TOWARDS MODELING THE DEVELOPMENTAL VARIABILITY OF HUMAN ATTENTION

Kuno Kim[1], **Megumi Sano**[1], **Julian De Freitas**[5], **Daniel L. K. Yamins**[1,2,3,*] **& Nick Haber**[4,*]
Departments of Computer Science[1], Psychology[2], Wu Tsai Neurosciences Institute[3], and
Graduate School of Education[4], Stanford University
Department of Psychology[5], Harvard University
{kunokim, megsano}@stanford.edu

ABSTRACT

Children exhibit extraordinary exploratory behaviors hypothesized to contribute to the building of models of their world. Harnessing this capacity in artificial systems promises not only more flexible technology but also cognitive models of the developmental processes we seek to mimic. Yet not all children learn the same way, and for instance children with autism exhibit characteristically different exploratory strategies early in life. What if we could, by developing artificial systems that learn through exploration, model not only typically development, but all its variations? In this work, we present a preliminary analysis of curiosity-driven agents in social environments that establishes links between early behavior and later acuity, with implications for the future of both diagnostics and personalized learning.

1 INTRODUCTION

Human infants exhibit a wide range of interesting, apparently spontaneous, visuo-motor behaviors — including navigating their environment, seeking out and attending to novel objects, and engaging physically with these objects in novel and surprising ways (Fantz, 1964; Twomey & Westermann, 2017; Hurley et al., 2010; Hurley & Oakes, 2015; Goupil et al., 2016; Begus et al., 2014; Gopnik et al., 2009). In short, young children are excellent at playing — "scientists in the crib" (Gopnik et al., 2009) who intentionally create events that are both fun and greatly informative for driving the self-supervised learning of sensorimotor and social planning capacities (Fantz, 1964; Sokolov, 1963; Goupil et al., 2016; Begus et al., 2014; Kidd et al., 2012). Harnessing this sort of capacity in artificial systems promises not only more flexible learning technologies but also cognitive models that will further elucidate early childhood learning.

Evidence suggests that Autism Spectrum Disorder (ASD) children exhibit characteristically *different* exploratory learning behaviors. Children with autism exhibit atypical, uncreative object play, (Beyer & Gammeltoft, 2000; Rettig, 1994), impaired predictive capacity (Sinha et al., 2014), lower facial gaze and mutual attention (Shic et al., 2014; Jones & Klin, 2013; Moriuchi et al., 2016), and abnormalities of sensory perception (Robertson & Baron-Cohen, 2017; CE et al., 2013). What if we could model not only typical development, but the full diversity of human developmental variability (Fig. 1)?

In this work, we analyze a "population" of agents put in an environment meant to loosely represent early childhood learning in social environments: stimuli look either animate or inanimate, and inanimate stimuli vary wildly (from static ones, to dynamic but predictable ones, to dynamic and unpredictable ones), and the agent simply looks about. To decide what to look at, the agent is curious (Schmidhuber, 2010; Oudeyer et al., 2007) — intrinsically motivated to take action as it tries to build a world model of its environment. Within this population, we take implementation differences (specifically, in choice of the agent's intrinsic motivation) to represent latent factors that drive developmental differences. By observing both the agent's behavior and downstream world modeling capacity, we establish a predictive model that, if translated to simulations that have the fidelity to capture early childhood learning, holds exciting implications for diagnostics, therapeutics, and personalized learning.

2 MODELING SOCIAL ATTENTION

Environment To simplify but faithfully capture key aspects of the algorithmic challenges children face, we work with a 3D virtual environment (Fig. 2d). Within the environment there are two main agent types:

