

A Unified Account of Gaze Following

Hector Jasso, Jochen Triesch, Gedeon Deák, and Joshua M. Lewis

Abstract—Gaze following, the ability to redirect one’s visual attention to look at what another person is seeing, is foundational for imitation, word learning, and theory-of-mind. Previous theories have suggested that the development of gaze following in human infants is the product of a basic gaze following mechanism, plus the gradual incorporation of several distinct new mechanisms that improve the skill, such as spatial inference, and the ability to use eye direction information as well as head direction. In this paper, we offer an alternative explanation based on a single learning mechanism. From a starting state with no knowledge of the implications of another organism’s gaze direction, our model learns to follow gaze by being placed in a simulated environment where an adult caregiver looks around at objects. Our infant model matches the development of gaze following in human infants as measured in key experiments that we replicate and analyze in detail.

Index Terms—Adaptive systems, artificial intelligence, autonomous mental development, behavioral science, cognition, cognitive science, computational and artificial intelligence, cybernetics, emergent phenomena, intelligent systems, learning systems, multiagent systems, systems, man, and cybernetics.

I. INTRODUCTION

DURING their first two years of life infants transform their earlier dyadic interactions with adults into new triadic interactions that also include objects [5], [65]. During triadic interactions, infants and adults sometimes simultaneously attend to the same object. Such episodes of “joint attention” [54] are a central component of many important social skills, including protodeclarative and protoimperative pointing [3], [4], social referencing [25], early vocabulary formation [64], early language development [8], and inferring other people’s beliefs, intentions, and plans, i.e., “theory-of-mind” skills [56]. All these skills are considered crucial for the infant’s entry into the social world, and for learning to interact and communicate with others in an increasingly sophisticated manner [52].

Scaife and Bruner first studied joint attention skills in a controlled experimental context [57]. They attempted to find the age of onset of gaze following, the ability to redirect one’s visual

attention to look at what another person is seeing. In their paradigm, an infant and experimenter sat facing each other. After the experimenter got the infant’s attention, he/she turned to either the right or the left, as if looking at an object. The researchers observed whether the infant then turned in the same direction as the experimenter. Starting at 8 to 10 months, they noted, infants turned in the same direction more often than chance. Moreover, age groups from 2 to 14 months showed gradually increasing reliability. This influential study established that infants are not completely egocentric, which was a common view at the time.

Other researchers adopted and modified Scaife and Bruner’s experimental paradigm to investigate more complex aspects of gaze following, and determine the mechanisms behind its development. For example, Butterworth recognized that it is a challenge for infants to ignore nearby targets that are salient but are nevertheless not the object of the adult’s gaze. Butterworth hypothesized that infants must eventually acquire a “geometric compensation mechanism” in order to extrapolate a line from the adult’s head and eye direction cues to a specific region of interest, with the target (salient or not) falling somewhere on this line [11], [13]. Another example is gaze following based on eye direction cues, which has been attributed to the acquisition of a new “shared attention mechanism.” By this, the infant understands that other people’s attentive states are determined not by head direction per se but by the direct line of sight from the eyes themselves [2]. A related idea is that this new representational mechanism allows the infant to understand other people as intentional agents [51]. These, and similar accounts, stipulate discrete changes in infants’ understanding of other people as entities with goals and beliefs. If these accounts are valid, however, the task remains to explicate the nature of infants’ understanding (e.g., the kinds of representations that are needed), to specify how development occurs, and to stipulate the role of experience in the environment as a factor in development.

In this paper, we offer a unified account of gaze following. Our account is based on a novel computational “infant model”, which we describe and test. The infant model learns gaze following as it interacts with its environment, deciding where to look next based on the location of objects as well as the apparent head and eye direction of an adult who is also looking at objects in the shared environment. First, the infant model learns to look at salient objects, and with time it learns to take into account the adult’s head/eye direction cues to direct its looking behavior. This basic process brings with the progression of gaze following behaviors described above. The model is based on a hypothesized “basic set”, which is a small number of generic phenotypes such as visual preferences, attention-shifting, habituation, and reinforcement learning [19],[61] plus a natural environment where other agents tend to look at interesting things. The core theory is that this basic set alone, without any specialized modules, is sufficient to explain the emergence of basic

Manuscript received September 28, 2011; revised March 30, 2012; accepted June 25, 2012. Date of publication July 18, 2012; date of current version December 07, 2012.

H. Jasso is with the Department of Computer Science and Engineering, University of California San Diego, La Jolla, CA 92093–0404 USA (e-mail: hjasso@cs.ucsd.edu).

J. Triesch is with the Frankfurt Institute for Advanced Studies, Johann Wolfgang Goethe University, Frankfurt 60438, Germany (e-mail: triesch@fiас.uni-frankfurt-de).

G. Deák and J. M. Lewis are with the Department of Cognitive Science, University of California San Diego, La Jolla, CA 92093–0515 USA (e-mail: deak@cogsci.ucsd.edu; josh@cogsci.ucsd.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.org>.

Digital Object Identifier 10.1109/TAMD.2012.2208640

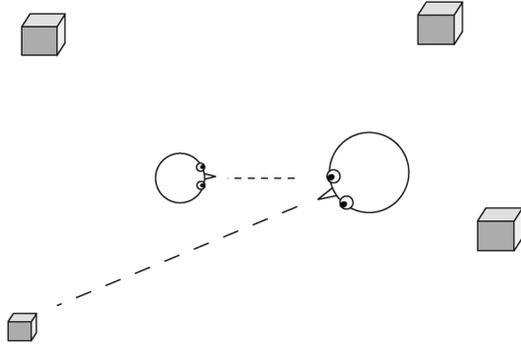


Fig. 1. Modeling the environment. Infant model and caregiver sit facing each other, with objects placed around them. Here, the caregiver has turned left to look at an object, while the infant model looks at the caregiver.

gaze following behavior [24].^{1,2} Overall, our approach is well suited to answer important questions that have until now begged explanation, such as “What is the nature of the mechanisms of gaze following?” and “What factors guide its developmental schedule?”

II. METHODS

In this section we describe our model in detail. First, we describe its three primary components: 1) a spatial representation of the environment; 2) the infant model’s visual system; and 3) a biologically plausible reinforcement learning algorithm [61], [19] used by the infant model to decide where to look next. After discussing these components, we review the training and testing procedures used.

A. Learning Environment

The environment is simulated as a 2-D plane containing the infant model, a simulated caregiver, and a number of objects. The infant model and caregiver are in fixed positions facing each other with a 40-cm separation between them (see Fig. 1), similar to some infant–adult interactions. Objects can be located anywhere on the plane except in the same location as the infant model or caregiver. The caregiver is not always present but sometimes “leaves the room.” Time is discretized into 1-s intervals.

The room contains N_o objects, where N_o is drawn from a geometric probability distribution with average \bar{N}_o . These objects are located randomly around the infant model with distances taken from a radially symmetric normal probability distribution with standard deviation of σ_o . The visual saliency Φ_{o_i}

¹Our model integrates ideas from two previous gaze following models using the basic set hypothesis. In the first model [14], [67] the environment is a set of regions occupied by the infant, the caregiver, and targets. But since the regions are discrete and hold no spatial relationship among each other, the model cannot be used to replicate spatial aspects of Butterworth’s experiments. The second model [45] used a spatial representation of the environment using a body-centered coordinate system. A Hebbian-like learning rule was used to strengthen the connections between visual inputs and the locations where visual saliency is encountered as a result of actions. Our model combines the reinforcement learning approach of the first model with the modeling of spatial aspects of the second. This has allowed us to replicate a wealth of gaze following phenomena, as discussed throughout the paper.

²Earlier versions of this model were presented in [37],[39], and [66], with newer versions successfully replicating results for more experimental setups.

for each object i has a value taken from an exponential probability with average $\bar{\Phi}_o$. The configuration of objects changes after a number of time steps taken from a geometric probability function with average \bar{T}_{objects} , with a new number of objects and their locations and saliencies drawn from probabilities as described earlier. This simulates a dynamic world of varied-saliency objects that the infant model and caregiver can look at. Additionally, the caregiver will be present only for a number of time steps corresponding to a geometric probability function with average \bar{T}_{present} , after which she will not be present in the room (i.e., “leaves the room”). She will return to continue interacting with the infant model after a number of time steps drawn from another geometric probability function with average \bar{T}_{absent} . This is meant to simulate the fact that infants carry out a certain amount of visual exploration of their environment without other people present.

During any time step, the infant model and the caregiver can change their gaze direction. φ_I captures the direction towards which the infant model is looking (where $\varphi_I = 0^\circ$ corresponds to the infant model looking towards the caregiver’s location), while φ_H and φ_E capture the caregiver’s head and eye direction respectively (where $\varphi_H = 180^\circ$ corresponds to the caregiver’s head turned towards the infant model, and similarly for φ_E). Variables φ_{o_i} hold the angle of object i from the infant model’s point of view, and do not change when the infant model’s heading changes.

Similar to Φ_{o_i} , Φ_C is the caregiver’s visual saliency as perceived by the infant model. This value is temporarily decreased to half when the caregiver is not looking at the infant model, modeling infants’ preference for looking at gaze directed at them rather than diverted elsewhere [23], [62]. Φ_I is the infant model’s visual saliency as perceived by the caregiver (Φ_I along with Φ_{o_i} will drive the caregiver’s visual attention, as described below). (Default values and ranges for all learning environment parameters as well as other parameters of the model to be described next are listed in the top section of Table I).

B. Infant Visual System

The infant model’s visual input is processed by three different systems (Fig. 2, left): a saliency map (\mathbf{s}); a head direction detector (\mathbf{h}); and an eye direction detector (\mathbf{e}).

The **saliency map** ($\mathbf{s} = [s_1, \dots, s_{96}]$) indicates the presence of visual saliency in a body-centered coordinate system with 96 different regions in space, along 24 heading ranges and 4 depth ranges. Our assumption of a body-centered representation is a convenient simplification that frees us from having to model how the infant brain computes various coordinate transformations.

Values for elements of \mathbf{s} are derived from saliency values of the caregiver and objects corresponding to the location (heading and depth) of these elements. Both the objects and the caregiver have to be within the infant model’s field of view (FOV) ($\varphi_I - \text{FOV}/2 \leq \varphi_{o_i} \leq \varphi_I + \text{FOV}/2$, where FOV is the extent of the visible area with respect to the infant model’s gaze direction) to be added to their corresponding s_i . Previously, these values are foveated and habituated, as discussed next.

