to segment continuous portions of images (Miller & Zucker, 1991), but the process of generating a symbolic description also requires that each node chooses a vocabulary for describing the signals that it sees in such a way that its vocabulary is an accurate description of both the current signals and the surrounding nodes' signals.

In contrast to traditional results in the evolution of symbolic communication, we found that several vocabularies of symbols may coexist in a Nash equilibrium of this network learning process. The overlap between vocabularies is high for nodes that communicate frequently and have similar local sources. This process provides a good account of the emergence of different languages in separated human groups and the greater frequency of words specialized to the local situation.

A concrete example is learning the sounds and meanings of words. Our 2002 *Cognition* paper (Roy & Pentland, 2002) showed that early word learning of human infants can be accounted for by quantization of the mutual information between their audio and visual perceptual streams. Later work showed that feedback from other people about the quality of audio encoding and appropriateness of referent is what guides the word learning process into its more-or-less final form.

Importantly, these sort of architectures are not specialized just for word learning. For instance, they can be also applied to action selection using affordance maps, as Gibson (1979), suggested. For example, if red circular blobs on a green background are associated with a map of affordance such as "edible," then outputs of the network can be used to trigger eye movements or further mental processing needed to facilitate picking red berries on a green bush.

This type of clustering-and-communication feedback may be happening when we train very large artificial networks by looking at all the relationships within the training corpus. Small networks have long been known to be quite good at finding local correlations, but in foundation models, longer-range structure is also very influential, leading to a balance between local description and long-range comparison, exactly the conditions that we have shown are sufficient for the emergence of symbolic representation. The generation of descriptive symbols for heterogeneous phenomena is a natural and likely common consequence of a cooperative communication for signal description, and the availability of symbols associated with environmental regularities pairs nicely with action selection networks such as the distributed Thompson sampling architecture discussed in section 4.

3 Rapid, Stable, and Generalizable Learning ---

Popular neural computation algorithms famously require huge amounts of data, and even specialization of the huge foundation models that are currently popular typically requires a great deal of data. Even so, most models

exhibit relatively poor generalization, not only because of overfitting but also because of hidden assumptions, such as the implicit assumption that data distributions are compact and stable. These are problematic assumptions because many social, financial, and biological phenomena have long-tailed and nonstationary distributions. In recent years the mathematics to address these limitations have been developed, and we can make use of these innovations to improve neural computation algorithms.

For example, one of the common frameworks for neural computation is Q-learning, which maximizes the expected value of the total reward over successive steps for any finite Markov decision process (FMDP). Q-learning works to discover an optimal action-selection policy for any given FMDP, given infinite exploration time and a partly random policy. In practice, the convergence of this optimization method is generally very slow. Moreover, this sort of reinforcement learning is unstable when a nonlinear function (such as neural network) is used to represent Q. This instability comes from the correlations in the sequence of observations and between Q and the target values, as well as the nonlinearity of Q.

An important advance associated with Q-learning is the ability to connect delayed rewards to actions. The technique of experience replay, a biologically inspired mechanism that uses a random sample of prior actions instead of the most recent action, helps suppress spurious correlations, and periodic update of Q reduces correlations with the target (Mnih et al., 2015). This allows Q-learning to be used, for instance, to develop a play policy for games like chess or Go.

Q-learning, however, does not take advantage of recent mathematical results concerning combination of evidence from different agents (or nodes) in a network to achieve optimal use of data, specifically, the mathematics of minimum regret decision making, except for very specific single-agent exploration policies. I refer to agents rather than nodes in the following to emphasize the potential for groups of notes to act as action selection alternatives. The criterion of minimum regret is the requirement that an agent or group of agents make the best action selection possible at each time period given the information and previous experience available at the time.

4 Networks and Thompson Sampling

Classic minimum-regret decision making is called Thompson sampling (TS), and the literature often refers to this as a *bandit problem*, because of the formal equivalence to the question of which slot machine (the “one arm bandit”) to try in a gambling casino. In the past decade, the mathematical solutions to such problems have been extended to networks of agents; for example, a gambler observes the payouts of other casino patrons and combines those observations with his or her personal knowledge to decide which slot machine to try next.

The classic TS strategy for optimal (minimum regret) decision making has demonstrated excellent performance in many domains, enough so that it is a standard approach in domains such as signal processing, medical decision making, and finance. It shows very fast convergence to optimal policies, good generalization to new and changing situations (Dubey, Ramanathan, Pentland, & Mahajan, 2021), and the ability to work with noisy and ill-conditioned data inputs.