[*]equal contribution

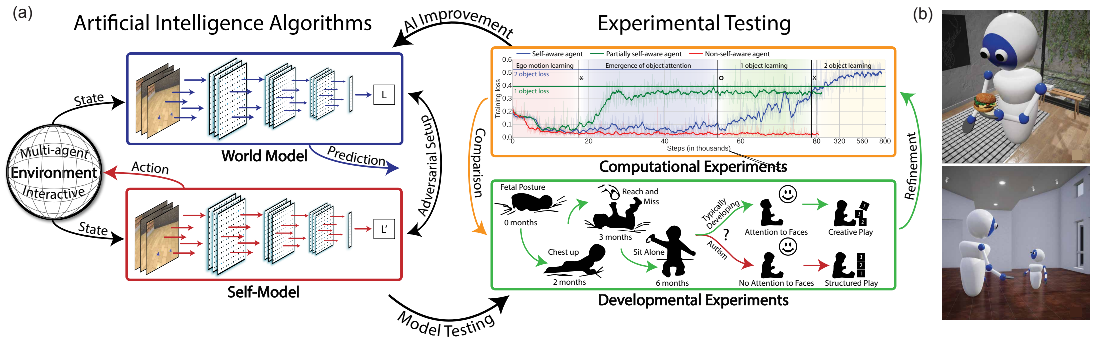

Figure 1: **Modeling human development.** (a) Our overall goal is to build a computationally quantitative model of the learning principles of early childhood, both for cases of typical and variable developmental trajectories. (b) We develop the learning principles by which they operate within virual environments.

a single "infant-agent", and various *external agents*. Just as very young human babies are not self-mobile but can control their gaze to visually explore their surroundings, our infant-agent is represented by an avatar fixed at the center of the room but with the ability swivel around, obtaining partial observations of whatever is in view at the given moment. The **external agents** are spherical avatars that each act under various hard-coded policies embodying abstract versions of the behaviors of real-world stimuli, both inanimate and animate. We experiment with external agent behaviors of increasing complexity and animacy, including static (no motion), periodic, noise, object-reaching, chasing, playing "peekaboo", and mimicry. See `https://bit.ly/2uf7lEY` for video descriptions of the environment and external agent behaviors. The task of the infant-agent is to learn to predict the behaviors of the external agents. Since external agents are devoid of surface features, the curious agent must understand behaviors based on spatiotemporal kinematics alone.

Learning predictive models of other agents. The infant-agent's neural network consists of two components: an *agent-interaction-centric world model*, which seeks to learns to predict dynamics of external agents, and a *curiosity-driven controller* which uses a novel variant of progress curiosity Schmidhuber (2010) to choose swivel actions (e.g. allocate attentional resources) that make world-model learning more effective. *Agent-centric World Model:* (Figure 2a) Our infant-agent learns to predict the dynamics of its environment via an agent-interaction-centric world model ω_θ. ω_θ consists of an ensemble of component networks $\{\omega_{\theta^k}\}_{k=1}^{N_{\text{cc}}}$ where each ω_{θ^k} independently predicts the forward dynamics separately, for each minimal group of interacting external agent(s). This agent-interaction-centric world model differs from a standard "joint" (non-agent-centric) model in that it allocates the parameters and learning gradients in a causally disentangled fashion. For example, one external agent reaching toward and pushing around static objects is allocated a separate world-model component than another external agent playing peekaboo with the infant-agent. If multiple external agents are causally interacting with each other (e.g. one agent chasing or mimicking another), they are allocated a single joint component of the world model. In this work, we manually feed the infant-agent knowledge of the causal graph of external agents; in future work estimating this from observations is a key goal. *Progress-driven Controller:* (Figure 2b) We propose γ-**Progress**, a scalable progress-based curiosity signal which approximates learning progress by the difference in the

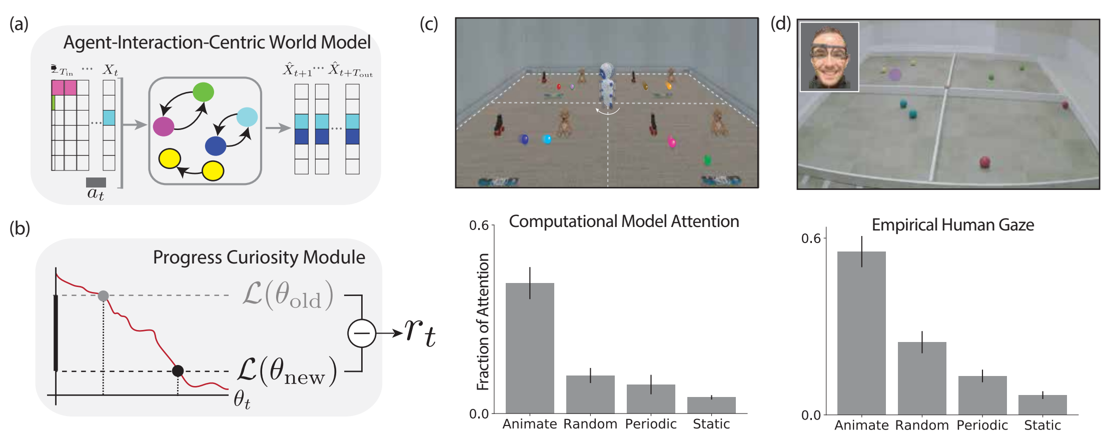

Figure 2: **Modeling social attention.** Proposed (a) agent-interaction-centric world model and (b) curiosity signal to facilitate learning in social environments. In a pilot study, we compared (c) simulation of our artificial agents with (d) data from a human experiment and observed similar attention patterns.

