Modeling the Learning and Use of Probability Distributions in Chimpanzees and Humans

Thomas R. Shultz[23] & Ardavan S. Nobandegani[13]
thomas.shultz@mcgill.ca, ardavan.salehinobandegani@mail.mcgill.ca
[1]Department of Electrical & Computer Engineering, McGill University
[2]School of Computer Science, McGill University
[3]Department of Psychology, McGill University

Abstract

We present a neural-network computational model of a recent experiment revealing that chimpanzees show some ability to reason probabilistically. Specifically, we show that the *neural probability learner and sampler* (NPLS) system can account for both success by chimpanzees and better performance by human controls. NPLS effectively combines learning probability distributions with sampling from those learned distributions to guide action choices. Because NPLS also simulates learning and use of probability distributions by human infants, this brings us closer to a unifying model of probabilistic reasoning, across various age groups and species.

Keywords: probabilistic reasoning; neural networks; neural probability learning and sampling; chimpanzees; Weber's law

Introduction

There has been intense recent interest in the learning and use of probability distributions, due in large part to the recognition that much of cognition involves making decisions in uncertain situations. A series of experiments has shown than even pre-linguistic infants show some amazing abilities in this domain: they learn simple binary probability distributions, use that knowledge to guide their intentional actions, and show differential surprise at seeing their probabilistic expectations violated (Denison, Reed, & Xu, 2013; Denison & Xu, 2010, 2014, 2019; Xu & Garcia, 2008).

Computational modeling of how these infant phenomena has attempted to solve the mystery of how such young infants could deal with probabilities long before they were able to explicitly count and divide (Shultz & Nobandegani, 2021). These infant experiments have also attracted the interest of researchers of animal cognition to see if other species could learn and use probability distributions (Eckert, Call, Hermes, Herrmann, & Rakoczy, 2018; Eckert, Rakoczy, & Call, 2017; Rakoczy et al., 2014). Initial experiments with monkeys and apes, using similar research designs from the infant studies, revealed some degree of success but also some divergence in results and interpretations. Here we focus on simulating the most recent of these non-human primate experiments, as it corrects some shortcomings of the earlier experiments and includes samples of both chimpanzees and adult humans (Eckert et al., 2018).

Experiment with Chimpanzees and Humans

Eckert and colleagues (2018) tested 24 chimpanzees in their home sanctuary using a paradigm inspired by the infant experiments and initial primate experiments. This paradigm required a chimp to choose between random binary samples drawn from populations of food items with different ratios of preferred (peanuts) and non-preferred items (carrot pieces). In a series of experimental conditions, the ratio of the two ratios to be discriminated (RoR) was varied from 1 to 16 in order to assess the Weber effect that the two ratios would be easier to distinguish with increasing RoR. Human adults (n = 144) were tested in a computerized version of an analogous task in which they were asked to blindly select a marble of a particular color from one of two containers, each containing two colors of marbles, where the ratios and RoRs were the same as those used with the chimpanzees (Eckert et al., 2018).

For the chimpanzees, the experimenter's hands were crossed on a random half of trials to control for proximity effects between food population and sample. As in some of the infant experiments (Denison & Xu, 2014), a positive correlation between frequencies of preferred items and probabilities was eliminated by ensuring that the more favorable population had fewer preferred items than the unfavorable population (see Table 1). Correct choices (of the preferred sample) increased with log RoR, taken as evidence for operation of Weber's law (Eckert et al., 2018).