TABLE I
OVERVIEW OF MODEL PARAMETERS, THEIR ALLOWED RANGES, AND DEFAULT VALUES

Symbol	Explanation	Range	Default
<i>Learning environment</i>			
Φ_I	infant model's saliency	$(-\infty, \infty)$	4.0
Φ_C	Caregiver's saliency	$(-\infty, \infty)$	4.0
$\bar{\Phi}_O$	Average object saliency	$(-\infty, \infty)$	1.0
\bar{N}_o	Average number of objects	$[0, \infty)$	4
σ_o	Object placement spread around infant model	$[0, \infty)$	1.0 m
$T_{present}$	Average caregiver interaction interval	$[0, \infty)$	60 s
T_{absent}	Average caregiver absence interval	$[0, \infty)$	60 s
$T_{objects}$	Average object replacement interval	$[0, \infty)$	5 s
<i>infant model visual system</i>			
FOV	Size of field of view	$[0^\circ, 360^\circ]$	180°
$\sigma_{H_{initial}}$	Initial σ_H value	$(0^\circ, \infty)$	50°
$\sigma_{H_{final}}$	Final σ_H value	$(0^\circ, \infty)$	1°
$\sigma_{H_{step}}$	Decrement in σ_H per 300,000 time steps	$(0^\circ, \infty)$	5°
$\sigma_{E_{initial}}$	Initial σ_E value	$(0^\circ, \infty)$	50°
$\sigma_{E_{final}}$	Final σ_E value	$(0^\circ, \infty)$	1°
$\sigma_{E_{step}}$	Decrement in σ_E per 300,000 time steps	$(0^\circ, \infty)$	2°
τ_H	Habituation rate	$[0, \infty)$	2.5
α_H	Target of habituation	$[1.0, \infty)$	1.0
d	Memory decay factor	$[0, 1]$	0.5
<i>Reinforcement Learning</i>			
η	Learning rate	$[0, \infty)$	0.005
γ	Discount factor	$[0, 1]$	0.1
β	Inverse temperature	$[0, \infty)$	30

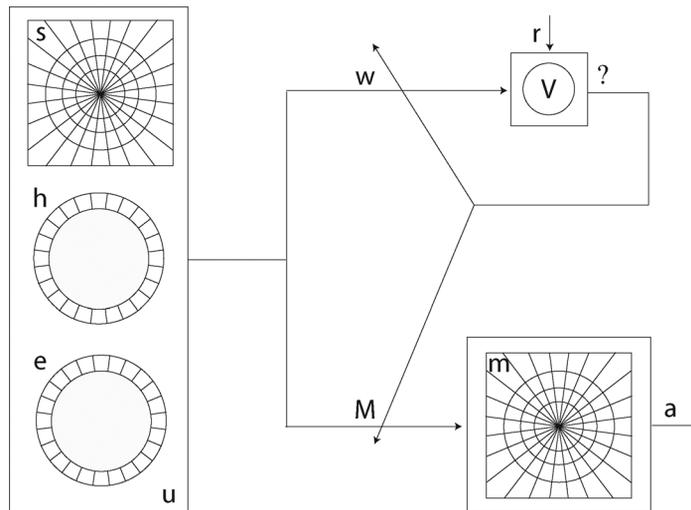


Fig. 2. (Left) Details of the infant model's visual system. (Right) Details of the actor-critic reinforcement learning model. Features calculated from the saliency map (s), caregiver head direction (h), and caregiver eye direction (e) are combined into \mathbf{u} [visual input vector, $\mathbf{u}(t) = (s(t), \mathbf{h}(t), \mathbf{e}(t))^T$], multiplied by \mathbf{w} [weight vector, $\mathbf{w}(t) = (w_1(t), w_2(t), \dots, w_{N_s}(t))$], and added into V [so that $V(t) = \mathbf{w}(t)\mathbf{u}(t)$], resulting in an estimate of the value of the present state. \mathbf{u} is also multiplied by \mathbf{M} (a matrix with as many rows as there are actions and as many columns as there are input features) and added into \mathbf{m} [so that $\mathbf{m} = \mathbf{M}(t)\mathbf{u}(t)$], where values in \mathbf{m} correspond to locations in space, as shown in the figure. The action a will consist of looking at a location in space, generally the one corresponding to the highest value in \mathbf{m} .

Foveation causes perceived saliency to decay as it falls outside the infant model's center of vision, according to the following formula (based on the contrast sensitivity function proposed by Daly *et al.* [18]):

$$S(\theta) = 0.2 + 0.8 \frac{1}{1 + k_{Ecc}\theta}$$

where θ is the eccentricity in visual angle of the object or caregiver, and k_{Ecc} is a constant that defines how the sensitivity diminishes with eccentricity. k_{Ecc} is set to 0.24 based on a fit on data sets from Virsu and Rovamo [68] and Johnston [42]. The offset of 0.2, added based on gaze-following experimental results where a distracter object at the periphery of vision captures the attention of the infant model, prevents values from decaying

to close to zero when objects are in peripheral vision (i.e., “in the corner of the eye”).

The infant model habituates separately to each object, with a formula proposed by Stanley [71] where perceived input decays in an exponential fashion with time

$$\tau_H \frac{d\phi_{o_j}(t)}{dt} = \alpha_H (\Phi_{o_j} - \phi_{o_j}(t)) - S_{o_j}(t)$$

where $\phi_{o_j}(t)$ is object j 's habituated saliency at time t and Φ_{o_j} its original dishabituated saliency; $S_{o_j}(t)$ is equal to Φ_{o_j} if the infant model is looking at object j at time t and 0 otherwise; τ_H is a time constant that specifies the rate of habituation (a smaller τ_H resulting in faster habituation); and α_H controls the level of long-term habituation. A similar formula applies for ϕ_C and

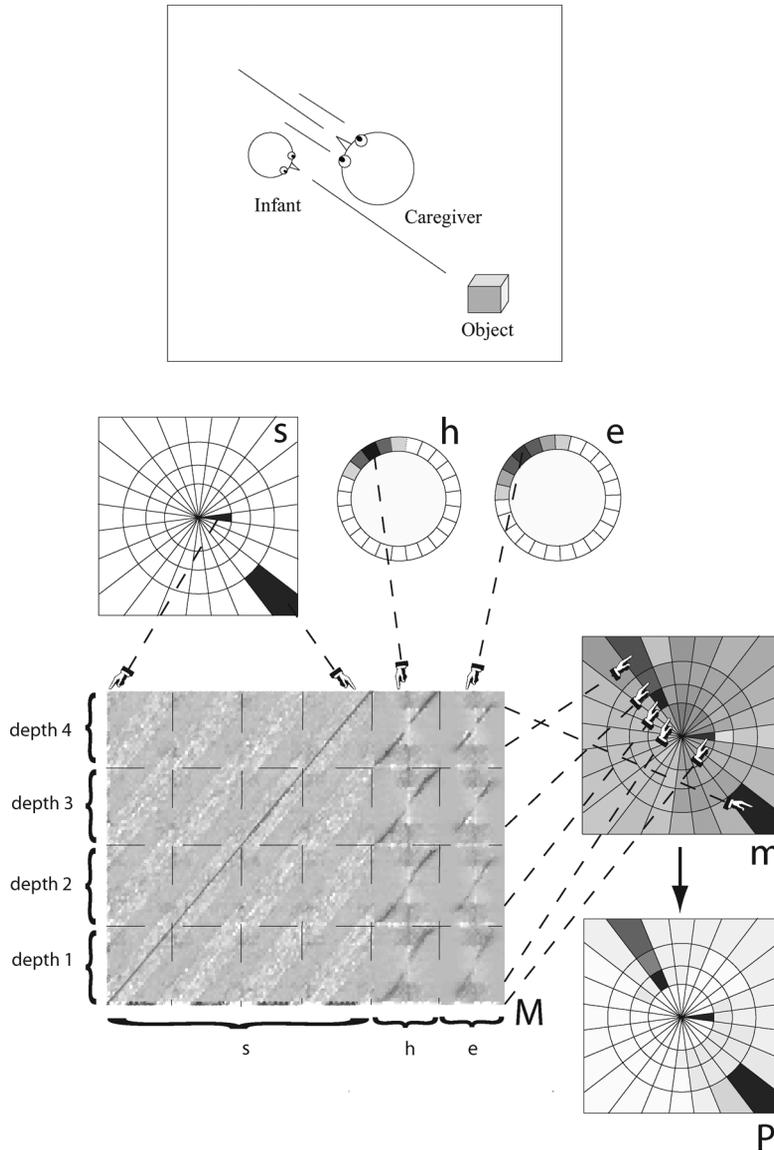


Fig. 3. Selecting an action. The room configuration as well as memory traces in the infant model’s visual system result in values for s (saliency map), h (head direction detector), and e (eye direction detector). s , h , and e contain $64 + 16 + 16 = 96$ values, that are multiplied with each row in M , resulting in values for m , where each value in m correspond to one of the 64 locations in space. The actual probability of looking at a particular location is stored in vector P . Probabilities thus correspond to a softmax selection of the highest value in m . The maximum value has the highest probability of being chosen, but with some probability of choosing other actions with high value in m .

ϕ_I , the habituated saliencies of the caregiver Φ_C and the infant model Φ_I , respectively.

The complete formula for calculating each element s_i of s within the infant model’s FOV is thus the sum of the habituated foveated saliencies of all objects that fall within s_i ’s range (heading and depth), plus the habituated foveated saliency of the caregiver if it falls within its range: $s_i = S_O + S_C$, where $S_O = \sum_{j=1}^N S_{o_j}$ and $S_{o_j} = \phi_{o_j} S(\theta_{o_j})$ if o_j lies within the range of s_i and zero otherwise, $\theta_{o_j} = |\varphi_I - \varphi_{o_j}|$ being the angular distance of the object from the center of vision, and $S_C = \phi_C S(\theta_C)$ if the caregiver is present and falls within the range of s_i and zero otherwise. θ_C will always be zero because the caregiver is positioned in front of the infant model at 0° .

For elements s_i of s outside the infant model’s FOV, the perceived saliency is zero, but the infant model’s visual system calculates the new value of s_i as a fraction of whatever previous

value it had. This is done by multiplying the previous value of s_i by a constant d ($0 < d < 1$; see Table I) and having the resulting number be the new value of s_i . This “memory decay” factor enables the model to temporarily remember recently observed states of the world. The top section of Fig. 3 shows an example setting and the resulting value of s .

The **head direction detector** ($h = [h_1, \dots, h_{24}]$) indicates 24 possible caregiver head directions as perceived by the infant model. If the caregiver is present and the infant model is looking at her, the value of each h_i is calculated according to an exponential decay, so that the closer h_i is to the caregiver’s heading (φ_H), the higher its value. If the infant model is not looking at the caregiver, then the values of h decay as in the calculation of s earlier.

The exact formula for calculating h when the caregiver is present and the infant model is looking towards her ($\varphi_I = 0^\circ$)

is: $h_i = \exp(-(\varphi_H - \theta_{I_i})^2 / \sigma_H^2)$, where σ_H is the decay factor, described below, and θ_{I_i} is the angle corresponding to heading i 's center ($\theta_{I_1} = 0^\circ, \theta_{I_2} = 15^\circ, \theta_{I_3} = 30^\circ, \dots, \theta_{I_{24}} = 345^\circ$). All h_i are normalized (using linear scaling), so that the sum of all h_i add to 1 ($\sum_{i=1}^{24} h_i = 1$). When the caregiver is not present or the infant model is not looking at her, the infant model's visual system calculates the new value of h_i as a fraction of whatever previous value it had, by multiplying h_i by the same d as above, resulting in a similar "memory decay" as above. The top section of Fig. 3 shows an example setting and the resulting value of \mathbf{h} .

Adults can detect differences in the gaze direction of others [1], [16], [28], with an acuity of just 1.4° at a distance of just over 1 m [16]. However, the development of this sensitivity during infancy has not been systematically studied. Very young infants can differentiate between direct and averted gaze [23]. In gaze following experiments [11], [13] infants can distinguish between right and left-facing head directions at around 6 months of age, and by 12 months they can discriminate a 25° difference in gaze direction to two objects [34]. This progression is captured in the model by making σ_H decrease linearly with time (see Table I). We do not explore the nature of the mechanism behind this increase in sensitivity, leaving this for further research. However, we do not discard a developmental learning-based process akin to the gaze following model presented here.

The **eye direction detector** ($\mathbf{e} = [e_1, \dots, e_{24}]$) is similar to \mathbf{h} , but computed with the caregiver's eye direction (φ_E) instead of head direction (φ_H), and with σ_E instead of σ_H (as above, this reflects the infant model's limited memory for this information). The values of σ_E and σ_H (see Table I) were chosen to reflect that infants improve their estimates of head direction faster than their estimates for eye direction, because the latter requires greater acuity. Additionally, when the infant model is looking at the caregiver but she is turning her head back ($\varphi_H < 90^\circ$ or $\varphi_H > 270^\circ$), all values e_i are set to zero. This reflects the fact that when the caregiver is facing backwards with respect to the infant model, the eyes are not visible. The top section of Fig. 3 shows an example setting and the resulting value of \mathbf{e} .