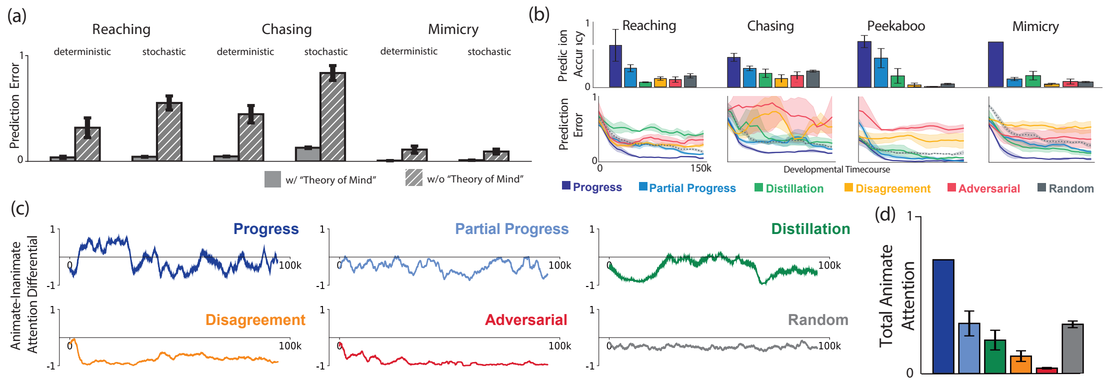

Figure 3: **Computational Models of “developmental variability.”** By varying (a) world model architecture, we see differences in end-state external-agent prediction performance. By varying curiosity signal, we see (b) difference in sample complexity and end-state performance, as well as in (c) behavior timecourses, specifically on an animate-inanimate attention differential. (d) Aggregated differences in animate attention.

losses of an old model and a new model. The old model weights, θ_{old}, lag behind those of the new model, θ_{new}, with a simple update rule: $\theta_{old} \leftarrow \gamma\theta_{old} + (1-\gamma)\theta_{new}$, where γ is scalar mixing constant. The curiosity reward is:

$$R = \mathcal{L}(\theta_{new}) - \mathcal{L}(\theta_{old}) \tag{1}$$

Our controller π_ϕ follows an ϵ-$greedy$ sampling scheme with respect to a Q-function Q_ϕ trained with the curiosity reward in Eq. 1 and updated with the DQN Mnih et al. (2013) learning algorithm.

Model Evaluation. We trained our proposed model on variety of situations with different compositions of external agent behaviors. Looking at the trained model’s behavior, we find that it appears to be able to learn to predict the behavior of a wide array of external agents, across the spectrum of animacy. See `https://bit.ly/31vg7v1` for visualizations of our model’s predictions.

We believe that the key reason that the infant-agent equipped with this two-component active-learning world model was able to capture a wide range of external agent behaviors is that it learned to properly allocate its attention, e.g. spending more time focusing on complex animate external agents, as seen in Figure 2c. This increased animate-inanimate attention differential corresponds to a characteristic attentional “bump” that occurs early as the γ-Progress curious agent focuses on animate external agents quickly before eventually “losing interest” as prediction accuracy is achieved.

Comparisons to human attentional allocation. Our model is a very rudimentary hypothesis for how babies allocate attention in social settings. To begin to examine the extent to which this hypothesis is accurate, we ran a simple pilot human subject experiment (Figure 2d) in which we conveyed static, periodic, animate, and noise stimuli to twelve human participants via spherical robots moving along a mat, while measuring patterns of attention via a mobile eye tracker. We find average fixation proportions favoring the animate stimuli, just as in the computational model. We also find a similar ordering of aggregate attentional fixation across multiple kinds of stimuli. In follow-up work, we aim to make a finer model comparison to the behavior of humans shown a diverse array of stimuli.