Table 1: Favorable and unfavorable populations, ground-truth probabilities, and ratios used in the empirical experiments and simulations. *Favorable* is represented here as *fav*, *unfavorable* as *unf*, *population* as *pop*, *preferred* item as *pre*, *not-preferred* as *not*, and *ratio-of-ratios* as *ror*.

fav pop		unf pop		probability		ratios		
pre	not	pre	not	fav	unf	fav	unf	ror
28	20	56	80	.58	.41	1.40	.70	2
28	14	56	112	.67	.33	2.00	.50	4
28	11	56	132	.72	.30	2.55	.42	6
28	10	56	160	.74	.26	2.80	.35	8
28	8	56	192	.78	.23	3.50	.29	12
28	7	56	224	.80	.20	4.00	.25	16

Methods

The same neural probability learner and sampler (NPLS) system (Shultz & Nobandegani, 2021) used to simulate the human infant experiments (Denison et al., 2013; Denison & Xu, 2010, 2014, 2019; Xu & Garcia, 2008) is used here. NPLS includes a modified version of the sibling-descendant cascade-correlation (SDCC) algorithm that has been used to simulate many deterministic phenomena in cognitive and language development (Nobandegani & Shultz, 2022; Shultz,

In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society.*

2003, 2017), and the widely used Markov chain Monte Carlo (MCMC) sampling algorithm that has been used to simulate a broad range of empirical findings in human probabilistic reasoning and decision making (Dasgupta, Schulz, & Gershman, 2017; Nobandegani & Shultz, 2020).

SDCC uses deterministic, feed-forward networks that learn from labelled examples by reducing overall prediction error (Baluja & Fahlman, 1994). The unit activations are passed forward from inputs that describe examples to hidden units that transform the inputs into more abstract representations, and finally to output units representing responses to that particular input. Network outputs can be regarded as predictions about what will happen, while target output represents what actually happens. There are two phases in SDCC processing: input phase and output phase. During output-phase learning, connection weights are adjusted to reduce sum-of-squared network error *E*:

$$E = \sum_o \sum_p \left(A_{op} - T_{op}\right)^2 \qquad (1)$$

where *A* is the actual output activation for unit *o* and pattern *p*, and *T* is the target output activation for this unit and pattern.

SDCC learning begins with only the input and the output layers of units, and then recruits hidden units one at a time, as needed, to solve the problem being learned. SDCC thus constructs its own network topology, as opposed to being designed by a programmer. In the input phase, weights entering randomly-generated candidate hidden units are trained to increase covariation of candidate hidden-unit output activation with network error. The highest correlating unit is then installed either on the highest layer of hidden units or on its own higher layer, whichever shows the better absolute covariation with network error. Input weights to each recruited hidden unit are frozen as soon as the recruited unit is installed. At this point, control is passed back to output phase, in which connection weights entering output units are adjusted one layer at a time, thus never requiring (unrealistic) propagation of error signals back through the network. The function to maximize in input phases is the covariance *C* between candidate-hidden-unit activation and overall network error:

$$C = \frac{\sum_o \left| \sum_p (h_p - \langle h \rangle)(e_{op} - \langle e_o \rangle) \right|}{\sum_o \sum_p \left(e_{op} - \langle e_o \rangle\right)^2} \qquad (2)$$

where h_p is activation of the candidate hidden unit for pattern p, $\langle h \rangle$ is the mean activation of the candidate hidden unit for all patterns, e_{op} is the residual error at output o for pattern p, and $\langle e_o \rangle$ is the mean residual error at output o for all the training patterns.

For probability learning, the networks use an asymmetric sigmoid activation function:

$$y_i = \frac{1}{1 + e^{-x_i}} \qquad (3)$$

In order to prevent the algorithm from recruiting new hidden units *ad infinitum*, NPLS monitors its progress in error reduction over learning cycles. SDCC already possessed the capacity to monitor its progress during both input and output phases, using parameters for threshold and patience (Baluja & Fahlman, 1994). During output phases, SDCC adjusts its connection weights to reduce error. But when error reduction stagnates, there is a shift to input phase to recruit a new hidden unit, adjusting the input weights to candidate units in order to increase covariation between candidate-unit activations and network error. In each of these two phases, stagnation is detected when progress no longer exceeds a threshold parameter for some number of training epochs, specified by a patience parameter.

