2016 Joint IEEE International Conference on Development and Learning
and Epigenetic Robotics (ICDL-EpiRob)
Cergy-Pontoise, Paris, France, Sept 19-22, 2016

Robert J. Powers, David H. Rakison, and David C. Plaut

Learning About Actions in Infancy Without a Rationality Principle*

Abstract— Between 6 and 9 months of age, infants begin to differentiate between the actions of others that are "rational" with respect to goals and those that are not. According to the teleological stance theory, this behavior is underpinned by an innate, naive rationality principle; according to a statistical learning account, experience alone is sufficient to explain this behavior. We present a recurrent neural network implementation of statistical learning that incorporates relevant pre-experimental experience that serves to shape expectations about goal-directed action. This model replicates the looking-time patterns of infants at different points in developmental time, and it demonstrates that a rationality principle is not necessary to account for the extant data.

I. Introduction

Infants are sensitive to the goal-directed nature of observed actions [1]. There is evidence that 9-mo-old infants expect inanimate moving objects to take the most efficient path to a goal [2]. With sufficient perceptual cues (e.g., human features or equifinal variation in motion), this expectation can be observed even in 6.5-mo-olds [3][4]. More complex abilities, such as inferring a goal from an action, emerge by 12 months of age [5]. These findings parallel those from other paradigms involving goal-directed reaching and pointing [6][7].

That infants expect agents to act "rationally" was first observed in a classic experiment by Gergely et al., [8]. Infants were habituated to events in which a small circle (the agent) moved towards a large circle (the goal). In one condition, a barrier blocked the most direct path, so the agent had to "jump" over the barrier. Other infants saw the same jumping motion but without a barrier. The infants habituated to the barrier-jumping event looked significantly less at a new, direct path to the goal compared to a jumping path with no barrier, indicating that they expected agents to move efficiently with respect to goals.

In interpreting these results, Csibra and Gergely [9] claimed that infants were adopting a *teleological stance*. On this view, infants can interpret a broad scope of behaviors by relying on a naive *principle of rationality.* Expectations about the outcomes of actions with respect to goals depend on events being well-formed according to a rationality principle: in the absence of barriers, agents should take the most efficient available path to a goal. This perspective has informed some computational models of goal-directed action perception in both infants and adults [10].

According to a *statistical learning* account, rationality is not directly involved in action-based inferences. Instead, infants expect agents to move "rationally" only to the extent that they have observed those kinds of actions in their environment. In support of this view, 9-mo-old infants in the lab expect objects to engage in non-rational actions if those actions were frequent in habituation [11].

Neural networks provide a natural means of exploring the implications of statistical learning in early development [12]. Van Overwalle [13] developed a simple network model of the Gergely et al. [8] task, but it lacked the ability to learn internal representations and did not take into account the role of infants' pre-experimental experience. We applied a Simple Recurrent Network (SRN) [14], to the same experimental task, including an approximation to infants' visual experiences prior to habituation to enable our model to account for the developmental trajectory observed in empirical studies. Importantly, our model did not include anything resembling a rationality principle as proposed by the teleological stance theory.

II. Methods

A. Network Architecture

Our simulations used a three-layer SRN. The input and output layers consisted of 49 units each, corresponding to a 7x7 spatial grid. The hidden and context layers consisted of 40 units each. There was full, feed-forward connectivity from input and context layers to the hidden layer, and from it to the output layer. Hidden activations were copied to the context layer after each processing step. Unit activations were computed using a standard logistic (sigmoid) function.

B. Training and Testing

The network was trained and tested on events representing simple motion of an agent in the presence of other, static objects (the barrier and the goal). To simulate infants' relevant experience outside of the laboratory, a set of 5000 training events was generated. Each event displayed agent-to-goal motion horizontally across the input grid, where vertical positions of both agent and goal were varied. A barrier with randomly chosen vertical extent was positioned in one of the three central columns of the input grid in 10% of these events. Agent motion was always as direct as possible. On most trials, the motion path was a straight line to the goal; motion around barriers was represented by two straight paths, with a discontinuity at the "turn" around the barrier.

