

To Care or Not to Care: Analyzing the Caregiver in a Computational Gaze Following Framework

Christof Teuscher and Jochen Triesch
University of California, San Diego, Department of Cognitive Science
9500 Gilman Drive, La Jolla, CA 92093-0515, USA
christof@teuscher.ch, triesch@ucsd.edu

Abstract

We first present a computational framework of the emergence of gaze following that is based on a generic basic set of mechanisms. Whereas much attention has been focused so far on the study of the infant’s behavior, we systematically analyze the caregiver and show that he plays a crucial role in the development of gaze following in our model, especially for virtual infants with “developmental disorders”.

We first create two nearly optimal infant parameter sets by means of an evolutionary algorithm and test their behavior with a simple standard caregiver. Based on these findings we then propose new caregiver models and evaluate them on normally developing and on infants with “developmental disorders”.

1 Introduction

Humans and many animals live in social groups, which confers a number of benefits and costs to its members. But living in groups also requires different cognitive abilities for social interactions, which often rely on visual inputs since they are less ambiguous than auditory and olfactory signals and allow for a much richer and more complex communication. One such cognitive skill—with an immediate benefit to the members of a group—is the capacity to imitate other members and to learn from them, reason why humans share with a number of non-human primates the ability to use the eyes, the head, and the body of others to orient to important objects and events in their environment [18, 20]. The ability to follow the direction of conspecifics’ visual gaze does not only help to localize interesting or dangerous entities in the environment, but also provides rich information about the group mates. For example, gaze following can be used

to determine the position of an individual in the dominance hierarchy of large groups, where each individual receives attention as a function of its social rank [7].

Compared to other mammals, human newborns are nearly helpless, but rather quickly begin to show burgeoning social responsiveness [8], and by 3–6 months, infant and caregiver typically engage in complex patterns of reciprocal interaction [14]. By their first birthday, normally developing infants show a robust gaze following [9]. The child then gets more and more adept at recognizing diverse social cues, such as eye direction and pointing gestures. These skills serve as a developmental basis for more complex social communication skills such as the development of the infant’s theory of mind [1] and the development of language, which starts around 13 months and is largely based on joint attentional interactions with adults and objects [4, 24]. Note that in the literature, a distinction is sometimes made between *shared attention* and *joint attention* [10]: in joint attention, two individuals are attending the same object only, whereas shared attention requires each having knowledge of the directions of the other’s attention. In this paper, we only deal with joint attention.

Gaze following has been addressed in a number of computational and embodied models (see Section 1.1) where most attention has been focused on the infant. The goal of this paper is to systematically investigate the role of the caregiver in a computational gaze following framework.

We first present a computational framework (Section 2) of the emergence of gaze following that is based on a generic basic set of mechanisms. In order to have a nearly optimal infant opposite to the fixed caregiver, we optimized the infant’s parameters by means of an evolutionary algorithm in Section 2.5. The outcome are two infants which perform almost equally well, but have different “personalities”. Section 2.6 analyzes the caregiver’s parameters of the original computational model

as presented in [6]. Section 3 describes and analyzes four new caregiver models with respect to our two reference infants, including versions with developmental disorders. Section 4 concludes the paper.

1.1 Related Work

One of the main questions is how shared attention develops and what the necessary and sufficient conditions are. There are basically two alternative hypotheses at the extremes, whereas a combination of both could also be imagined: (1) it is hard-wired into the infant’s brain or (2) it emerges through a learning process while the infant interacts with its environment. The most prominent exponent of the hard-wired hypothesis is Baron-Cohen’s theory of social-cognitive modules [2], which consists of four modules. This nativist and modularist description provides little explanatory and predictive power and does not explain how the modules develop.

In 2001, Matsuda and Omori [16] proposed a reinforcement learning model for acquiring joint visual attention in infants. In their model, the infant’s behavior is supervised and rewarded when it follows the mother’s gaze. Further, they concluded that three abilities are required for joint visual attention to emerge: (1) the recognition of the mother’s face, (2) detection of the eye gaze direction, and (3) a reward system.