This idea was extended by adding an outer loop having its own threshold and patience parameters to monitor progress over learning cycles, where each cycle is an input phase and the next output phase (Shultz & Doty, 2014). This allows NPLS to stop when learning stagnates, which happens to coincide with accurate estimates of the probability distribution being learned. Thus, NPLS can learn an unnormalized multivariate probability distribution from examples specifying whether or not an output occurs in the presence of a particular input (Kharratzadeh & Shultz, 2016).

We run 20 NPLS networks in each of the six conditions representing the six RoR values for the chimpanzee and human populations studied by Eckert et al. (2018). Networks are trained on event sequences with an input unit arbitrarily coding for the identity of the source container (1 or 2) and an output unit coding 1 for presence and 0 for absence of an object type (food item for chimpanzees, or marble color for humans).

With this deterministic binary coding, corresponding directly to the visual stimuli experienced by the empirical experimental participants, networks learn to output the probability of drawing a preferred item from a favorable or an unfavorable population. Importantly, ground-truth probabilities are not used as learning targets. Probability estimates are instead an emergent property of NPLS learning, represented by learned network output activations (Shultz & Nobandegani, 2021).

Table 2 shows an example of the coding scheme for a much simpler binary distribution with frequency ratios of 4:1 vs. 1:4. This example requires five training patterns for each ratio. In 4 of 5 examples for container 1, a focal object appears. For container 2, a focal object appears in only 1 of 5 examples. Our simulations use the exact ratios from the empirical experiments, realistically representing what a participant sees in the containers. An asymmetric sigmoid activation function on the output unit constrains outputs to the 0-1 range of probabilities.

Table 2: Schematic coding of a binary probability distribution

4:1		1:4	
Input	Output	Input	Output
1	1	2	1
1	1	2	0
1	1	2	0
1	1	2	0
1	0	2	0

A parameter called score-threshold is used to distinguish deeper learning in humans from the apparently shallower learning of chimpanzees. Formally, score-threshold represents the allowable distance between actual and expected outputs that are considered as correct. The default score-threshold in SDCC is .4, allowing for an uncertain, buffer zone from outputs of .4 to .6. We set score-threshold to .5 for human simulations and .6 for chimpanzee simulations, reflecting hypothesized imperfect learning for humans and slightly less perfect learning for chimpanzees.

Following recent advances in how to probabilistically generate examples from learned categories (Nobandegani & Shultz, 2017), NPLS then uses an MCMC algorithm to simulate how participants select a container, essentially converting a deterministic neural network into a probabilistic generative network. NPLS learns categories from examples and then generates examples from those categories using the learned weights (Nobandegani & Shultz, 2018). Such backward inferences can be mathematically characterized as a form of sampling from the underlying, learned probability distribution. A participant could mentally draw a sample, cued by the higher probability of a preferred object, and thus identifying the more favorable sample for obtaining that object.

Formally, NPLS induces a probability distribution $p(\mathbf{X}|\mathbf{Y})$ on the deterministic input-output mapping $f(\mathbf{X}; W^*)$ learned by an NPLS network, and uses MCMC to sample from that induced distribution. The induced distribution is given by:

$$p(\mathbf{X}|\mathbf{Y} = Y) \propto exp(-\beta||Y - f(\mathbf{X}; W^*)||_2^2) \quad (4)$$

where $||\cdot||_2$ is the l_2-norm, W^* the set of weights for a network after training, and β a damping factor. For an input instance $\boldsymbol{X} = X$ belonging to the desired class Y, the network output $f(X; W^*)$ is expected to be close to Y in the l_2-norm sense. Equation 4 adjusts the probability of input instance X to be inversely proportional to the base-e exponentiation of the l_2 distance. The NPLS system can handle any MCMC method, including Metropolis-Adjusted Langevin, a gradient-based MCMC method, which could be implemented in a biologically-plausible way (Moreno-Bote, Knill, & Pouget, 2011; Savin & Denève, 2014).