3 Sources of Variability

What happens to external agent prediction performance and animate attention when we vary components in the model? To evaluate the effect of varying the structure of the agent-interaction-centric architecture, independently of controller choice, we produce datasets for offline training for each task and train the world model to convergence. We compare the performance of the agent-interaction-centric world model to a parameter-matched *joint* LSTM architecture that takes as input and predicts all external agents together, with no agent-centric disentangling. As seen in Figure 3a, the agent-interaction-centric (disentangled) architecture significantly outperforms the entangled model on final external-agent prediction.

To evaluate dependence on curiosity signal, we measure both end-state prediction performance and sample complexity (rate of reduction in loss with respect to the number of environment interactions). We compare performance of γ-**Progress** to a range of potential variants, including: δ-**Progress** (Achiam & Sastry, 2017), **RND** (Burda et al., 2018), **Disagreement** (Pathak et al., 2019), **Adversarial** (Stadie et al., 2015), and a simple **Random** policy. Fig. 3b shows end performance (first row) and sample complexity (second row). γ-Progress has higher end performance on all baselines and all tasks. γ-Progress has lower sample complexity than Disagreement, Adversarial, and Random baselines on all behaviors, and RND and δ-Progress on all but one behavior, tying on stochastic chasing.

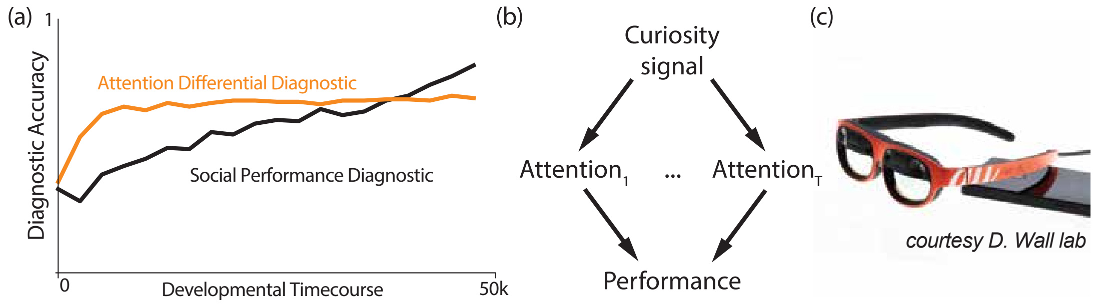

Figure 4: **Early diagnostic analysis** (a) Attention-differential diagnostic achieves better diagnostic accuracy than direct social performance measurement in critical early phase — in computational modeling experiment. (b) Factor analysis hypothesis: curiosity signal determines attention, which determines final performance. (c) Lightweight wearable suitable for measuring response to diagnostic stimuli.

Baselines display two distinct modes in failing to exhibit animate attention (Fig 3c). The first is *attentional indifference*, in which it finds no particular external agent interesting. The second failure mode is *white noise fixation*, where the observer is captivated by the noise external agents. Non-progress-based curiosity signals exhibited both kinds of failure mode but were more dominated by white noise. RND, a novelty measure, exhibited both types of failure at a lower rate. δ-Progress, a direct information gain measure, had no white noise failure but frequently led to attentional indifference as the new and old world model, separated by a fixed time difference, were often too similar to generate a useful curiosity signal. We also found (data not shown) that γ-Progress exhibited indifference when γ was too small, but robustly succeeded across all behaviors for sufficiently large γ. Overall, emergence of animate attention is highly correlated with prediction performance, suggesting that γ-Progress succeeds because its improved ability to flexibly estimate information gain allows it to focus on more informative interactions.

Obviously this is a very preliminary investigation of sources of variability in social prediction performance and attentional allocation. Our hypothesis is that variability in a computational model, whether this one or some future better model, will describe the underlying mechanisms behind social behavior variability.

4 Towards model-based diagnostics.

As mentioned above, Autism Spectrum Disorder is characterized by both differences in low-level attention (Jones & Klin, 2013; Constantino et al., 2017) and high-level social acuity (Hus & Lord, 2014). Yet currently, ASD diagnosis is done by expert clinicians, using only observations of high-level behaviors (Hus & Lord, 2014). This method is subjective, expensive, and too late — the average diagnosis comes after 4 years of age, often preventing interventions during a critical period of development.