The output of the infant model's visual system at time t combines the saliency map and the head and eye direction representations into one vector $\mathbf{u}(t) = (\mathbf{s}(t), \mathbf{h}(t), \mathbf{e}(t))^T$. This vector has $N_s = \dim(\mathbf{s}) + \dim(\mathbf{h}) + \dim(\mathbf{e}) = 144$ dimensions.

C. Reinforcement Learning

At each time step, the output $\mathbf{u}(t)$ of the visual system described above is used as input to a standard actor-critic reinforcement learning system [19], [61] that decides where the infant model will look next.

The **critic** part of the actor-critic reinforcement learning algorithm (see Fig. 2, upper right) estimates the value of being in the present state of the world. We use a simple linear model $V(t) = \mathbf{w}(t)\mathbf{u}(t)$, where $\mathbf{w}(t) = (w_1(t), w_2(t), \dots, w_{N_s}(t))$ is a weight vector. It is updated each time step according to the standard gradient descent rule [19]

$$\mathbf{w}(t+1) = \mathbf{w}(t) + \eta\delta(t)\mathbf{u}(t).$$

In this, $\delta(t)$, also called temporal difference error, is defined as

$$\delta(t) = r(t) + \gamma V(t+1) - V(t)$$

where $r(t)$ is the reward at time t which corresponds to the saliency of the new location being looked at, $V(t+1)$ the estimated value of the new state reached after taking the action, and γ is a parameter, also called a "discount factor," that specifies how much into the future do we want the agent to look into when measuring the effect of the action being taken. A value of 0 means that only the reward obtained immediately after the action will be taken into consideration. A value of 1 is only valid in tasks where there is a start state and an end state. Since our infant model does not have an end state, and the environment is fairly nonstatic, we choose a value of 0.1.

The **actor** part of the reinforcement learning algorithm (see Fig. 2, lower right) specifies the motor action to be taken that will direct the infant model's attention to one of 24 possible different headings (i.e., directions) H and one of four different depths (i.e., distances) D , with a total of $N_a = 96$ different possible actions ($A = (H, D)$, $H \in \{0^\circ, 15^\circ, 30^\circ, \dots, 345^\circ\}$, $D \in \{0.4 \text{ m}, 1.0 \text{ m}, 1.45 \text{ m}, 2.0 \text{ m}\}$), where A is the action, and H and D are the heading and depth, respectively, where attention is directed. At each time step, the action is chosen probabilistically according to the softmax decision rule [19]

$$P[a] = \frac{\exp(\beta m_a)}{\sum_{a'=1}^{N_a} \exp(\beta m_{a'})}$$

with m_a being the action value parameter for action a for the present state and calculated every time step from $\mathbf{m} = \mathbf{M}(t)\mathbf{u}(t)$, where \mathbf{M} is a matrix with as many rows as there are actions and as many columns as there are input features (see Fig. 3 for an illustration of \mathbf{M}), and \mathbf{u} is the vector with the input features from the visual system as described above. In the formula, the higher the value of m_a , the higher the chances of selecting action a . The softmax selection formula works as follows. For very high values of β , one always chooses the action with the maximum m_a value; for low β , one chooses the uniformly at random among all actions; and for intermediate β , the selection is something in between, thus, it is a "soft" form of selecting the maximum.

After an action is taken, \mathbf{M} is updated according to a formula that is similar to the one used to update the critic's weights

$$M_{a'b}(t+1) \leftarrow M_{a'b}(t) + \eta(\delta_{aa'} - P[a'])\delta(t)u_b(t)$$

for all elements of \mathbf{M} , where a' and b correspond to the particular element of \mathbf{M} being updated (a' corresponding to an action and b to a feature; see Fig. 3 for an illustration), a is the action that was taken, η is the same learning rate as above, $\delta(t)$ is the value of the temporal difference error as calculated by the critic above, $P[a']$ is the probability of taking action a' at this time step as calculated in the softmax formula, $\delta_{aa'}$ being defined as 1 if $a = a'$, 0 otherwise (so that it will only be 1 if a' matches the actual action a taken), and $u_b(t)$ is the value of $\mathbf{u}(t)$ corresponding to b .

D. Training and Testing

The simulation is run for 15×10^6 training steps, where each time step corresponds to one second, as noted above. During this process, the infant model selects an action at the end of each time step, based on the state produced by the simulation environment and the caregiver, and as mediated by the reinforcement learning algorithm described above. The caregiver in turn always looks at the most salient point in the room (involving habituation-mediated values of Φ_{o_i} and Φ_I). The caregiver's perceived saliencies are mediated by the same foveation and habituation mechanisms (with identical parameters) as in the infant model's visual system. The caregiver's head direction is slightly offset from that of the eyes according to a Gaussian distribution with $\sigma = 5^\circ$ and $\mu = 0^\circ$. This offset is recalculated for every gaze shift of the caregiver. This reflects the fact that eyes and head are not always perfectly aligned, and corresponds to values observed in naturalistic settings [32]. During this training phase, the weights of the critic and actor are updated every time step, as described above.

Training is stopped after a number of time steps, to observe the infant model's reaction to different configurations of the room. (These "experiments" are described in the following sections in detail). One of the advantages of the model is that it can be tested extensively, and we typically did 200 repetitions for each different test setup so as to properly establish results. During testing, learning (i.e., updating of weight values) is "frozen", otherwise the extensive number of tests we perform would bias the infant model towards the test setups used (this is not an issue for real experiments, where infants are tested with a low number of repetitions). After all testing is carried out, training is resumed by returning the environment to its previous configuration and continuing with the training schedule described above.

III. RESULTS

A. Basic Emergence of Gaze Following

The first thing the model learns is to look preferentially at locations with high saliency. Although simple, this relationship between the saliency of specific locations and the reward resulting from looking at those locations is not a prior state of the model, but instead must be learned. The large diagonal, at the left part of \mathbf{M} in Fig. 3, shows the result of this learning. A nonzero saliency at location i (s_i) will mostly add, through \mathbf{M} , to the element of \mathbf{m} corresponding to action i (m_i), eventually increasing the probability of looking at precisely that location i . Before learning, however, all values are zero. This reflects the fact that human infants also take time to develop saliency preferences. For example, they consistently saccade to stimulus contours at 14 weeks, but not at 2 weeks of age [6]. Although there are a few findings of weak visual stimulus preferences in infants within hours or several days of birth [41], [59], there is no evidence that these preferences are spatially specific; rather, stimuli that fall into the visual field might elicit different looking times. Such studies therefore demonstrate slightly differentiated values for different stimuli within several hours or days of postuterine visual experience, but not differentiated spatial bindings. Moreover, even a low-power machine vision system can

learn to discriminate classes of salient stimuli within a few minutes of visual experience [9], so even the initial stimulus values can be set very rapidly. For these reasons, the choice to start the initial spatial-value matrix at zero is not inconsistent with available developmental evidence.

The smaller diagonals, on the right part of \mathbf{M} in Fig. 3, correspond to the learning of gaze following. Shown are the final states of the diagonals, when gaze following has been fully learned. Although these diagonals start to form from the very first training steps, they take longer to develop than the saliency-related larger diagonal described above. The reason for this is that the caregiver's head and eye direction predict reward for looking at a given location, but with much less certainty than direct saliency itself. Instead, there are several possible locations with visual reward for a particular head or eye direction, which are collectively located in the path consistent with the caregiver's head and eye direction.

Learning to look at saliencies is therefore easier to learn than gaze following. However, it is a prerequisite for it. If the infant model did not know about the value of looking at salient locations, then gaze following could only be learned in episodes where the infant model incidentally looks at the caregiver, and immediately afterwards incidentally looks at the object the caregiver is looking at. But once looking at salient objects has been learned, then the caregiver will naturally attract the infant model's visual attention, and after habituation has taken place, the infant model will likely look at salient objects in the room, making it more likely that the next target will be the same one that the caregiver is looking at, thus facilitating the acquisition of an abstract gaze-following policy. Note too, that this does not apply to situations where the caregiver is looking at an object behind the infant model, which would be outside the infant model's FOV since it is looking at the caregiver. Instead, after habituating, the infant model would likely either look at any object within view, or explore to a random location in the room. This makes gaze following to objects behind the infant model take longer to develop, as described in the next sections.

Fig. 4 illustrates two examples of \mathbf{m} at the end of training, for different states of the world. Here, the infant model, caregiver, and objects are superimposed on a representation of the world, divided into all possible looking locations. In the left-side illustration, two objects within the infant model's FOV create dark tones (high values) in their corresponding locations, with the bigger one corresponding to a darker tone because of its higher saliency. Since the infant model is looking at the caregiver, locations in the caregiver's line of sight are of a darker tone, including locations outside the infant model's FOV, even when there are no objects there. Other locations not corresponding to these two objects or the caregiver's line of sight are lighter-colored, and not entirely uniform because of small fluctuations in the stochastic learning of \mathbf{m} . The infant model will thus likely choose to look at either of the objects or one of the locations within the caregiver's line of sight, with a small probability (how small depending on temperature) of looking elsewhere in the room. The right illustration shows two objects outside the infant model's FOV not activating \mathbf{m} , and two objects with similar saliency in the infant model's FOV activating two locations in \mathbf{m} . Since one of these locations is also within the caregiver's

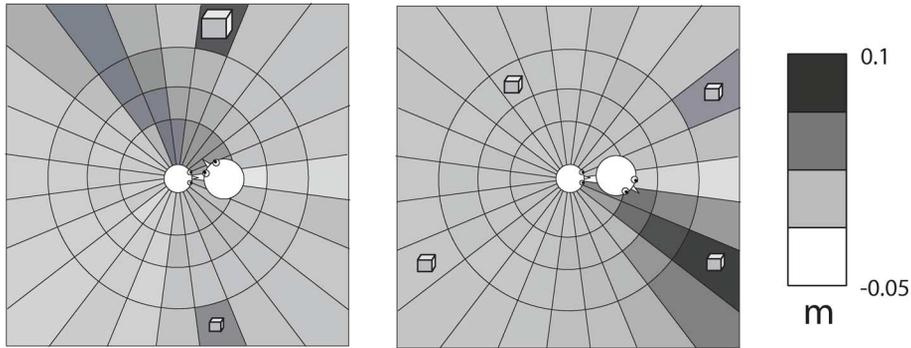


Fig. 4. Illustration of values of m for two different states of the world.

line of sight, it has a darker tone than the other, and will likely be selected over the other one, resulting in gaze following. The scale for m , from -0.05 to 0.1 , is not absolute. In other phases of learning, m 's values could have other ranges. The range used, thus, is only for illustration purposes.

B. Mechanisms Improving Spatial Aspects of Gaze Following

Butterworth's theory of gaze following development [10] posited progressive and discrete accumulation of processing mechanisms, including a salience-based ecological mechanism, a spatially allocentric geometric mechanism, and a mentalistic representational mechanism. In this section, we describe in detail the two gaze following experiments that Butterworth used to postulate the existence of these mechanisms. By replicating these experiments in our simulation, we show that our infant model elicits the succession of infant behaviors that suggested these mechanisms, but without ever requiring any qualitatively new mechanisms.

1) *Emergence of the Geometric Compensation Stage*: We begin with the second mechanism proposed by Butterworth, because the first is a reactive exogenously generated response to salient stimuli, corresponding roughly to the starting state of our model. Butterworth carried out a series of experiments, the results of which led him to propose a putative "geometric compensation mechanism," which would allow the infant to disregard distracter objects between the caregiver and the object to which she attends [11], [13]. In these experiments, Butterworth and colleagues placed salient targets at specific locations in an otherwise boring room. During each trial the caregiver looked at the infant and then turned toward one of the targets. Four targets were displayed during every trial, with two on each side of the room along the wall, as shown in the part of Fig. 5 labeled "4-Target." In the first three variations of the experiment (top row) the correct target is *first along the infant's scan path* when compared to the other target on the same side of the room. The correct target could be at either 30° , 60° , or 90° (absolute values) from infant's midline. In the last three variations (bottom row) the target is *second along the scan path*, at either 90° , 120° , or 150° from midline. We will refer to these settings, as "30", "60", "90-first", "90-second", "120", and "150". (Note that only the 90° target occurs in both first and second position; this becomes an important distinction. All other locations are only first or only second on the infant's scan path when rotating from midline.)