A fundamental difference between TS and Q-learning is that TS incorporates an exploration strategy to help it find the highest reward actions, whereas Q-learning is an optimization algorithm. This means that TS combines exploitation (optimization by computing the posterior) with exploration (by sampling from the posterior to pick an action or “bandit arm” that may yield greater rewards). In contrast, Q-learning simply tells us how to estimate the Q values but still requires an exploration strategy such as UCB. In the context of neural networks, typical actions are adjusting network weights.

TS can provide an exploration strategy for Q-learning that has strong optimality guarantees. If we were to compute Q-values assuming some parametric form of the environment (e.g., gaussian rewards), and then select actions to explore or exploit based on samples previously drawn from this parametric form, the resulting strategy could be thought of as being similar to TS.

This example highlights a second basic difference between Thompson sampling and Q-learning. TS typically operates on a discrete, symbolic representation of actions, whereas Q-learning is model free. Essentially, to compute the posterior reward of an action, TS assumes a parametric model for the prior distribution and a likelihood function that is updated with observations. Q-learning instead starts from zero and simply keeps updating its parameters without trying to model the environment.

Sections 2 and 3 and this one have described how cooperative communication for quantization naturally segments input signals (which can be, for example, either affordances or percepts) into a set of symbols that can be used for accurate communication between agents or neurons. The output of such a cooperative quantization computation is exactly the sort of symbolic parameterization of the environment required for TS to produce an optimal, minimum-regret action policy.

5 Networks of Smart Neurons

The mathematics of TS can easily be extended to networks of agents, so that they may collectively develop an optimal, minimum regret policy for action. The core of the distributed Thompson sampling (DTS) strategy is for each decision-making agent to use the experience of other agents to form an estimate of the prior probability for each potential action and then multiply

this prior by each action's likelihood distribution as determined from their personal knowledge. This produces a posterior distribution of what reward each action is expected to produce.

Importantly, the symbolic representations that naturally emerge from cooperative signal description can be used as a discrete, symbolic representation for DTS. This allows the posterior distribution estimated by DTS to efficiently and robustly drive the cooperative symbolic encoding. For example, if cooperative signal description distinguishes small red blobs on a green background, then DTS can efficiently discover that a good reward can be obtained by using a red blob as a trigger for eye movements or further mental processing. In this way the encoding for "red blob" can be mapped from a purely descriptive representation to a functional or intentional representation—for example, an affordance map that relates input values to possible actions. This is the sort of computation required to implement the automatic releasing-stimuli behaviors seen in simple animals.

Several research groups (including my own) have recently extended DTS to cover communication-limited networks analogous to the networks seen in both biological systems and human social behavior (Dubey et al., 2021; Dubey & Pentland, 2020a, 2020b). We have also extended the DTS framework to adversarial situations and situations where different agents have different goals. This is accomplished by having each agent compare the choices he or she makes to the choices of others, and from that comparison estimate the similarity of the utility function of the other agents to one's own utility function, as in Dubey and Pentland (2020a).

These results allow agents to select the best subset of agents with which to trade experiences and identify agents acting in an adversarial manner or reacting to different signals. A concrete image example is deciding which image patches have the same source distribution, thus segmenting the image into regions that are "the same." Application of these techniques to neural networks can provide a learning rule for selecting and reinforcing connections between neurons in order to achieve minimum regret action selection even in the presence of inhomogeneous input signals and costly or limited interneuron communication.

In plainer English, these extensions of DTS provide an iterative estimation framework that allows networks of agents or nodes with different inputs and different output connections (e.g., different action selection) to modulate their network connectivity in order to maximize their reward (Dubey & Pentland, 2020a). The result is a minimum regret decision sequence for each agent or node and an analogous optimality result for connected subgroups within the group as a whole. Thus, DTS offers a way to prune large neural networks and discover core sets that are the most important for high-reward behaviors. Moreover, under relatively mild assumptions, the stability of the network and its output can be guaranteed as shown in our 2020 paper in the *Proceedings of the National Academy of Science* (Lera, Pentland, & Sornette, 2020).

6 Toward a Computational Theory of Network Intelligence

Developing a computational theory of how intelligence evolved in neural networks is problematic given our current state of knowledge about the early evolution of life. However, we can say much more about the sort of computational theory that applies to modern biological networks that have many individual members, such as those of social animals (including everything from cooperative single-cell organisms to humans), because the action selection and reward functions are easier to observe experimentally. For instance, we know that local environments change, and so virtually all species, from single cells on up, have developed ways of finding new food sources. The problem of exploring for new food sources while maintaining sufficient exploitation of current sources is thus a nearly universal problem for biological networks.