Results

Mean output activations are plotted across RoRs, with SDs, for 20 NPLS networks in the human simulations (Figure 1) and the chimpanzee simulations (Figure 2). Comparing these network probability estimates to the ground-truth probabilities calculated in Table 1 reveals better accuracy and less variation for the human than for the chimpanzee simulated agents. Mean network output activations correlate highly with ground-truth probabilities across the twelve conditions of the six experiments, in both chimpanzees ($r(10) = .9709$, 95% CI = [.8966, .9920], $p < 10^{-6}$) and humans ($r(10) = .9999$, 95% CI = [.9997, .9999], $p < 10^{-15}$). Human simulated agents recruit more hidden units ($M = 3.54$) than do chimpanzee simulated agents ($M = 1.54$), independent $t(139.53) = 11.986$, $p = 3.36\text{E-}23$, Cohen's $D = 1.55$. This reflects deeper learning in human simulated agents than in chimpanzee simulated agents.

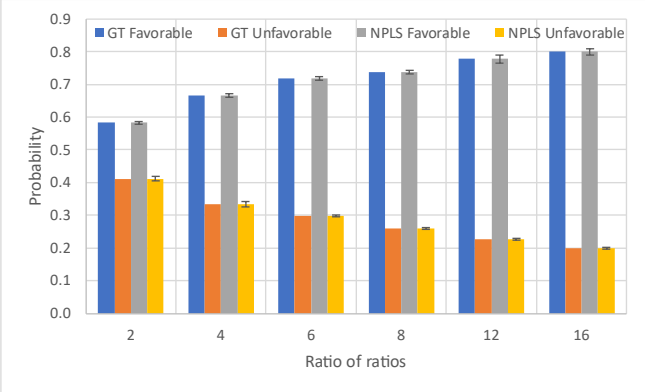

Figure 1: Comparison of mean network probability estimates (output activations), along with SDs, to ground-truth probabilities for 20 NPLS networks in human simulations.

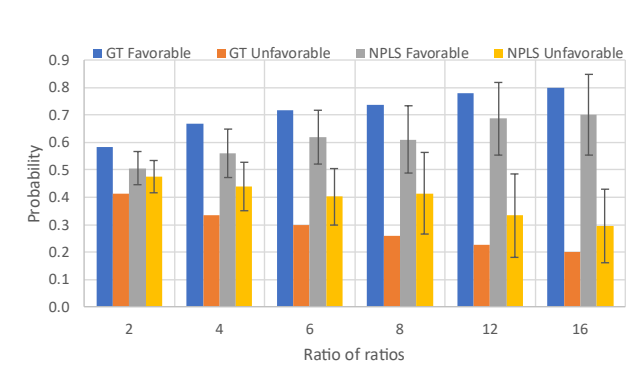

Figure 2: Comparison of mean network probability estimates (output activations), along with SDs, to ground-truth probabilities for 20 NPLS networks in chimpanzee simulations.

Mean sampling probabilities (with SDs) are shown in Figure 3. They correspond well with the empirical experiments (Eckert et al., 2018) in that simulated agents select correctly with probability close to .5 with a lower RoR of 2, and approach .7 for the chimpanzees and .8 for humans at a higher RoR of 16. Mean sampling probabilities correlate highly with empirically observed choice behavior across the

twelve conditions of the six experiments, in both chimpanzees ($r(10) = .9532$, 95% CI = [.8374, .9871], $p < 10^{-5}$) and humans ($r(10) = .998$, 95% CI = [.993, .999], $p < 10^{-12}$).

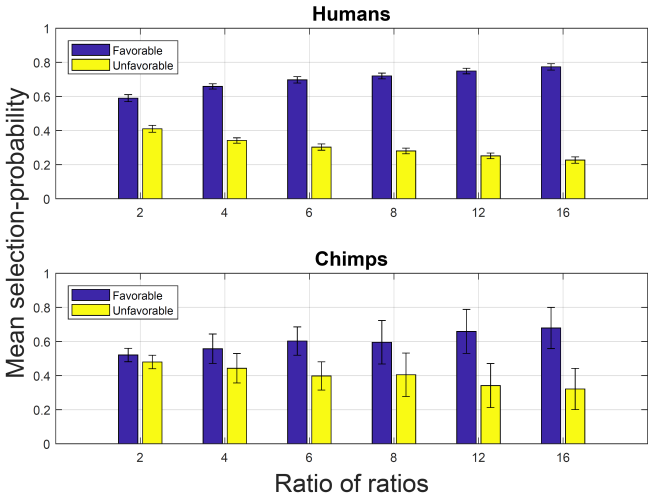

Figure 3. Mean sampling probabilities and SDs for the human and chimpanzee simulated agents.