In these studies, each trial type, as defined by a particular target location, was repeated twice, once for each side of the room (four times for the 90° locations: two trials for 90-first; two for 90-second). They defined a *correct* response as looking at plus or minus 30° around the correct target, and a *wrong* response as looking within 30° of the incorrect target on the same side of the room. *Non-codable* responses were coded if the infant made no response within six seconds, or looked at up or down. A final category of responses, omitted from the calculation of scores, included looking at the opposite side of the room or (though not described by Butterworth) looking at no-target locations (i.e., the wall) on the correct side of the room.

Fig. 6 shows a plot of the accuracy scores we calculated based on Butterworth's reported results (the origin was set to 0 months, for comparison with the figure on the right, described below). This score is calculated as the number of correct responses over the sum of correct plus wrong responses. Results show that at all ages, infants reliably follow gaze when the correct target is positioned first along the scan path (30° and 60° trials). When the correct target is first along the scan path at 90° (but still within the FOV when looking at the caregiver), 6-month-olds stop at the distracter object about half the time, 12-month-olds disregard the distracter more often, and 18-month-olds follow gaze correctly, disregarding the distracter altogether. When the correct target is second along the scan path at 120° or 150° (now outside the visual field), at no age did infants reliably follow gaze. Butterworth interpreted these results as indicating that infants by 6 months follow gaze only via an "ecological mechanism," which also causes them to be distracted by other objects along the scan path. With time, they are able to disregard the distracter object through a "geometric compensation mechanism," but only in the "90 second along the scan path" variation. In the other second-target trials, however, where the target is outside the infant's visual field, infants have even more difficulty disregarding the distracter within their visual field.

2) *Modeling the "Geometric Stage"*: Our simulation was intended to replicate the spatial layout and trials of Butterworth's experiments. The saliency of all objects Φ_o was set to 1, corresponding to objects of an average saliency. At the beginning of every trial, the infant model looked at the caregiver ($\phi_I = 0^\circ$), and the caregiver's heading (ϕ_h and ϕ_e in tandem) was set to look towards the object specified by the particular trial. Each trial was run for 6 time steps, equivalent to the six seconds used in the experimental paradigm described earlier. During this time

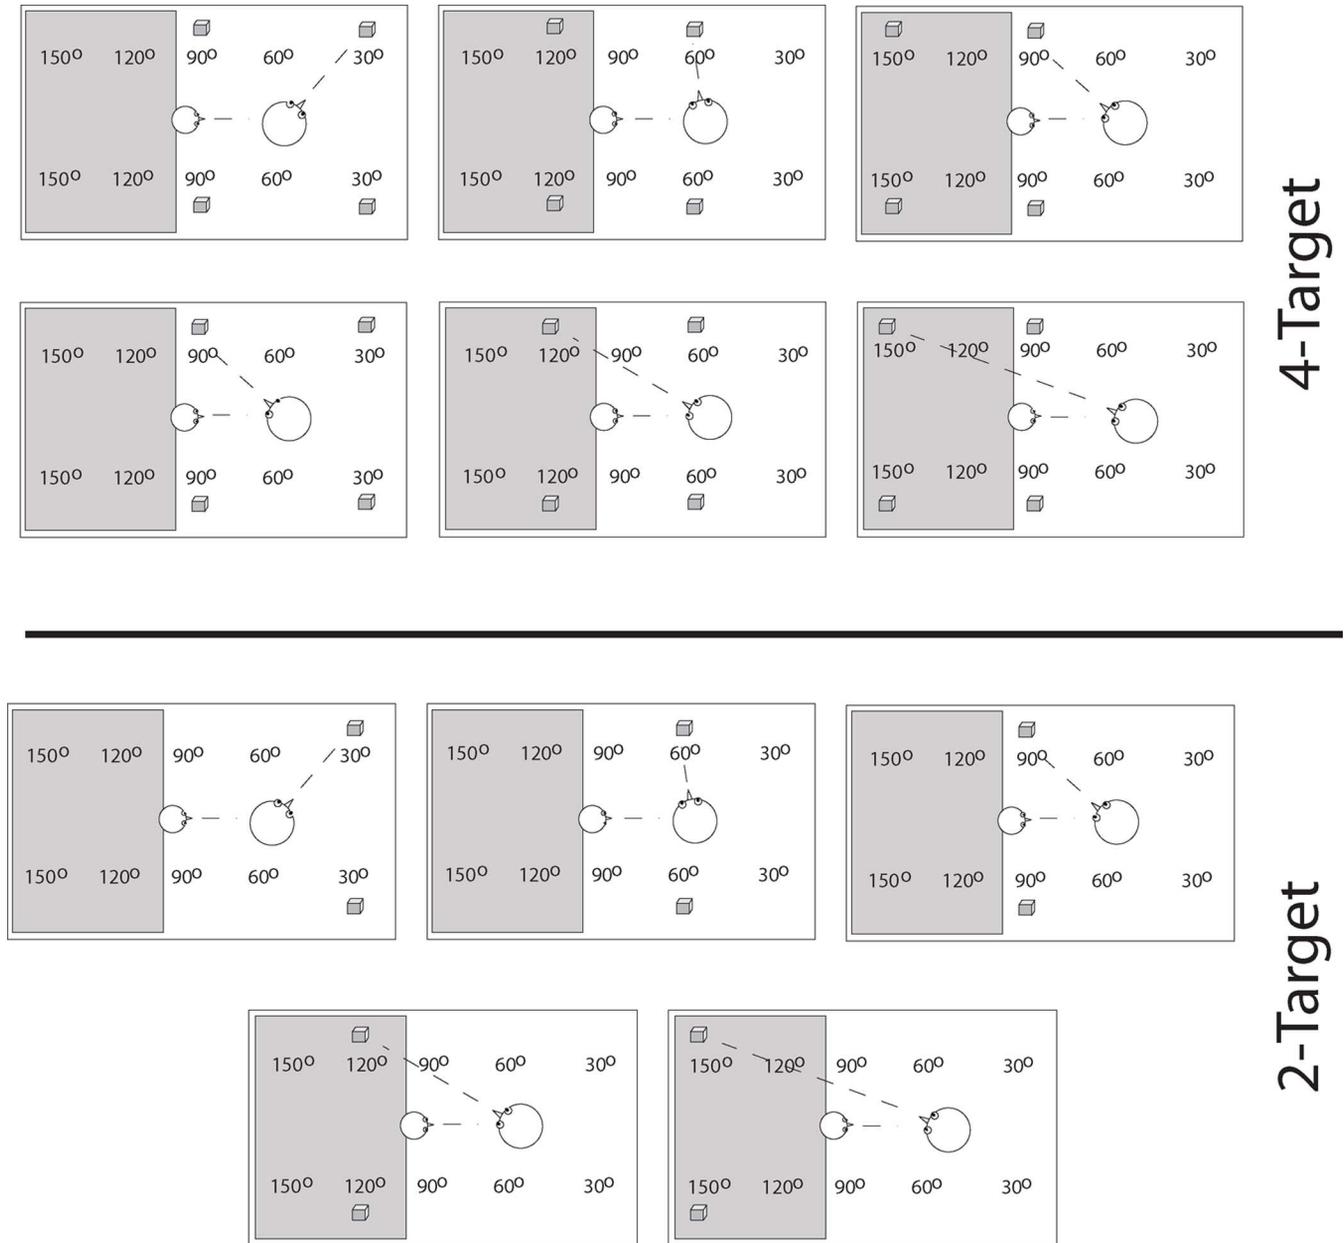


Fig. 5. Butterworth's 4-target and 2-target settings. Gray area represents space outside the infant's FOV. (Top row) 4-target setting: Target *first* along the scan path, at 30° (left), 60° (middle), or 90° (right). (Bottom row) 4-target setting: Target *second* along the scan path, at either 90° (left), 120° (middle), or 150° (right). (Top row) 2-target setting: Target *within* the infant's FOV, at 30° (left), 60° (middle), or 90° (right) from the infant's midline. (Bottom row) 2-target setting: Target *outside* the infant's FOV, at 120° (left), or 150° (right).

the caregiver's heading did not change, and the infant model followed the actions specified by its action-selection algorithm, as described earlier.

Trials were scored as follows. *Correct* responses were those where the infant model's attention (heading) shifted from initially looking at the caregiver directly to looking at the target, or to a heading immediately to the right or the left of the target. This corresponds to looking at the target plus or minus 22.5°, which is slightly stricter than the $\pm 30^\circ$ used by Butterworth. *Wrong* responses were those where the infant model looked at the correct side of the room but at the incorrect target. If during the trial's 6 time steps the infant model did not shift gaze, the response was considered *noncodable*. Re-

sponses where the infant model either looked at the wrong side of the room or at empty space on the correct side of the room were omitted, as in Butterworth's studies. These responses were quite infrequent in the infant model, as in the original studies of real infants.

After every 1 500 000 time steps, learning was suspended (to avoid over-training) and each of the six 4-target trial types was performed 200 times (100 times on each side of the room). Fig. 6 shows the results for the complete training, i.e., for the whole 1500 000 time steps (for more detailed results, see [36]). These results correspond to five models that were trained and tested separately, each with its own random action-selection seed, and with an initially blank state of the world (i.e., all values

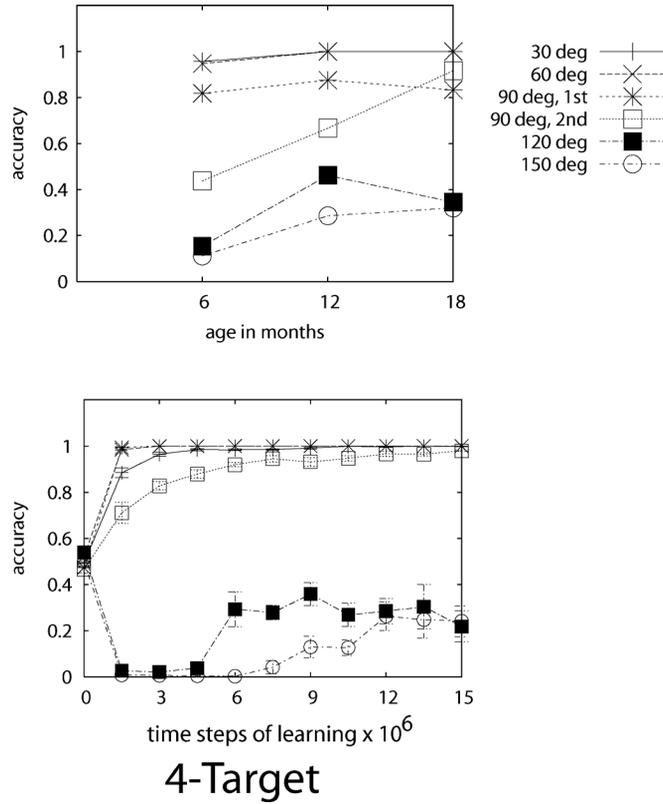


Fig. 6. Results for the 4-target setting. (Top) Butterworth’s results. (Bottom) Model’s results (error bars indicate standard errors after 5 repetitions).

in array \mathbf{M} were zero). At 0 time steps, no learning has taken place, so that the results reflect random actions.

We can see that the accuracy of following to the 30°, 60°, and 90°-*first* targets is close to 100%, similar to what Butterworth reports for those setups (compare with Fig. 6, top). The infant model progressively increases the accuracy for the 90°-*second* targets, qualitatively similar to the increase of accuracy reported by Butterworth. However, note that this target is learned somewhat sooner, relative to the other locations, than it was by the infants in Butterworth’s studies. Finally, for the 120° and 150° settings, the infant model never reaches the point of reliably following gaze, similar to the findings reported by Butterworth.