Examples like this exploration-exploitation trade-off can guide the development of a computational theory of how biology solves everyday but critical problems like feeding and reproduction in an uncertain and changing environment, problems that are certainly an important part of all species' survival fitness. Consequently, a natural hypothesis is that the mathematical patterns we see in today's biological organisms (including primitive single-cell organisms) may be a common networking pattern selected by evolution and thus may plausibly apply to neural networks as well.

It is encouraging that “distributed Thompson sampling” is a good description of the group foraging behavior of many social animals (Berger-Tal, Nathan, Meron, & Saltz, 2014). Group foraging may be viewed as a “portfolio strategy” where the animals do not just choose the maximum posterior likelihood action (e.g., the action that has been yielding the most food); instead, they make the frequency of different actions proportional to the posterior likelihood. The frequency with which various actions occur can thus incorporate information from the observable experience of all the animals in the social group.

Sampling and integration of group experience provide a solution to the exploration-exploitation dilemma, where animals must allocate some effort to actions that have been reliably producing good results, while at the same time exploring for new resources, with the consequence that their portfolio of frequent actions continuously evolves over time. The fact that this behavior is ubiquitous, even in single-cell animals, means that DTS is a good hypothesis for understanding the computational problems that shaped the evolution biological neural networks.

7 Human Network Behavior

Studies of human financial decision making, as well as human mobility and shopping behavior over many days, show the exploration-exploitation

pattern characteristic of DTS (Adjodah et al., 2021; Krumme, Lorente, Cebrian, Moro, & Pentland, 2013). For instance, in Adjodah et al. (2021) we found that financial experts use observations about the actions of other experts to control their portfolio risk in a manner consistent with DTS (Adjodah et al., 2021). An important function of the exploration actions is to avoid behavioral rigidity, where only a limited number of very familiar actions are chosen within a group. The phenomenon of insufficient exploration, resulting in a static portfolio of actions, is familiar in human social networks as echo chambers and group-think.

Exploration of novel actions is therefore critical for avoiding unforeseen risks, finding new opportunities, and adapting to changing conditions. In the context of neural networks, this means continually experimenting with weights and pruning unrewarding connections. One of the important strengths of the DTS framework is that it provides a formal method of determining when there is insufficient exploration, and thus risk of group-think, as well as a formula for derisking decisions by accounting for rare outcomes and sampling error (Dubey & Pentland, 2020b).

Inspired by these mathematical decision systems results, my research group has worked on the cognitive science version of this decision-making literature and believe that we have made interesting progress toward relating human decision making to Thompson sampling style optimal decision making. This work began with our 2010 *Science* paper (Woolley, Chabris, Pentland, Hashmi, & Malon, 2010), which showed that the collective intelligence of small human groups is strongly correlated with entropy of the group communication and that it is different from, and often more effective than, individual human intelligence. This insight was sharpened in our 2013 *Scientific Reports* paper (Krumme et al., 2013), which showed that the day-to-day behavior of human populations exhibits the same exploration-exploitation behavior seen in social animal species.

More recently, our 2021 *Cognition* paper (Krafft, Shmueli, Griffiths, Tenenbaum, & Pentland, 2021) showed that commonly observed individual-level social heuristics closely approximate DTS group decision making and accurately model human small-group behavior in consumer financial markets. Our 2021 *Entropy* paper (Adjodah et al., 2021) showed that financial experts show the same DTS behavior and, in particular, use the social information of others to control portfolio risk. These results were extended by in our 2020 *Proceedings of the National Academy of Sciences* paper (Almaatouq et al., 2020), which showed that group performance can be dramatically improved by selecting agents with similar utility functions, further motivating the idea of limiting inputs to action selection decisions via segmentation of the surrounding network by similarity in source distribution and symbolic representation.

We have also seen that the Thompson sampling framework provides a good model of some important examples of multimodal learning and decision making. For instance, our 2002 *Cognitive Science* paper on infant

language learning (Roy & Pentland, 2002) showed that audiovisual cues from adults provided critical cues allowing infants to locate and disambiguate nouns within the audio signal. In more recent work, we have shown that the DTS sampling framework provides a useful model for humans to create and learn category boundaries in visual stimuli (Epstein, Groh, Dubey, & Pentland, 2021) and for highly efficient domain generalization in image classification (Dubey et al., 2021).