Both training and test events were represented as sequences of 10 discrete states of the agent's motion. Continuous spatial locations and trajectories were approximated by linear interpolation across contiguous units. The network's task was to predict only the agent's location 3 time-steps into the future. A total of 1200 weight updates were performed using back-propagation, with cross-entropy error derivatives accumulated over batches of 100 randomly

*Research supported by Award Number T32GM081760 from the National Institute of General Medical Sciences. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIHMS or the NIH. RJP is a graduate student with Carnegie Mellon University and the Center for the Neural Basis of Cognition, Pittsburgh, PA 15213 USA (412-268-3151; e-mail: powers@cmu.edu). DHR and DCP are in the Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: rakison@andrew.cmu.edu; plaut@cmu.edu).

selected training events. Habituation and test events designed to match the Gergely et al. [8] study were performed at fixed intervals during training to simulate testing at different ages. After testing, the weights were set to their pre-habituation values to prevent earlier tests from interfering with learning. Habituation was simulated by reducing error to half of the starting error on each set of stimuli.

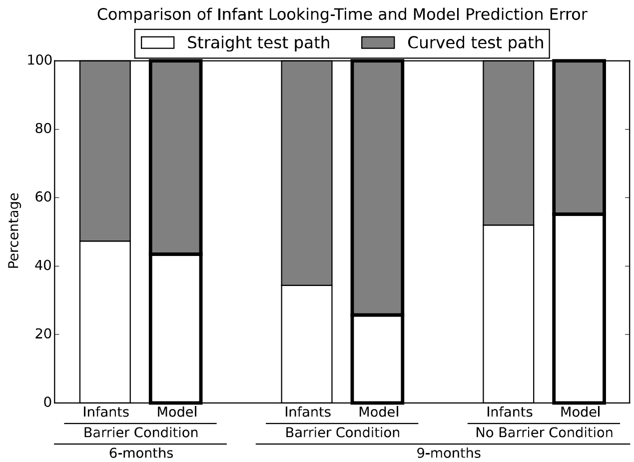

Fig. 1 Model performance (error) compared to infant looking-time at two time points. The shaded portion of each bar represents the relative percentage of real or simulated looking at the curved test path, and the light portion represents looking at the straight test path. Infant data were reported in [2]. The model is compared to 6-mo-olds at 400 weight updates of pre-experimental training, and to 9-mo-olds at 1200 updates.

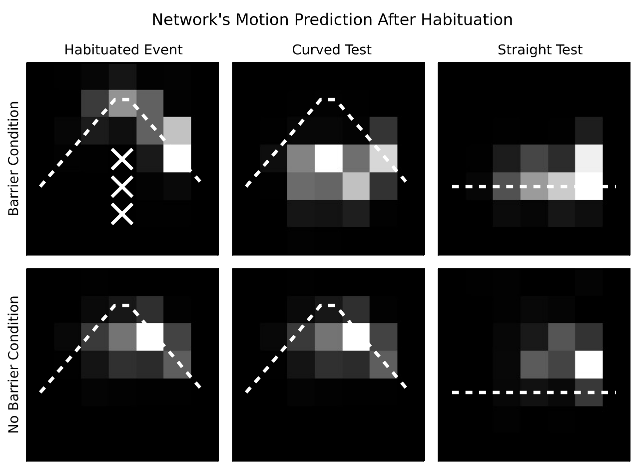

Fig. 2 The model's motion predictions when habituated with a barrier (top row) and without (bottom row). Each cell represents an output unit, where lighter colors indicate higher activation. Output vectors at each time step were summed to produce a trace of the full predicted path. White dashed lines indicate the agent's actual trajectory, as presented in the input, and white Xs indicate barrier locations. The agent moved from left to right in each event shown, although habituation also included a right-to-left version.