The infant model’s progress in the “90°-second” setting, which Butterworth attributes to the incorporation of the geometric compensation mechanism, is explained as follows. Initially, foveation makes the 30° target more salient than the 90° target. But because the caregiver is looking at the 90° target, the emerging association of the head/eye state eventually balances the probability of selection between the two targets. What then causes the preference for the 90° target is the decay in σ_H and σ_E , the parameters that define how “fuzzy” the representation of the caregiver’s head/eyes direction is in \mathbf{h}/\mathbf{e} . As σ_H and σ_E diminish, the values of \mathbf{h} and \mathbf{e} reflect the caregiver’s head and eye direction more precisely. The increasing probability of gaining reward by highly valuing the head and eye cues eventually “wins” over the saliency of the 30° target; thus, the infant model shifts to prefer looking at the 90° target. Note that the accuracy at 12 months in [11] is around

30%, contrasting with the results in Butterworth and Jarrett [13] of around 70% at 12 months, and 50% at 6 months for a similar setup. Our model could explain this disparity in results as due to differences in target saliency. In fact, Butterworth and Cochran’s targets were blue stars on a yellow background, whereas Butterworth and Jarrett’s were simply yellow squares. It is reasonable to presume that the former were more salient, consistent with our explanation. (However, without an experimental comparison we cannot assume that the former were *functionally* more salient to infants than the former.)

The model also can explain why in the 120° and 150° trials gaze is followed seldom and late. The saliency of the distracter can override the cue value of the caregiver’s head and eyes direction, because the foveation offset does not let saliency drop to values too close to zero when the distracter is positioned at 30° and 60°, respectively. This is illustrated by repeating the 4-target setting with the target at 120° (Fig. 5, bottom row, middle) with different values for the object saliencies. Although both target and distracters vary in saliency, the target (and the distracter on the other side of the room) is not visible to the infant model. The saliency of the visible distracter changes the results of the experiments, as shown in Fig. 7. Decreasing object saliencies helps the infant model disregard the distracter (object saliencies of 0 and 0.5), whereas increasing the object saliencies causes the infant model to always look at the distracter instead (object saliencies of 1.5 and 3.0). This was observed by [31], who noted an almost 100% likelihood of attending to the first target along the scan path (the distracter on the same side of the room as the target) when the saliency of targets was increased by setting them in motion.

Another possible contributor to this effect is the informativeness of the caregiver’s head and eye cues. When the infant model turns too far away from the caregiver (i.e., to the 120° or 150° targets), the caregiver’s head and eyes are out of view, and their information value, relative to the infant model’s learning state, is minimized. Experiment 2 in [20] found that infants’ tendency to ignore targets behind them is partly due to the reduced saliency of the caregiver’s head and eye cues. In addition, [13] found that 6-month-olds’ gaze-following drops off when they must turn more than 90° to 135° away from their mother’s face. These findings also are consistent with the model.

3) *Emergence of the Representational Stage*: In Butterworth’s 2-target setting experimental paradigm [11], [13], targets were positioned two at a time, one on each side of the room along the wall, as shown in the part of Fig. 5 labeled “2-Target.” In the first three variations of the experiment (top row), targets were set at either 30°, 60°, or 90° from the infant’s midline. These locations are *within the infant’s FOV*. In two other variations (bottom row), the targets were set at either 120° or 150°. These locations are *outside the infant’s FOV*. Butterworth posited that the three first setups (30°, 60°, and 90°) do not require the infant any kind of representation of space not immediately seen, and that following gaze at the last two setups (120° and 150°) would only make sense when the infant realizes that there is space besides the immediately seen, and that an object might be there.

Scoring for this setting is similar to the one used for the 4-target experiments, except that instead of wrong responses

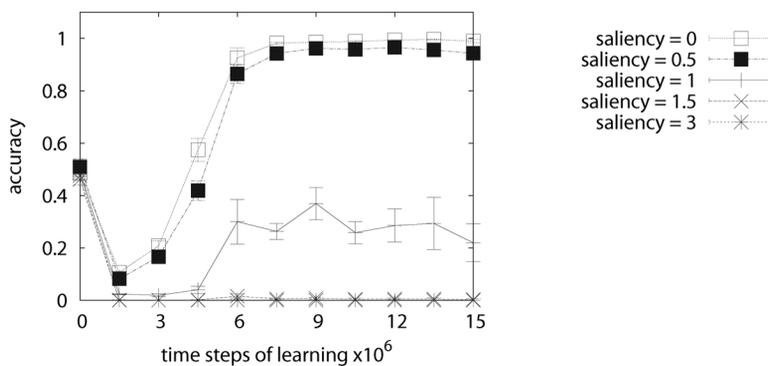


Fig. 7. Effect of varying object saliencies in the 4-target, 120° setting. Increasing the saliency causes the infant model to look at the distracter, resulting in a low score (accuracy) for this setting. But with a low target saliency, the infant model is able to disregard the distracter and achieve a high accuracy. Note that the result for object saliency = 1 is the same as that of Fig. 6, 120°.

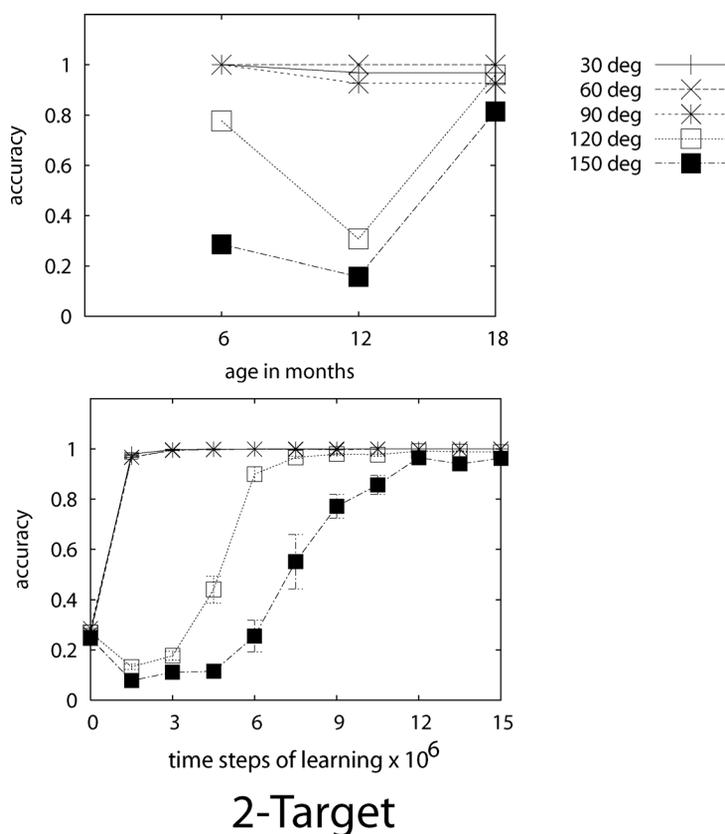


Fig. 8. Results for the 2-target setting. (Top) Butterworth's results. (Bottom) Model's results (error bars indicate standard errors after 5 repetitions).

being defined as looking at plus or minus 30° around the incorrect target on the same side of the room, they were defined as looking at another location on the same side of the room as the target.

The results are presented in Fig. 8. Infants age 6 months and older turned when the target was within their FOV (angles of 30°, 60°, and 90°). However, not until 18 months of age did they consistently follow gaze to targets outside their FOV (targets at 120° and 150°).

Butterworth interpreted these results as the infant incorporating a “representational mechanism” on top of the earlier ecological and geometric mechanisms. This new mechanism allows it to understand that objects can exist even when the infant

cannot see them. By eliminating the distracter object, Butterworth tested this understanding more cleanly. The results suggest that either a representational mechanism, or some other effect, renders infants unable to follow gaze to unseen objects until some months after they have started to follow gaze to peripheral locations.

4) *Matching the Representational Stage:* We conducted new simulations to determine if the infant model could, through the same single learning mechanism, replicate the results from Butterworth's 2-target experiments. Each of the five trial types were repeated 200 times (100 times on each side of the room). The scoring system was the same as was used earlier. In Fig. 8, we can see that the accuracy of the 30°, 60°, and 90° trials in the

2-target setting is always close to 100%, similar to what Butterworth reports for those settings (compare with Fig. 8, top). The model progressively increases the accuracy for the 120° and 150° trials, with 120° trials always slightly more accurate than 150° trials, but both approaching 100% by the end of training, as in the results reported by [11]. Thus, the infant model exhibited similar behavior as the human infants, which Butterworth explained as coming from a new mentalistic “representational” mechanism. However, no additional mechanism was added to the model.

The infant model’s progression in the 120° and 150° settings is explained as follows. Initially the infant model will learn to look at salient locations (see Section II). After this, its typical behavior will be to shift attention among the salient objects within view, including the caregiver. In this process, when the infant model looks at the caregiver and eventually habituates, it will look at another object within its FOV. These objects will lie in the white area of Fig. 5. After many repetitions of this behavior, the infant model will learn to follow gaze to objects in the white area. Gaze following to objects in the gray area will typically be learned when the infant model, after habituating to the caregiver, “explores”, shifting its gaze to a random location in the room, which in some cases turns out to be an object within the caregiver’s line of sight. This, however, will happen rarely when there are other objects within the infant model’s FOV that could capture its attention, and these “lost opportunities” for gaze following make this kind of gaze following slower to learn than gaze following to the white area.

5) *Discussion*: Butterworth designed these experiments to highlight spatial aspects of gaze following, positing geometric and mentalistic representational mechanisms emerging progressively to supplement the more basic, saliency-driven ecological mechanism. Our model explains the gradual improvements attributed to those later mechanisms through a unified biologically plausible learning process. Butterworth’s three-mechanism explanation is less parsimonious in that it requires additional explanations of how the mechanisms emerge and how they become integrated during development. Moreover, the model we propose assumes only a set of phenotypes that are very general, and, critically, easily demonstrable in fairly young infants. Those phenotypes are therefore known to be available to infants as they start learning the significance of social cues. They include, for example, foveation, saliency-based preferences, attention-shifting, reinforcement learning, and habituation. Our simulations show that these general and known traits, in the context of a structured environment, are sufficient to explain complex social behaviors, without positing additional specialized social-cognitive spatial or mentalistic mechanisms.

C. Mechanisms Supporting “True Gaze Following”

Infants gradually incorporate the caregiver’s eye direction cue in addition to the head direction cue during gaze following. This transition is said to be an indication of a qualitatively different kind of gaze following stemming from a “true understanding” of others as agents with intentions. In this section we simulate one of the most sophisticated experiments designed to detect this transition, and show that our system goes through the same transition without the need for additional mechanisms.

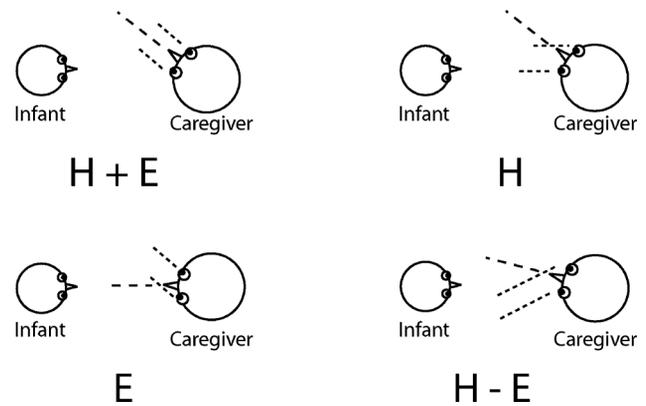


Fig. 9. Experimental conditions set up by Corkum and Moore [17]. (H + E) Both head and eyes turn. (H only) Head turns. (E only) Eyes turn. (H – E) Head turns in one direction, eyes in the opposite direction.