8 Summary

Neural computation today is dominated by excitement over the new-found power of deep networks. To move to the next level of performance, the field will have to address some hard questions about the purpose and nature of intelligence and how neural computation fits into these larger questions. This paper makes three main points that address these issues:

1. Cooperative signal description between different node assemblies naturally generates stable symbolic representations that support accurate, shared communication. This may in part be responsible for the surprising flexibility of large foundation models.
2. Thompson sampling (TS) can extend Q-learning to symbolic computation, providing efficient and robust learning of minimum regret action policies, including improved network connectivity update rules, tolerance of erroneous and inhomogeneous inputs, and improved response to nonstationary and long-tailed distributions. Thus, TS offers a principled way to prune large neural networks in order to retain only the most useful nodes and features.
3. Network minimum regret behavior is very common in real-world biological networks and systems such as human social networks. It also appears to apply to social animals more generally, including some simple and primitive organisms, and thus is plausible as a computational theory for much of neural computation.

The goal of this article has been to provide a framework to guide research into how neural computation can produce intelligent behavior. It is my hope that the research program, scientific papers, experiments and theorems that I have described here will help future researchers answer these important questions.

Acknowledgments

This research is the thesis work of several generations of my graduate students, who have been supported by the MIT Media Lab's Industry Consortium and my Trust Data Consortium within MIT's Institute for Data Systems and Society.

References

- Adjodah, D., Leng, Y., Chong, S., Krafft, P., Moro, E., & Pentland, A. (2021). Accuracy-risk trade-off due to social learning in crowd-sourced financial predictions. *Entropy*, 23(7), 801. 10.3390/e23070801, PubMed: 34202445
- Almaatouq, A., Noriega-Campero, A., Alotaibi, A., Krafft, P., Moussaid, M., & Pentland, A. (2020). Adaptive social networks promote the wisdom of crowds. In *Proceedings of the National Academy of Sciences*, 117(21), 11379–11386. 10.1073/pnas.1917687117
- Berger-Tal, O., Nathan, J., Meron, E., & Saltz, D. (2014). The exploration-exploitation dilemma: A multidisciplinary framework. *PLoS One*. <https://doi.org/10.1371/journal.pone.0095693>
- Dubey, A., & Pentland, A. (2020a). Private and byzantine-proof cooperative decision-making. In *Proceedings of the 20th Conference on Autonomous Agents and Multiagent Systems* (pp. 357–365).
- Dubey, A., & Pentland, A. (2020b). Kernel methods for cooperative multi-agent contextual bandits. In *Proceedings of the International Conference on Machine Learning* (pp. 2740–2750).
- Dubey, A., Ramanathan, V., Pentland, A., & Mahajan, D. (2021). Adaptive methods for real-world domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 14340–14349). Piscataway, NJ: IEEE.
- Epstein, Z., Groh, M., Dubey, A., & Pentland, A. (2021). An experimental study of social influence and information design. In *Proceedings of the ACM on Human-Computer Interaction* (pp. 1–18). New York: ACM.
- Gibson, J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.
- Krafft, P., Shmueli, E., Griffiths, T., Tenenbaum, J., & Pentland, A. (2021). Bayesian collective learning emerges from heuristic social learning. *Cognition*, 212, 104469. 10.1016/j.cognition.2020.104469, PubMed: 33770743
- Krumme, C., Lorente, A., Cebrian, M., Moro, E., & Pentland, A. (2013). The predictability of consumer visitation patterns. *Scientific Reports*, 3, 1645. 10.1038/srep01645, PubMed: 23598917
- Lera, S., Pentland, A., & Sornette, D. (2020). Prediction and prevention of disproportionately dominant agents in complex networks. In *Proceedings of the National Academy of Sciences*, 117(44), 27090–27095. 10.1073/pnas.2003632117
- Mani, A, Varshney, L., & Pentland, A. (2021). Quantization games on social networks and language evolution. *IEEE Transactions on Signal Processing*, 69, 3922–3934. 10.1109/TSP.2021.3090677
- Marr, D., & Poggio, T. (1976). From understanding computation to understanding neural circuitry. AI MIT memo AIM (37). <http://hdl.handle.net/1721.1/5782>
- Miller, D., & Zucker, S. (1991) Copositive-plus Lemke algorithm solves polymatrix games. *Operations Research Letters*, 10, 285–290. 10.1016/0167-6377(91)90015-H
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A., Veness, J., Bellemare, M., ... Fiedland, Andreas K. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533. 10.1038/nature14236, PubMed: 25719670
- Roy, D., & Pentland, A. (2002). Learning words from sights and sounds: A computational model. *Cognitive Science*, 26(1), 113–146. 10.1207/s15516709cog2601_4

Woolley, A., Chabris, C., Pentland, A., Hashmi, N., & Malone, T. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330 (6004), 686–688. 10.1126/science.1193147, PubMed: 20929725

Received September 10, 2021; accepted January 8, 2022.