Simple linear regressions reveal that mean favorable selection probabilities increase linearly as a function of log(RoR) for both simulated species (Table 3), consistent with Eckert et al.'s (2018) empirical results. Equivalently, these probabilities increase logarithmically as a function of raw RoR.

Table 3: Linear regression results for sampling.

Species	β	95% CI	$F(1,4)$	$p <$	R^2_{adj}
Human	.087	[.079,.096]	806	.001	.994
Chimp	.077	[.053,.101]	81	.001	.941

Discussion

Our simulation results show that NPLS provides accurate coverage of the empirical results presented by Eckert et al. (2018), including some success by chimpanzees and even better performance by human controls. Although Eckert et al. (2018) dealt only with selection and not with learning, it is likely that individuals would have to somehow register the probability distributions before being able to use them to guide their selections. To our knowledge, this is the first successful computer simulation consisting of a precise and plausible set of causal mechanisms for both learning and acting on probability distributions that matches empirical data across various age groups and species.

In the NPLS model, the superior performance of humans is due to more accurate learning (due to a smaller score-threshold setting).

Our familiarity with SDCC allowed us to estimate the score-threshold values that could simulate the difference success levels of humans and chimpanzees, illustrated by asymptotic performance of humans (.8 correct) and chimpanzees (.7 correct). Our initial estimate of effective score-threshold values proved to be correct, requiring no data fitting by comprehensive parameter variation.

Depth of learning could involve any of several different underlying factors including, e.g., attention, motivation, and ability. Such factors are effectively summarized by manipulation of the score-threshold parameter, which controls depth of learning. Future empirical and modeling work could perhaps sort out the essential causal factors in more detail.

Alternatively, it is possible that humans are better able to leverage their knowledge to select an advantageous action. Future empirical and computational work could further clarify the precise causal mechanisms for the human advantage on this task.

Like some other researchers contemplating successful performance in probability experiments (Denison & Xu, 2014; McCrink & Birdsall, 2015), Eckert et al. (2018) attributed the success of their participants to use of the Approximate Number System (ANS). The ANS (sometimes called the Analog Magnitude System) is described as a nonverbal system that allows approximate numerical estimation of collections of items at a glance, yielding magnitude values (Carey, Shusterman, Haward, & Distefano, 2017; Dehaene, 2009; Feigenson, Dehaene, & Spelke, 2004; Gallistel & Gelman, 1992).

However, there are several problems with the ANS hypothesis for explaining probabilistic reasoning. First, raw frequencies suffice for probability judgments only when the two are positively correlated. When that positive correlation is disentangled, participants focus on probabilities and ignore raw frequencies (Denison & Xu, 2014; Eckert et al., 2018). Accurate conversion of raw frequencies into probabilities is arguably too difficult for chimpanzees and human infants, both of whom lack explicit division skills. Also, the ANS is known to be relatively imprecise, particularly with large numbers, small Weber fractions, and infants (Carey et al., 2017; Feigenson et al., 2004; Xu & Spelke, 2000). In contrast, recent empirical experiments have revealed that infants are surprisingly precise in probability matching, even in tasks involving large numbers of items (Denison & Xu, 2014). Finally, the ANS has not yet, to our knowledge, been fully implemented as a precise computational mechanism, making it difficult to determine how much power it actually has in explaining probabilistic reasoning, but see Dehaene and Changeux (1993) for a preliminary ANS model.