1) *Emergence of Use of Eye Direction Cues*: While the head direction of others is easier to discern over eye direction because of the relative size of the head over the eyes, the eyes give the correct indication of where others are seeing. Corkum and Moore [17] performed a set of experiments designed to determine at what age infants from 6 to 19 months give priority to eye direction cues over head direction cues, thus indicating the transition to “true gaze following”. In their experiments, infant and caregiver sat facing each other, with no targets present. Four different setups were used: 1) H + E; 2) H; 3) E; and 4) H – E (see Fig. 9). In the H + E condition, the caregiver turned both head and eyes 60° from midline. In the H condition, only the caregiver’s head turned, while her eyes remained directed towards the infant. In the E condition, only the caregiver’s eyes turned while the head remained pointed towards the infant. Finally, in the H – E condition, the caregiver’s head turned 30° to one side of the room, and her eyes turned 30° to the other side (this held constant the 60° displacement of head and eye direction). Gaze following scores were calculated by adding +1 for each turn to the *correct* side, –1 for each turn to the *incorrect* side, and 0 for continuing to look at the caregiver (i.e., *nonresponses*). The “correct side” was defined in the H – E condition as the side towards which the caregiver’s *head* turned, to assess preference for head direction over eye direction cues. However, E condition scores were based on eye-following, to assess sensitivity to eye direction alone. Corkum and Moore’s observations were as follows: 1) from 6 to 10 months of age, no gaze following was found; 2) at 12 to 13 months, some infants followed gaze, based primarily on head direction; 3) at 15 months infants followed gaze, based primarily on head direction but with some sensitivity to eye direction; 4) at 18 to 19 months, an effect of the eye cue was found, but not when head and eye orientations were in opposition. Corkum and Moore argue that the discrepancy between their results at 18 months and those of [47] and [13], who found gaze following based on eye direction alone, can be explained by procedural differences. The previous studies presented eye-only trials in separate blocks, which might have enhanced their saliency, relative to randomized blocks (where head cues are available in many trials), such as in Corkum and Moore’s design. Also, although Corkum and Moore did not find a difference between the H + E and the H variations before 18

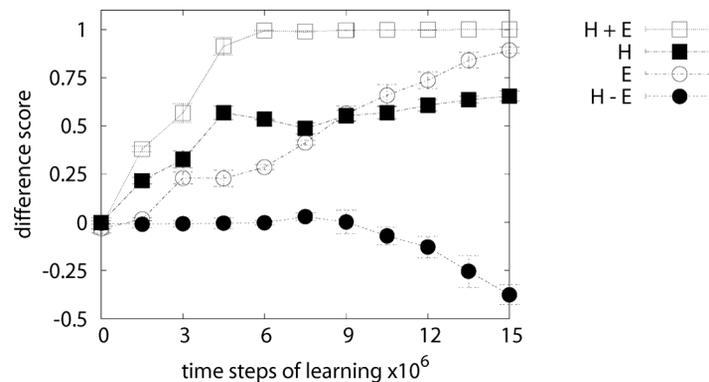


Fig. 10. Simulation results for head versus eye cue experiments.

months, [15] found a difference at 14 months, perhaps due to testing more infants (32 versus 12), or because they used a different procedure.

2) *Matching the Use of Eye Direction Cues*: Corkum and Moore’s simulations were replicated as follows. As in their experiments, no targets were used. At the beginning of every trial, the infant model looked at the caregiver and the caregiver’s head and eye orientations were set as specified by the particular condition. Each trial was run for 6 time steps, during which the adult’s head and eye direction did not change, and the infant model responded based on its action-selection algorithm. Trials were scored as in Corkum and Moore’s study.

These tests were run after every 1 500 000 time steps, immediately after the ones described above for the replication of Butterworth’s 4-target and 2-target setups. Each of the four trial types was repeated 200 times (100 times on each side of the room). Fig. 10 shows the results. At 4.5 million time steps, we can see an above-zero score in the H + E condition, this score being somewhat higher than in the H and clearly higher than in the E condition. This matches Corkum & Moore’s observation that at 15 months of age infants follow gaze based primarily on head direction and with some sensitivity to eye direction, with the distinction that while they did not find a difference between the H + E and the H condition, we did. This might be explained by our larger number of trials (200 versus 4), but it also seems like the “age range” in our simulation went much further than 18 months. Note that around 12×10^6 (12 million) steps, eye-direction begins to be used more reliably than head direction, which is of course what humans eventually learn.

There are some differences between the human and model performance. For example, the rise in effect of head-only cues in the model, around the time that H + E response reaches asymptote, was not seen in infants. Infants also never showed the distinction between E and H - E conditions that emerges in the model around 3-m steps. However, it should be noted that studies of infants have yielded somewhat different results, depending on methods, so there are clearly some salience and attention-recruiting factors that are not yet understood. Overall, however, our infant model shows a progressive influence, and eventual dominance, of eye direction cues, without requiring additional qualitatively different mechanisms.

3) *Measuring Looking Time and Checking Behavior*: Tomasello [63] suggested that as infants progress in their

gaze following skills, they should develop a mentalistic understanding of gaze following. He further argued that the most compelling evidence of this would be: 1) looking longer at objects looked at by the caregiver; and 2) alternate gaze between caregiver and her/his object of attention. Although Tomasello did not test these claims, Flom and Pick [26] found that 18-month-olds look for longer periods at an object that the caregiver is looking at, when the caregiver verbally instructs the infant to do so. Because the caregiver’s voice is itself an auditory stimulus that might attract the infant’s attention, the fact that they look during her verbalizations suggests that infants are influenced by the “meaning” of her words: the message to redirect attention. Related findings were reported by [21]. Moreover, [46] and [49] found that toddlers, as well as enculturated adult chimpanzees, tend to alternate gaze to an adult if they are also pointing at something to show or request it. This evidence that apes and human children alternate gaze when trying to direct another’s attention is consistent with Tomasello’s claim that alternation is related to mentalistic uses of attention-sharing.

An alternative possibility is that longer shared-gaze episodes, and gaze alternation, emerge from the learning process of agents with the basic set of phenotypes specified in the model, within a structured social environment. To investigate this, we included in the tests a setup similar to Corkum and Moore’s H + E experiment but using a target of saliency of 1 (as opposed to no target). For trials where the infant model looks at the object after habituating to the caregiver (i.e., trials with successful gaze following), we observed how long it looks at the object before looking elsewhere. Fig. 11 left shows that with time the infant model looks at the target longer, with an eventual leveling off.

The model explains this increase in looking time from an early bias towards exploration. This bias causes the infant model to “keep exploring” even when it is already looking at a salient object. With time, as exploitation kicks in (because values of m diverge from their initial zero values), the infant model will more effectively exploit learned rewards (i.e., salient targets). The infant model will always eventually look away, thereby restricting looking time. This is due to the habituation routine, which systematically reduces the salience of the attended target over time. Therefore, the infant model’s exploration versus exploitation dynamics is the main causation of the increase in the time that the infant model looks at the object, and not the

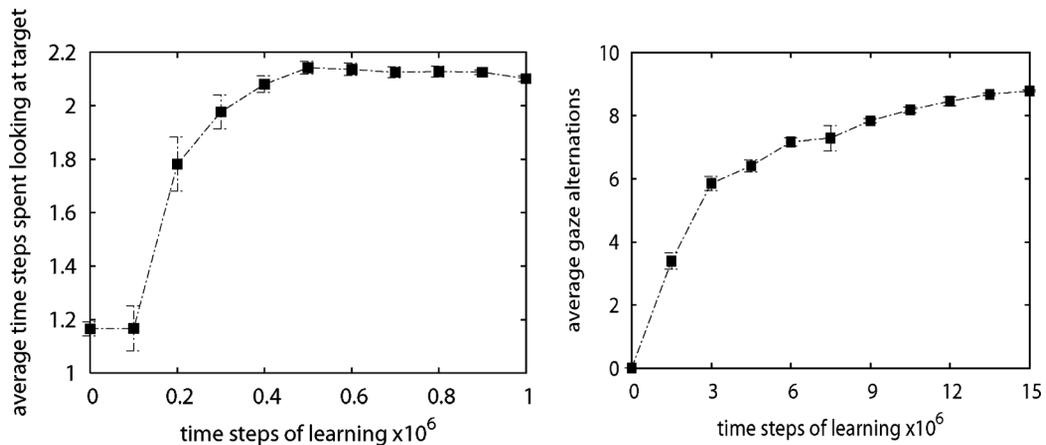


Fig. 11. (Left) Average time looking at target. (Right) Average number of gaze alternations.

fact that the caregiver is looking at the object. In other words, the observed behavior does not require the additional claim by Tomasello that infants take a special interest in an object if someone else is looking at it. Thus, our model predicts that during some period as infants have learned gaze-following, their looking time to salient target will increase, without regard to a caregiver’s intentions.

We also examined, in the same experiment, whether the infant model would alternate gaze between the caregiver and object. For this, we measured the average number of gaze alternations from the caregiver to the target and back, with no intermediate looks to other directions (i.e., if the infant model looks elsewhere, we stopped counting). The right side of Fig. 11 shows that over time, the infant model does more gaze alternations, and this number eventually levels off (note the different time scale with respect to the left side of Fig. 11).

The model explains gaze alternations as stemming from the dynamics of habituation. When the infant model gets “bored” of looking at the caregiver, it looks in the direction of the object (because it is salient and because the caregiver is looking at it). As the infant model habituates to the object, it also dishabituates to the caregiver (because the caregiver is no longer in the infant model’s focus of attention), so that once it gets “bored” of looking at the object, it looks in the direction of the caregiver again (note that in this setup, when looking at one object the caregiver is within sight but the other object is out of sight). This process would repeat itself indefinitely, but since actions are chosen probabilistically, there is always a chance that the infant model will look somewhere else (i.e., explore instead of exploit), thus breaking the cycle. The number of alternations increases over time for the same reason that looking times to objects increase. Exploration diminishes as the infant models learned action-policies permit exploitation, so the alternation cycle become longer and/or less frequent.

4) *Discussion:* As in the previous section, we have shown that specific gaze following behaviors seen in infants also emerge in a fairly simple model, without adding any new mechanisms beyond a small set of infant phenotypes and ecological conditions that are known to exist for some months before infants begin following gaze. The findings of growing sensitivity to eye cues, and increasing gaze-alternation, are

particularly relevant because previous explanations for these behaviors have incorporated a “theory-of-mind” [56]. Of course, these results in no way suggest that primitive mentalistic inferences do not occur during the first two to three years. However, they suggest that there is no compelling reason to attribute even late-emerging gaze-following behaviors to such inferences.

IV. CONCLUSION

We presented a computational model of the development of gaze following. Our model rests on the “basic set hypothesis,” which states that a small number of generic ingredients such as visual preferences, habituation, reinforcement learning, and a proper learning environment are sufficient to explain the emergence of gaze following skills. It is not necessary to postulate an ever-increasing set of new mechanisms to explain the progressive refinement of infant behavior. A small set of generic ingredients is sufficient to explain major aspects of the development of gaze following during the first 18 months.