Motivated by these observations, we sought to determine, using computational models whether the easily-measurable low-level attention could be used as an early indicator of high-level social prediction performance. In this interpretation, the attention a readily observable behavioral metric, and performance represents some more difficult-to-obtain measure of social acuity. Variation in curiosity signal would, in this account, be a latent correlate of developmental variability. To perform an early indicator analysis, we thus train two statistical regression models to predict the final end-state social performance of each variant of our computational agents: (1) $\text{PERF}_{\leq T}$, which takes performance before time T as input, and (2) $\text{ATT}_{\leq T}$, which takes attention before time T as input. As seen in Figure 4a, $\text{ATT}_{\leq T}$ is an effective predictor of late social performance, and in fact, throughout most of the timecourse, a more accurate indicator than direct measurement of early-stage model performance itself. The overall situation is conveyed by the factor diagram Figure 4b.

This analysis, which establishes a link between early behaviors and downstream social acuity differences via latent factors, is obviously just a toy model. Its translation into higher-fidelity simulations of early childhood learning, however, holds exciting possibilities in diagnostics, therapeutics, and personalized learning. If variations in simulated agents can represent the learning process of diverse human populations, then such a link enables us to search, entirely in simulation, for stimuli that elicit easily measurable behavioral responses that differentiate between underlying factors of variation which in turn predict differences later in life. Such diagnostics could be performed cheaply using lightweight AR/VR devices (such as those shown in Figure 4c) that will soon emerge. In the long run, if computational modeling approaches to developmental variability are able to correctly describe patterns of behavior across the population and produce effective diagnostics, a natural extension will be the development of model-based therapeutics tested in simulation.

REFERENCES

Joshua Achiam and Shankar Sastry. Surprise-based intrinsic motivation for deep reinforcement learning. *arXiv preprint arXiv:1703.01732*, 2017.

Katarina Begus, Teodora Gliga, and Victoria Southgate. Infants learn what they want to learn: Responding to infant pointing leads to superior learning. *PLOS ONE*, 9(10):1–4, 10 2014.

Jannik Beyer and Lone Gammeltoft. *Autism and play*. Jessica Kingsley Pub, 2000.

Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation, 2018.

Robertson CE, Kravitz DJ, Freyberg J, Baron-Cohen S, and Baker CI. Slower rate of binocular rivalry in autism. *The Journal of Neuroscience*, 33(43):16983–16991, 2013.

J. N. Constantino, S. Kennon-McGill, C. Weichselbaum, N. Marrus, A. Haider, A. L. Glowinski, S. Gillespie, C. Klaiman, A. Klin, and W. Jones. Infant viewing of social scenes is under genetic control and is atypical in autism. *Nature*, 547(7663):340–344, Jul 2017.

Robert L. Fantz. Visual experience in infants: Decreased attention to familiar patterns relative to novel ones. *Science*, 146(3644):668–670, 1964.

A. Gopnik, A.N. Meltzoff, and P.K. Kuhl. *The Scientist In The Crib: Minds, Brains, And How Children Learn*. HarperCollins, 2009.

Louise Goupil, Margaux Romand-Monnier, and Sid Kouider. Infants ask for help when they know they don't know. *Proceedings of the National Academy of Sciences*, 113(13), 2016.

K. B. Hurley and L. M. Oakes. Experience and distribution of attention: Pet exposure and infants' scanning of animal images. *J Cogn Dev*, 16(1):11–30, Jan 2015.

K. B. Hurley, K. A. Kovack-Lesh, and L. M. Oakes. The influence of pets on infants' processing of cat and dog images. *Infant Behav Dev*, 33(4):619–628, Dec 2010.

Vanessa Hus and Catherine Lord. The autism diagnostic observation schedule, module 4: revised algorithm and standardized severity scores. *Journal of autism and developmental disorders*, 44(8): 1996–2012, 2014.

Warren Jones and Ami Klin. Attention to eyes is present but in decline in 2–6-month-old infants later diagnosed with autism. *Nature*, 504(7480):427–431, 2013.

Celeste Kidd, Steven T. Piantadosi, and Richard N. Aslin. The goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLOS ONE*, 7(5):1–8, 05 2012.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

Jennifer M Moriuchi, Ami Klin, and Warren Jones. Mechanisms of diminished attention to eyes in autism. *American Journal of Psychiatry*, 174(1):26–35, 2016.

Pierre-Yves Oudeyer, Frdric Kaplan, and Verena V Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2):265–286, 2007.