III. Results

The network's motion prediction error was used to approximate infant looking time, under the assumption that error reflects a violation of expectation. The model shows the same qualitative pattern observed in [8] and [2] (Figure 1). As pre-experimental experience increases, the network looks longer at (i.e. produces more error for) the curved path when it was habituated with a barrier, but not when it was habituated to the same curved path without a barrier.

The effects of habituation on prediction after 1200 updates of pre-training are shown in Figure 2. The model predicts similar motion for both test events within each condition. When habituated with a barrier, the habituated event is consistent with prior experience. In this case, the weights remain more similar to their pre-habituation values, and an expectation for straight motion—the more frequent motion in pre-training—is preserved. When habituated without a barrier, the event is inconsistent with prior experience, and the weights are altered so that the prediction more closely resembles the curved path. This makes the network better at predicting curved motion at the expense of predicting straight motion.

These findings support a statistical learning account of infants' action interpretation, and suggest that a naive theory of rational action is unnecessary to explain previous findings. One limitation of this model is that predictions are limited to future actions of the agent given the current state of the environment. Future work should extend this model to account for related inferences, such as inferring the goal state from an observed action [5][6][7].

References

[1] J. Sommerville, M. Upshaw and J. Loucks, “The Nature of Goal-Directed Action Representations in Infancy”, Rational Constructivism in Cognitive Development, pp. 351-387, 2012.

[2] G. Csibra, G. Gergely, S. Bíró, O. Koós and M. Brockbank, “Goal attribution without agency cues: the perception of ‘pure reason’ in infancy”, Cognition, vol. 72, no. 3, pp. 237-267, 1999.

[3] G. Csibra, “Goal attribution to inanimate agents by 6.5-month-old infants”, Cognition, vol. 107, no. 2, pp. 705-717, 2008.

[4] K. Kamewari, M. Kato, T. Kanda, H. Ishiguro and K. Hiraki, “Six-and-a-half-month-old children positively attribute goals to human action and to humanoid-robot motion”, Cognitive Development, vol. 20, no. 2, pp. 303-320, 2005.

[5] G. Csibra, S. Bıró, O. Koós, and G. Gergely. “One-year-old infants use teleological representations of actions productively”, Cognitive Science, vol. 27, no. 1, pp. 111-133, 2003.

[6] I. Király, B. Jovanovic, W. Prinz, G. Aschersleben and G. Gergely, “The early origins of goal attribution in infancy”, Consciousness and Cognition, vol. 12, no. 4, pp. 752-769, 2003.

[7] A. Woodward and J. Guajardo, “Infants’ understanding of the point gesture as an object-directed action”, Cognitive Development, vol. 17, no. 1, pp. 1061-1084, 2002.

[8] G. Gergely, Z. Nádasdy, G. Csibra and S. Bíró, “Taking the intentional stance at 12 months of age”, Cognition, vol. 56, no. 2, pp. 165-193, 1995.

[9] G. Csibra and G. Gergely, “‘Obsessed with goals’: Functions and mechanisms of teleological interpretation of actions in humans”, Acta Psychologica, vol. 124, no. 1, pp. 60-78, 2007.

[10] C. Baker, R. Saxe and J. Tenenbaum, “Action understanding as inverse planning”, Cognition, vol. 113, no. 3, pp. 329-349, 2009.

[11] M. Paulus, S. Hunnius, C. van Wijngaarden, S. Vrins, I. van Rooij and H. Bekkering, “The role of frequency information and teleological reasoning in infants' and adults' action prediction.”, Developmental Psychology, vol. 47, no. 4, pp. 976-983, 2011.

[12] Y. Yermolayeva and D. Rakison, “Connectionist modeling of developmental changes in infancy: Approaches, challenges, and contributions.”, Psychological Bulletin, vol. 140, no. 1, pp. 224-255, 2014.

[13] F. Van Overwalle, “Infants' teleological and belief inference: A recurrent connectionist approach to their minimal representational and computational requirements”, NeuroImage, vol. 52, no. 3, pp. 1095-1108, 2010.

[14] J. Elman, “Finding structure in time”, Cognitive Science, vol. 14, no. 2, pp. 179-211, 1990.