Eckert et al. (2018) show that both humans' and chimpanzees' choice behavior depends logarithmically on RoR, taking that as evidence for use of Weber's law. However, in Eckert et al.'s experiments, ground-truth probabilities also depend logarithmically on RoR (see Figures 1-2). Because the ground-truth probabilities used here almost exactly add up to 1, the normalized ground-truth probabilities are almost identical to the ground-truth probabilities. Therefore, as an alternative explanation, the observed logarithmic trend in humans' and chimpanzees' choice behavior could be instead attributed to probability-

matching, an empirically well-supported behavioral pattern found in a wide variety of animal species, e.g., bees (Greggers & Menzel, 1993), fish (Behrend & Bitterman, 1961), turtles (Kirk & Bitterman, 1965), humans (Denison & Xu, 2014), and apes and monkeys (De Petrillo & Rosati, 2019; Eckert et al., 2018; Tecwyn, Denison, Messer, & Buchsbaum, 2017). Future experimental work should adjudicate between these two competing explanations of observed choice behavior (i.e., Weber's effect and probability matching), by using a set of normalized ground-truth probabilities that do not depend logarithmically on RoR.

In the meantime, NPLS modeling suggests a relatively simple set of neural mechanisms that explain how a wide variety of learners could internalize and employ probability distributions to guide their action choices, without explicit counting and dividing of large numbers. So far, NPLS predicts and explains probabilistic learning and reasoning in adult chimpanzees, adult humans, and human infants (Shultz & Nobandegani, 2021), all of which approaches a unifying model of probabilistic learning and reasoning across various ages and species.

However, several additional challenges remain to be explored in the study of probabilistic reasoning in non-human animals. One is the extensive, classic literature on the matching law in individual operant conditioning. Individual animals (often a pigeon or a rat) allocated choices in direct proportion to the rewards that the choices provided (Herrnstein, 1970). Coherent interpretation of this literature is currently difficult because these animals sometimes optimized probabilistic choices rather than matching probabilities. This is often referred to as the explore vs. exploit issue, and is subject to active research efforts in several disciplines, including psychology (Gaissmaier & Schooler, 2008; Koehler & James, 2009), computer science (Agrawal & Goyal, 2012), and behavioral economics (Uotila, Maula, Keil, & Zahra, 2009).

Interestingly, NPLS allows for simulating both probability-matching and probability-maximization by modulation of its β parameter (see Eq. 4). Concretely, a gradual increase of the β parameter results in a smooth transition from probability-matching, at one end of a spectrum, to probability-maximization, on the other end. As such, variation in β allows for reconciling two ostensibly distinct behaviors, probability-maximization and probability-matching, with the former being normatively justified while the latter is often interpreted as a sign of irrationality. Future research should investigate whether and how variation in the NPLS β parameter could account for observed discrepancies in probabilistic reasoning.

Another promising comparative literature deals with foraging strategies in animal groups. In ecology, Ideal Free Distribution theory (IFD) describes how the individuals in a group distribute themselves across multiple patches of resources in their environment. Presumably, they would do this in order to minimize resource competition and maximize reproductive fitness. The IFD theory predicts that the number of individual animals that aggregate in various patches is proportional to the richness of resources available in each patch. For this to work, presumably individual animals would have to estimate the probabilities of finding food in each of the patches to decide where to spend their foraging time. As with the operant conditioning literature, there is evidence both for (Dreisig, 1995) and against (Godin & Keenleyside, 1984) this prediction.

Additionally, there are other probabilistic reasoning studies of apes and monkeys that have not yet been modeled (Eckert et al., 2017; Rakoczy et al., 2014; Tecwyn et al., 2017). Careful and extensive interpretations of these three empirical literatures would be required in order to prepare for the quest of achieving a fully unified model and theory of probabilistic reasoning across species. Because decision making is often complicated by context and individual differences, NPLS provides a starting point, and perhaps a solid foundation, for understanding such behavior. For now, it is encouraging to see that a unified model can already predict and explain the learning and use of probability distributions across several infant experiments (Shultz & Nobandegani, 2021) and, as we show here, experiments with human adults and chimpanzees.