Deepak Pathak, Dhiraj Gandhi, and Abhinav Gupta. Self-supervised exploration via disagreement. *arXiv:1906.04161*, 2019.

Michael A Rettig. Play behaviors of young children with autism: Characteristics and interventions. *Focus on Autistic Behavior*, 9(5):1–6, 1994.

Caroline E Robertson and Simon Baron-Cohen. Sensory perception in autism. *Nature Reviews Neuroscience*, 18(11):671, 2017.

J. Schmidhuber. Formal theory of creativity, fun, and intrinsic motivation (1990 – 2010). *IEEE Transactions on Autonomous Mental Development*, 2(3):230–247, Sept 2010.

F. Shic, S. Macari, and K. Chawarska. Speech disturbs face scanning in 6-month-old infants who develop autism spectrum disorder. *Biol. Psychiatry*, 75(3):231–237, Feb 2014.

Pawan Sinha, Margaret M Kjelgaard, Tapan K Gandhi, Kleovoulos Tsourides, Annie L Cardinaux, Dimitrios Pantazis, Sidney P Diamond, and Richard M Held. Autism as a disorder of prediction. *Proceedings of the National Academy of Sciences*, 111(42):15220–15225, 2014.

E.N. Sokolov. *Perception and the conditioned reflex*. Pergamon Press, 1963.

Bradly Stadie, Sergey Levine, and Pieter Abbeel. Incentivizing exploration in reinforcement learning with deep predictive models. *arXiv preprint arXiv:1507.00814*, 2015.

K. E. Twomey and G. Westermann. Curiosity-based learning in infants: a neurocomputational approach. *Dev Sci*, Oct 2017.

Appendix

Here we provide details of the early indicator analysis and a regression of what factors (curiosity signal, architecture, external agent behavior) best predict animate/inanimate attention ratios.

.1 DETAILS OF EARLY INDICATOR ANALYSIS

We look to predict final performance P_{final} of a given agent, which we take to be the average of the final four validation runs. To make the modeling problem simple, we discretize this into a classification task by dividing validation performance into 3 equal-sized classes ("high", "medium", and "low", computed separately for each external agent behavior), intuitively chosen to reflect performance around, at, and below that of random policy.

We consider two predictive models of final performance, one that takes as input early attention of the agent, and the other, early performance. Early performance may be quantified simply: given time T ("diagnostic age") during training, let $P_{\leq T}$ be the vector containing all validation losses measured up to time T. Early attention, however, is very high-dimensional, so we must make a dimensionality-reducing choice in order to tractably model with our modest sample size. Hence, we "bucket" average. Given choice of integer B, let

$$A_{\leq T,B} = (f^{\text{anim}}_{0:\frac{T}{B}}, f^{\text{rand}}_{0:\frac{T}{B}}, f^{\text{anim}}_{\frac{T}{B}:\frac{2T}{B}}, f^{\text{rand}}_{\frac{T}{B}:\frac{2T}{B}}, \ldots f^{\text{anim}}_{\frac{(B-1)T}{B}:T}, f^{\text{rand}}_{\frac{(B-1)T}{B}:T}), \tag{2}$$

where $f^{\text{anim}}_{a:b}$ and $f^{\text{rand}}_{a:b}$ are the fraction of the time $t = a$ and $t = b$ spent looking at the animate external agent and random external agents respectively (so $A_{\leq T,B}$ is the attentional trajectory up to time T discretized into B buckets).

Finally, both models must have knowledge of the external agent behavior to which the agent is exposed — we expect this to both have an effect on attention as well as the meaning of early performance and expected final performance as a result. Let χ_{BHR} be the one-hot encoding of which external animate agent behavior is shown.

We then consider models

1. PERF$_{\leq T}$, which takes as input $P_{\leq T}$ and χ_{BHR}, and
2. ATT$_{\leq T}$, which takes as input $A_{\leq T,B}$ and χ_{BHR}.

Figure **??**b shows the plot of PERF$_{\leq T}$ and ATT$_{\leq T}$ accuracy as T varies. We see that, up to a point, ATT$_{\leq T}$ makes a better predictor of final performance, and then PERF$_{\leq T}$ dominates. This confirms the intuition that attention patterns precede performance improvements. Intuitively, early attention predicts performance by being able to predict the sort of curiosity signal the agent is using, which predicts the full timecourse of attention, which in turn predicts performance.