Learning as a general mechanism for the development of gaze following was first proposed by [53], and our model backs up this conjecture. The first model based on the basic set hypothesis was developed by [67], and pertained to the basic gaze following skill. But space was modeled as discrete locations, with no particular relation between them. By using a 2-D representation of space, which also implied simulating a basic visual input system, our model replicates the geometric and representational stages of gaze following. We do not posit additional mechanisms for that. Instead, we further exploit learning as an explanatory basic mechanism. Preprogrammed models of gaze following such as Cog [7] and Infanoid [44], [43] cannot take advantage of learning. Instead, additional programming would be needed to implement more complex forms of gaze following. Of the computational/robotic models of gaze following that are based on learning, such as [27], [33], [55], [30], and [60], learning is used only to explain the basic skill of gaze following. Ideally, their models should be able to replicate, without modification, further stages of gaze following as our model does. However, this has not been attempted successfully yet. Maybe it was assumed, as Butterworth and others did, that learning only pertained to basic gaze following, and

that subsequent stages did not indicate an improvement of gaze following per se, but of other mechanisms instead. In contrast, our approach has been to assume that *where development in a skill happens, learning should be explored as a central explanation for its appearance as well as for its development*. In other words, introduction of additional “mechanisms” should be avoided in favor of the use of learning to explain the further development of the basic skill. For example, Moore [72] believed that the use of eye direction cues over head direction cues could not be explained by learning, but instead belonged to the set of “theory-of-mind” skills. Our model showed that learning indeed can explain the development of this skill, without trying to integrate some kind of “theory-of-mind” module. And, further in this direction, why not use the basic set hypothesis to model “theory-of-mind” skills? In fact, we have already used it to model social referencing [38], where infants learn to consult the expression of adults before interacting with novel objects. The development in the attribution of false beliefs to others in the false-belief task, for example, might also be subjected to modeling using the same approach. Other visual attention skills observed in infants could also be amenable to similar learning-based approaches, such as the active information selection observed in infants, a strategy that helps them cope with uncertainty and ambiguity in their environment [70].

Beyond this, the value of the model lies in the predictions it makes. Our model previously predicted the emergence of a new class of mirror neurons for gaze-following emerging in its premotor layer [66].³ Such kinds of mirror neurons have indeed been recently found in macaque monkeys⁴ [58]. This lends strong support to our model as a viable candidate for a mechanistic account of the emergence of gaze following.

We originally posed two questions, and will now try to address these. Our response to the question “What is the nature of the mechanisms of gaze following?” is that we believe that the gradual reinforcement-based learning of our model is a satisfactory explanation. Our model not only acquires basic gaze following, but also shows progressive refinement that fits human behavioral findings. Our simulations replicate some classic findings of cumulative changes in gaze following ability. Although such changes may seem categorical and have been associated with different “mechanisms,” our model produces these categorical changes via gradual adaptations in its connection weights. Such an account based on well-studied generic learning mechanisms is attractive due to its simplicity and parsimony.

To the question “What factors guide the developmental schedule of gaze following?” our answer is that detectable structure of events in the social environment affords opportunities for learning. Note that while our model makes some aspects of this learning explicit (changes in how sensory states get mapped onto gaze shifts), other aspects have been treated only implicitly (such as the progressive improvement of the

infant’s ability to estimate the orientation of the caregiver’s eyes and head).

The proof of possibility that a “basic set” of general-purpose elements can lead to a variety of gaze following behaviors also obviates some assumptions about representation made by prior models. Our model does not understand others as intentional agents, or analogize its own actions and the caregiver’s (as in simulation theory [50]). In fact there is no explicit representation of “self,” of intentions, of the caregiver as a separate entity, or of the caregiver’s intentions or thoughts or beliefs. Nonetheless, our model does eventually follow gaze to out-of-sight targets, learn to weight eye direction cues over head direction cues, and do gaze alternation—all of which have been interpreted as signs of a new ability to represent “intersubjectivity.” Given our results, then, we can conclude that these phenomena do not necessarily reveal a new “intersubjective” understanding of others. Of course, we cannot disconfirm that possibility, but they are not strictly *necessary* to explain the behavior. In addition, we do not rule out that the gaze-following behaviors that emerge in our model could be foundational for more sophisticated understanding of others. After all, reacting to others’ gaze and establishing joint attention *is* a kind of understanding of other’s attention, albeit not a discretely conceptual or theoretical one. It is an embodied understanding—embodied, that is, in behavioral tendencies that gradually emerge from learned values of associations between observations and subsequent actions. How this implicit embodied knowledge could be utilized in the development of a more sophisticated understanding of others is a topic for future research. However, we would emphasize the necessity of a structured environment in this process.

It should be noted that our model does not include as input the caregiver’s body, hands, and voice, which play a part in gaze following. However, [21], [20], and [26] found that parents’ verbal exhortations modestly add to the efficacy of gaze- and point-following, but do not qualitatively change any patterns. Also, [48] added hand-and-arm motions to the parent model, and the initial findings suggested that a reinforcement learning model could learn gaze-following routines, even in this more complex environment.

While the model was primarily tested on Butterworth’s and Corkum and Moore’s experimental setups, it can also be tested on similar setups, such as the varied configurations of target setups used by [22], or the ones done by [12] using targets with small degrees of separation. This would lend validity to the general architecture used, making sure that we are not over-generalizing with our solution. And while we believe that our model makes a strong case for eliminating additional unnecessary mechanisms in favor of a learning-based explanation, many of its parameters and corresponding values were based on their utility for replicating experimental results of interest. These parameters could be tuned using results from new studies based on eye-tracking software that have given more detailed measurements of gaze following behavior [34], [29].

REFERENCES

- [1] S. M. Anstis, J. W. Mayhew, and T. Morley, “The perception of where a television portrait is looking,” *Amer. J. Psychol.*, vol. 82, pp. 472–489, 1969.

³The model used in [66] corresponds to an earlier version, with the same architecture and learning mechanism as the one presented here but with minor differences (σ_H and σ_E were not yet introduced, and parameter values differ slightly) The same mirror neuron-like properties found in that model are present in the model shown here.

⁴Compare Fig. 6 in [66] with Fig. 2 in [58].

- [2] S. Baron-Cohen, *Mindblindness: An Essay on Autism and Theory of Mind*. Cambridge, MA: MIT Press, 1995.
- [3] E. Bates, "The emergence of symbols," in *Intentions, Conventions, and Symbols*, L. Camaioni, E. Bates, L. Benigni, I. Bretherton, and V. Volterra, Eds. New York: Academic, 1979, pp. 33–44.
- [4] E. Bates, L. Camaioni, and V. Volterra, "The acquisition of performatives prior to speech," in *Developmental Pragmatics*, E. Ochs and B. B. Schieffelin, Eds. New York: Academic, 1979, pp. 111–129.
- [5] T. B. Brazelton, B. Koslowski, and M. Main, "The origin of reciprocity: The early mother-infant interaction," in *The Effect of the Infant on Its Caretaker*, M. Lewis and L. A. Rosenblum, Eds. New York: Wiley, 1974, pp. 49–76.
- [6] G. W. Bronson, "Changes in infants' visual scanning across the 2- to 14-week age period," *J. Exp. Child Psychol.*, vol. 49, pp. 101–125, 1990.
- [7] R. A. Brooks, C. Breazeal, M. Marjanovic, B. Scassellati, and M. M. Williamson, "The COG project: Building a humanoid robot," in *Computation for Metaphors, Analogy, and Agents of Springer Lecture Notes in Artificial Intelligence*, C. L. Nehaniv, Ed. New York: Springer-Verlag, 1999, vol. 1562, pp. 52–87.
- [8] J. S. Bruner, *Child's Talk: Learning to Use Language*. New York: Norton, 1983.
- [9] N. K. Butko, I. Fasel, and J. R. Movellan, "Learning about humans during the first 6 minutes of life," in *Proc. 5th Int. Conf. Develop. Learning*, Bloomington, IN, Jun. 2006.
- [10] G. Butterworth, "Origins of mind in perception and action," in *Joint Attention: Its Origins and Role in Development*, C. Moore and P. J. Dunham, Eds. Hillsdale, NJ: Erlbaum, 1995, pp. 29–40.
- [11] G. Butterworth and E. Cochran, "Towards a mechanism of joint visual attention in human infancy," *Int. J. Behav. Develop.*, vol. 3, pp. 253–272, 1980.
- [12] G. Butterworth and S. Itakura, "How the eyes, head and hand serve definite reference," *Brit. J. Develop. Psychol.*, vol. 18, pp. 25–50, 2000.
- [13] G. Butterworth and N. Jarrett, "What minds have in common is space: Spatial mechanisms serving joint visual attention in infancy," *Brit. J. Develop. Psychol.*, vol. 9, pp. 55–72, 1991.
- [14] E. Carlson and J. Triesch, "A computational model of gaze following," in *Proc. 8th Neural Comput. Psychol. Workshop, Connectionist Models Cognition Perception II*, Canterbury, U.K., Aug. 2003.
- [15] A. J. Caron, S. C. Butler, and R. Brooks, "Gaze following at 12 and 14 months: Do the eyes matter?," *Brit. J. Develop. Psychol.*, vol. 20, pp. 225–239, 2002.
- [16] M. G. Cline, "The perception of where a person is looking," *Amer. J. Psychol.*, vol. 80, pp. 41–50, 1967.
- [17] V. Corkum and C. Moore, "Development of joint visual attention in infants," in *Joint Attention: Its Origins and Role in Development*, C. Moore and P. J. Dunham, Eds. Hillsdale, NJ: Erlbaum, 1995, pp. 61–83.
- [18] S. Daly, K. Matthews, and J. Ribas-Corbera, "Visual eccentricity models in face-based video compression," in *Proc. IS&T/SPIE Conf. Human Vision Electron. Imaging*, San Jose, CA, Jan. 1999.
- [19] P. Dayan and L. F. Abbott, *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. Cambridge, MA: MIT Press, 2001.
- [20] G. O. Deák, R. Flom, and A. D. Pick, "Perceptual and motivational factors affecting joint visual attention in 12- and 18-month-olds," *Develop. Psychol.*, vol. 36, pp. 511–523, 2000.
- [21] G. O. Deák, T. A. Walken, M. Yale, and A. Lewis, "Driven from distraction: How infants respond to parents attempts to elicit and re-direct their attention," *Infant Behav. Develop.*, vol. 31, pp. 34–50, 2008.
- [22] B. D'Entremont, "A perceptual-attentional explanation of gaze following in 3- and 6-month-olds," *Develop. Sci.*, vol. 3, pp. 302–311, 2000.
- [23] T. Farroni, G. Csibra, F. Simion, and M. H. Johnson, "Eye contact detection in humans from birth," in *Proc. Nat. Acad. Sci. USA*, 2002, vol. 99, pp. 9602–9605.
- [24] I. Fasel, G. O. Deák, J. Triesch, and J. Movellan, "Combining embodied models and empirical research for understanding the development of shared attention," in *Proc. IEEE Int. Conf. Develop. Learning*, Cambridge, MA, Jun. 2002, pp. 21–27.
- [25] S. Feinman, "Social referencing in infancy," *Merrill-Palmer Quart.*, vol. 28, pp. 445–470, 1982.
- [26] R. Flom and A. D. Pick, "Verbal encouragement and joint attention in 18-month-old infants," *Infant Behav. Develop.*, vol. 26, pp. 121–134, 2003.
- [27] M. Fujita and H. Kitano, "Development of an autonomous quadruped robot for robot entertainment," *Auton. Robots*, vol. 5, pp. 7–18, 1998.
- [28] J. J. Gibson and A. Pick, "Perception of another person's looking behavior," *Amer. J. Psychol.*, vol. 76, pp. 386–394, 1963.
- [29] G. Gredeback, C. Theuring, P. Hauf, and B. Kenward, "The microstructure of infants' gaze as they view adult shifts in overt attention," *Infancy*, vol. 13, pp. 533–543, 2008.
- [30] S. Grossberg and T. Vladusich, "How do children learn to follow gaze, share attention, imitate their teachers, and use tools during social interactions?," *Neural Netw.*, vol. 23, pp. 940–965, 2010.
- [31] L. Grover, "Comprehension of the Manual Pointing Gesture in Human Infants," Unpublished Ph.D. dissertation, Dept. Psychol., University of Southampton, Southampton, U.K., 1988.
- [32] M. Hayhoe, M. Land, and A. Shrivastava, "Coordination of eye and hand movements in a normal environment," *Investigat. Ophthalmol. Visual Sci.*, vol. 40, p. S380, 1999.
- [33] M. W. Hoffman, D. B. Grimes, A. P. Shon, and R. P. N. Rao, "A probabilistic model of gaze imitation and shared attention," *Neural Netw.*, vol. 19, pp. 299–310, 2006.
- [34] C. Hofsten, E. Dahlstrom, and Y. Fredriksson, "12-month-old infants' perception of attention direction in static video images," *Infancy*, vol. 8, pp. 217–231, 2005.
- [35] R. J. Itier and M. Batty, "Neural bases of eye and gaze processing: The core of social cognition," *Neurosci. Biobehavioural Rev.*, vol. 33, pp. 843–863, 2009.
- [36] H. Jasso, "A Reinforcement Learning Model of Gaze Following," Ph.D. dissertation, Dept. Comput. Sci. Eng., Univ. of California, San Diego, 2007.
- [37] H. Jasso, J. Triesch, and G. O. Deák, "Using eye direction cues for gaze following—A developmental model," in *Proc. 5th Int. Conf. Develop. Learning*, Bloomington, IN, May 2006.
- [38] H. Jasso, J. Triesch, and G. O. Deák, "A reinforcement learning model of social referencing," in *Proc. IEEE 7th Int. Conf. Develop. Learning*, Monterey, CA, Aug. 2008.
- [39] H. Jasso, J. Triesch, C. Teuscher, and G. O. Deák, "A reinforcement learning model explains the development of gaze following," in *Proc. 7th Int. Conf. Cognitive Model.*, Trieste, Italy, Apr. 2006, pp. 148–153.
- [40] R. Jenkins, J. D. Beaver, and A. J. Calder, "I thought you were looking at me! direction-specific aftereffects in gaze perception," *Psychol. Sci.*, vol. 17, pp. 506–514, 2006.
- [41] M. H. Johnson, S. Dziurawiec, H. D. Ellis, and J. Morton, "Newborn's preferential tracking of face-like stimuli and its subsequent decline," *Cognition*, vol. 40, pp. 1–19, 1991.
- [42] A. Johnston, "Spatial scaling of central and peripheral contrast sensitivity functions," *J. Opt. Soc. Amer. A*, vol. 4, pp. 1583–1593, 1987.
- [43] H. Kozima, "Infanoid: A babybot that explores the social environment," in *Socially Intelligent Agents: Creating Relationships With Computers and Robots*, K. Dautenhahn, A. H. Bond, L. Canamero, and B. Edmonds, Eds. Amsterdam: Kluwer Academic, 2002, pp. 157–164.
- [44] H. Kozima and H. Yano, "A robot that learns to communicate with human caregivers," in *Proc. Int. Workshop Epigenetic Robotics*, Lund, Sweden, Sep. 2001.
- [45] B. Lau and J. Triesch, "Learning gaze following in space: A computational model," in *Proc. 3rd. Int. Conf. Develop. Learning*, La Jolla, CA, Oct. 2004.
- [46] D. A. Leavens and W. D. Hopkins, "The whole-hand point: The structure and function of pointing from a comparative perspective," *J. Compar. Psychol.*, vol. 113, pp. 417–425, 1999.
- [47] J. D. Lempers, "Young children's production and comprehension of nonverbal deictic behaviors," *J. Genetic Psychol.*, vol. 135, pp. 93–102, 1979.
- [48] J. Lewis, G. Deák, H. Jasso, and J. Triesch, "Building a model of infant social interaction," in *Proc. CogSci*, Portland, OR, Aug. 2010, pp. 278–283.
- [49] E. F. Masur, "Gestural development, dual-directional signaling, and the transition to words," *J. Psycholinguistic Res.*, vol. 12, pp. 93–109, 1983.
- [50] A. N. Meltzoff, "'Like me': A foundation for social cognition," *Child Develop.*, vol. 10, pp. 126–134, 2007.
- [51] C. Moore, "Gaze following and the control of attention," in *Early Social Cognition*, P. Rochat, Ed. Hillsdale, NJ: Erlbaum, 1999.
- [52] C. Moore, *The Development of Commonsense Psychology*. Mahwah, NJ: Erlbaum, 2006.
- [53] C. Moore and V. Corkum, "Social understanding at the end of the first year of life," *Develop. Rev.*, vol. 14, pp. 349–372, 1994.
- [54] C. Moore and P. J. Dunham, *Joint Attention: Its Origins and Role in Development*. Hillsdale, NJ: Erlbaum, 1995.
- [55] Y. Nagai, M. Asada, and K. Hosoda, "Learning for joint attention helped by functional development," *Adv. Robotics*, vol. 20, pp. 1165–1181, Sept. 2006.
- [56] D. Premack and G. Woodruff, "Does the chimpanzee have 'a theory of mind'?", *Behav. Brain Sci.*, vol. 4, pp. 515–526, 1978.
- [57] M. Scaife and J. S. Bruner, "The capacity for joint visual attention in the infant," *Nature*, vol. 253, pp. 265–266, 1975.
- [58] S. V. Shepherd, J. T. Klein, R. O. Deaner, and M. L. Platt, "Mirroring of attention by neurons in macaque parietal cortex," in *Proc. Nat. Acad. Sci.*, 2009, vol. 106, no. 23, pp. 9489–9494.