Although NPLS is not the only algorithm capable of learning probability distributions, it possesses several features that make it particularly suitable for simulating and thus explaining learning and use of such distributions across ages and species: network construction, learning cessation, and sampling. Network construction allows a clear distinction between learning (via weight adjustment) and development (via hidden-unit recruitment). Learning cessation prevents fruitless attempts to reduce error that cannot be further reduced while allowing efficient and accurate probability matching. And sampling allows reverse inferencing of example generation and action selection from learned probability distributions, with the potential ability to maximally exploit knowledge of probability distributions.

Acknowledgements

This research was supported in part by an operating grant to TRS from the Natural Sciences and Engineering Research Council of Canada. Thanks to Marcel Montrey for acquainting us with the literature on the IFD theory.

References

Agrawal, S., & Goyal, N. (2012). Analysis of thompson sampling for the multi-armed bandit problem. *Journal of Machine Learning Research*, *23*, 1–26.

Baluja, S., & Fahlman, S. (1994). *Reducing network depth in the cascade-correlation learning architecture*. Carnegie Mellon University. Pittsburgh, PA. Retrieved from dtic.mil

Behrend, E., & Bitterman, M. (1961). Probability-Matching in the Fish. *The American Journal of Psychology*, *74*(4), 542–551.

Carey, S., Shusterman, A., Haward, P., & Distefano, R. (2017). Do analog number representations underlie the meanings of young children's verbal numerals? *Cognition*, *168*, 243–255.

Dasgupta, I., Schulz, E., & Gershman, S. (2017). Where do hypotheses come from? *Cognitive Psychology*, *96*, 1–25.

De Petrillo, F., & Rosati, A. (2019). Rhesus macaques use probabilities to predict future events. *Evolution and Human Behavior*, *40*(5), 436–446.

Dehaene, S. (2009). Origins of mathematical intuitions: The case of arithmetic. *Annals of the New York Academy of Sciences*, *1156*, 232–259.

Dehaene, S., & Changeux, J. (1993). Development of elementary numerical abilities: A neuronal model. *Journal of Cognitive Neuroscience*, *5*(4), 390–407.

Denison, S., Reed, C., & Xu, F. (2013). The emergence of probabilistic reasoning in very young infants: evidence from 4.5- and 6-month-olds. *Developmental Psychology*, *49*(2), 243–249.

Denison, S., & Xu, F. (2010). Twelve- to 14-month-old infants can predict single-event probability with large set sizes. *Developmental Science*, *13*(5), 798–803.

Denison, S., & Xu, F. (2014). The origins of probabilistic inference in human infants. *Cognition*, *130*(3), 335–347.

Denison, S., & Xu, F. (2019). Infant Statisticians: The Origins of Reasoning Under Uncertainty. *Perspectives on Psychological Science*, *14*(4), 499–509.

Dreisig, H. (1995). Ideal Free Distributions of Nectar Foraging Bumblebees. *Oikos*, *72*(2), 161–172.

Eckert, J., Call, J., Hermes, J., Herrmann, E., & Rakoczy, H. (2018). Intuitive statistical inferences in chimpanzees and humans follow Weber's law. *Cognition*, *180*, 99–107.

Eckert, J., Rakoczy, H., & Call, J. (2017). Are great apes able to reason from multi-item samples to populations of food items? *American Journal of Primatology*, *79*(10), 1–14.

Feigenson, L., Dehaene, S., & Spelke, E. (2004). Core systems of number. *Trends in Cognitive Sciences*, *8*(7), 307–314.

Gaissmaier, W., & Schooler, L. (2008). The smart potential behind probability matching. *Cognition*, *109*(3), 416–422.

Gallistel, C. R., & Gelman, R. (1992). Preverbal counting and computation. *Cognition*, *44*, 43–74.

Godin, J., & Keenleyside, M. (1984). Foraging on patchily distributed prey by a cichlid fish (Teleostei, Cichlidae): A test of the ideal free distribution theory. *Animal Behaviour*, *32*(1), 120–131.

Greggers, U., & Menzel, R. (1993). Memory dynamics and foraging strategies of honeybees. *Behavioral Ecology and Sociobiology*, *32*(1), 17–29.

Herrnstein, R. (1970). On the law of effect. *J Exp Anal Behav*, *13*(243–266).

Kharratzadeh, M., & Shultz, T. (2016). Neural implementation of probabilistic models of cognition. *Cognitive Systems Research*, *40*, 99–113.

Kirk, K., & Bitterman, M. (1965). Probability-learning by the turtle. *Science*, *148*(3676), 1484–1485.

Koehler, D., & James, G. (2009). Probability matching in choice under uncertainty: Intuition versus deliberation. *Cognition*, *113*(1), 123–127.

McCrink, K., & Birdsall, W. (2015). Numerical abilities and arithmetic in infancy. In R. Cohen Kadosh & A. Dowker (Eds.), *The Oxford Handbook of Numerical Cognition* (pp. 258–274). Oxford, England: Oxford University Press.

Moreno-Bote, R., Knill, D., & Pouget, A. (2011). Bayesian sampling in visual perception. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(30), 12491–12496.

Nobandegani, A., & Shultz, T. (2017). Converting cascade-correlation neural nets into probabilistic generative models. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.), *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (pp. 1029–1034). Austin, TX: Cognitive Science Society.

Nobandegani, A., & Shultz, T. (2018). Example generation under constraints using cascade correlation neural nets. *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*, 2385–2390. Retrieved from http://mindmodeling.org/cogsci2018/papers/0456/0456.pdf

Nobandegani, A., & Shultz, T. (2020). A Resource-Rational, Process-Level Account of the St. Petersburg Paradox. *Topics in Cognitive Science*, *12*, 417–432.

Nobandegani, A., & Shultz, T. (2022). Computational approaches to cognitive development: Bayesian and artificial-neural-network models. In O. Houdé & G. Borst (Eds.), *The Cambridge Handbook of Cognitive Development* (pp. 318–338). Cambridge, UK: Cambridge University Press.

Rakoczy, H., Clüver, A., Saucke, L., Stoffregen, N., Gräbener, A., Migura, J., & Call, J. (2014). Apes are intuitive statisticians. *Cognition*, *131*(1), 60–68.

Savin, C., & Denève, S. (2014). Spatio-temporal representations of uncertainty in spiking neural networks. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems 27* (pp. 2024–2032). Curran Associates, Inc. Retrieved from http://papers.nips.cc/paper/5343-spatio-temporal-representations-of-uncertainty-in-spiking-neural-networks.pdf

Shultz, T. (2003). *Computational developmental psychology.* Cambridge, MA: MIT Press.

Shultz, T. (2017). Constructive artificial neural-network models for cognitive development. In N. Budwig, E. Turiel, & P. D. Zelazo (Eds.), *New Perspectives on Human Development* (pp. 13–26). Cambridge: Cambridge University Press.

Shultz, T., & Doty, E. (2014). Knowing when to quit on unlearnable problems: another step towards autonomous learning. In J. Mayor & P. Gomez (Ed.), *Computational Models of Cognitive Processes* (pp. 211–221). London: World Scientific.

Shultz, T., & Nobandegani, A. (2021). A computational model of infant learning and reasoning with probabilities. *Psychological Review*. https://doi.org/http://dx.doi.org/10.1037/rev0000322

Tecwyn, E., Denison, S., Messer, E., & Buchsbaum, D. (2017). Intuitive probabilistic inference in capuchin monkeys. *Animal Cognition*, *20*(2), 243–256.

Uotila, J., Maula, M., Keil, T., & Zahra, S. (2009). Exploration, exploitation, and financial performance: Analysis of S&P 500 corporations. *Strategic Management Journal*, *30*(2), 221–231.

Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences*, *105*(13), 5012–5015.

Xu, F., & Spelke, E. (2000). Large number discrimination in human infants. *Cognition*, *74*, B1–B11.