- [59] A. Slater, D. C. Earle, V. Morison, and D. Rose, "Pattern preferences at birth and their interaction with habituation-induced novelty preferences," *J. Exp. Child Psychol.*, vol. 39, pp. 37–54, 1985.
- [60] H. Sumioka, Y. Yoshikawa, and M. Asada, "Reproducing interaction contingency toward open-ended development of social actions: Case study of joint attention," *IEEE Trans. Auton. Mental Develop.*, vol. 2, pp. 40–50, 2010.
- [61] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 1998.
- [62] L. A. Symons, S. M. J. Hains, and D. W. Muir, "Look at me: Five-month-old infants sensitivity to very small deviations in eye-gaze during social interactions," *Infant Behav. Develop.*, vol. 21, pp. 531–536, 1998.
- [63] M. Tomasello, "Joint attention as social cognition," in *Joint Attention: Its Origins and Role in Development*, C. Moore and P. J. Dunham, Eds. Hillsdale, NJ: Erlbaum, 1995, pp. 103–130.
- [64] M. Tomasello and J. Todd, "Joint attention and lexical acquisition style," *First Lang.*, vol. 4, pp. 197–212, 1983.
- [65] C. Trevarthen and P. Hubley, "Secondary intersubjectivity: Confidence, confiding and acts of meaning in the first year," in *Action, Gesture, and Symbol*, A. Lock, Ed. London: Academic, 1978, pp. 183–229.
- [66] J. Triesch, H. Jasso, and G. O. Deák, "Emergence of mirror neurons in a model of gaze following," *Adapt. Behav.*, vol. 15, pp. 149–165, 2007.
- [67] J. Triesch, C. Teuscher, G. O. Deák, and E. Carlson, "Gaze following: Why (not) learn it?," *Develop. Sci.*, vol. 9, pp. 125–147, 2006.
- [68] V. Virsu and J. Rovamo, "Visual resolution, contrast sensitivity, and the cortical magnification factor," *Exp. Brain Res.*, vol. 37, pp. 475–494, 1979.
- [69] C. von Hofsten, E. Dahlström, and Y. Fredriksson, "12-month-old infants' perception of attention direction in static video images," *Infancy*, vol. 8, pp. 217–231, 2005.
- [70] C. Yu, L. B. Smith, H. Shen, A. F. Pereira, and T. Smith, "Active information selection: Visual attention through the hands," *IEEE Trans. Auton. Mental Develop.*, vol. 1, pp. 141–151, 2009.
- [71] J. C. Stanley, "Computer simulation of a model of habituation," *Nature*, vol. 261, pp. 146–148, 1976.
- [72] C. Moore, "Gaze following and the control of attention," in *Early Social Cognition*, P. Rochat, Ed. Hillsdale, NJ: Erlbaum, 1999.

Jochen Triesch received the Diploma and Ph.D. degrees in physics from the Ruhr-University Bochum, Germany, in 1994 and 1999, respectively.



He spent two years as a postdoctoral Fellow in the Computer Science Department of the University of Rochester, NY. In 2001, he joined the faculty of the Cognitive Science Department of the University of California San Diego, La Jolla, CA, as an Assistant Professor. In 2005, he became a Fellow of the Frankfurt Institute for Advanced Studies (FIAS), Frankfurt am Main, Germany. Since 2007 he has been the Johanna Quandt Research Professor for Theoretical Life Sciences at FIAS, where he also serves as the Vice Chairman of the Board of Directors. He also holds Professorships in the Department of Physics and the Department of Computer Science and Mathematics at the Goethe University in Frankfurt am Main, Germany. His research interests span computational neuroscience, cognitive science, developmental robotics, and biologically inspired computer vision.

Prof. Triesch has been helping to organize the IEEE International Conference on Development and Learning several times as either Program Chair or General Chair. He is Associate Editor of the IEEE TRANSACTIONS ON AUTONOMOUS MENTAL DEVELOPMENT, *Neural Networks*, and the *International Journal of Social Robotics* as well as Review Editor of *Frontiers in Computational Neuroscience*. In 2006, he received a Marie Curie Excellence Center Award of the European Union. From 2009–2011 he coordinated the Bernstein Focus Neurotechnology Frankfurt, a large center grant focusing on biologically inspired vision systems.

Gedeon Deák received the B.A. degree in psychology from Vassar College, Poughkeepsie, NY, in 1990, and the Ph.D. degree in child psychology from the University of Minnesota-Twin Cities, Minneapolis, MN, in 1995.



He is an Associate Professor of cognitive science at the University of California San Diego, where he has been director of the Cognitive Development Laboratory (<http://www.cogsci.ucsd.edu/~deak/cdlab/>) since 1999. From 1995 to 1999, he was an Assistant Professor of psychology and human development at Vanderbilt University. In 2011 he was appointed visiting summer faculty at the School of Psychology, Southwest University, Chongqing, China. His and his colleagues' empirical and theoretical work has been described in more than 40 peer-reviewed articles and chapters. He and his collaborators study the dynamics of infant social interactions and social learning. Their studies use experimental behavioral paradigms, ethnographic investigations, physiological studies, and computer simulations. Recently, he and his collaborators conducted the first study of real-time interactions between parents and toddlers while capturing full-spectrum EEG from both participants, synchronized to one another's actions via motion capture. He and his colleagues also study how children learn and use words, and how they flexibly shift their attention, representations, and inferences. Recently, he began a collaborative study with three other labs, tracking the development of language from infancy to early childhood, in bilingual and monolingual children learning English, French, and Spanish.

Prof. Deák previously received a National Academy of Education/Spencer Foundation Postdoctoral Fellowship, and a Junior Faculty Fellowship from the Hellmans. He and his colleagues have received research grants from the National Science Foundation, Spencer Foundation, M.I.N.D. Institute at UC-Davis, National Alliance for Autism Research, Kavli Institute for Mind and Brain, and Nicholas Hobbs Foundation. He serves on the editorial board of the *Journal of Cognition and Language and Child Development Research*, and on the IEEE Technical Committee for Autonomous Mental Development.

Joshua M. Lewis received the B.A. degree in cognitive science and philosophy from Pomona College, Claremont, CA, in 2006, and the Ph.D. degree in cognitive science from the University of California San Diego (UC San Diego), La Jolla, in 2011.

He is currently a postdoctoral student at UC San Diego in the Natural Computation Lab under the direction of P. I. Virginia de Sa. His research interests include the interplay between humans and machine learning algorithms during data analysis.



Hector Jasso received the B.A. degree in computer science from Instituto Tecnológico y de Estudios Superiores de Monterrey, Monterrey, Mexico, in 1988, the M.S. degree in information technology from the University of Edinburgh, Scotland, U.K., in 1991, and the Ph.D. degree in computer science and cognitive science from the University of California San Diego, La Jolla, in 2007.

He has worked as a Researcher at the Instituto de Investigaciones Eléctricas, Cuernavaca, Mexico, as a Data Mining Team Leader at Bancomer,

S.A., Mexico City, Mexico, and most recently as a Research Scientist at the San Diego Supercomputer Center and as a Proprietary Software Developer.