Recently Carlson *et al.* [6] proposed a new and more realistic dynamical systems approach which takes into account that complex behaviors can emerge from simple learning mechanisms and which is based on a different and more complete basic set of hypotheses (see Section 2). It basically relies on Moore’s suggestion [17] that gaze following might emerge because infants learn that the caregiver’s direction of gaze is a reliable predictor of where interesting things are located. The proposed approach is the outcome of a more general framework seeking to combine embodied models and empirical research [11]. In contrast to the work of Matsuda and Omori, the infant learns without supervision and only receives a reward when he sees an interesting object or looks at the caregiver’s face (see Section 2.3), i.e., the model therefore learns how to use the caregiver’s face as a predictor for where interesting things are.

Nagai *et al.* [19] present a constructive model by which a real robot acquires the ability of joint attention with a human caregiver without supervised feedback. Although the model develops according to the three developmental stages as proposed by Butterworth and Jarret [5], it cannot deal with ambiguous object situations, therefore Lau and Triesch [15] recently proposed

a new approach which uses the infant’s depth perception to resolve such ambiguities.

A number of other researchers (see for Example [3, 21]) proposed and implemented mechanisms for joint attention for their robots, however, the behaviors were usually fully programmed in advance.

2 A Simple Computational Model

The computational model proposed in [6] relies on a *basic set* of plausible mechanisms [11], which were shown to be sufficient for gaze following to emerge:

- **Perceptual preferences:** Normally developing infants enjoy looking at faces in general and their caregivers in particular. In contrast, many children with autism do not show such a social preference.
- **Habituation:** Infants tend to shift gaze from one target to another after some time.
- **Reward driven learning mechanisms:** The infant shifts gaze between social stimuli and interesting targets in order to maximize internal rewards resulting from visual stimuli.
- **Contingent interactions and a structured environment:** There must be a correlation between where the caregiver is looking and where interesting things are.

The basic set of mechanisms is then formalized within the framework of a biologically plausible temporal difference reinforcement learning algorithm [23]. The framework shall be briefly summarized in the following.

2.1 Environment and Object

The environment is represented as a set of N distinct spatially unattributed regions where exactly one interesting object is present at any discrete time step t in one of the N locations. After an initial fixation time T_{fix} , the object has a relocation probability of p_{shift} at each upcoming time step t . Once relocated to a new location it will remain fixed again for at least T_{fix} time steps before the cycle restarts. Carlson *et al.* [6] used an environment with $N = 10$ regions, $T_{\text{fix}} = 4$, and $p_{\text{shift}} = 0.5$.

2.2 The Caregiver (CG)

The caregiver can either look at the object, at one of the $N - 1$ remaining empty regions, or at the infant. Each time the object is relocated to a new location the caregiver decides to look at either the infant

with probability $\frac{1}{N+1}$, at the object with probability $\frac{N}{N+1}p_{\text{valid}}$, or at an empty location with probability $\frac{N}{N+1}(1 - p_{\text{valid}})$. He then fixates the location or the infant until the target is relocated again as described in the previous section. Carlson *et al.* [6] set p_{valid} to 0.75, which means that the CG spends about 10% of its time looking at the infant, 68% at the object and 22% at an empty location in an $N = 10$ region environment.

2.3 The Infant (INF)

The infant is modeled as a pleasure-driven temporal difference (TD) reinforcement learning agent [23], which tries to maximize the rewards it receives for looking at interesting things. At each time step, the infant can look at one of the N regions or at the caregiver, whereby, the infant can see four possible things: (1) the object, (2) no object, (3) the caregiver’s frontal view, or (4) the caregiver’s profile. Associated with these views are four base reward values $R_{\text{fix}} \in \{R_{\text{object}}, R_{\text{nothing}}, R_{\text{frontal}}, R_{\text{profile}}\}$. The infant can only tell where the caregiver is looking when he directly gazes at him. When the infant is looking at something he habituates to it. For each location N , the caregiver, and the object the infant has a habituation value. As the infant continues to fixate on something, this value decreases, likewise, the infant dishabituates to all other possible looking locations. The habituation change at each time step is given by $h(t+1) = \frac{h(t)}{e^\beta}$, the dishabituation by $h(t+1) = 1 - \frac{1-h(t)}{e^\beta}$. In our experiments, all habituation values are set to 1 at initialization and $\beta = 1$. The actual or instantaneous reward received by the infant at time t is given by $r_{\text{inst}}(t) = R_{\text{fix}}(t)h(t)$.

Based on the work of Findlay and Walker [12], the decision of when and where to shift gaze is implemented in two separate agents in our infant model. The *when-agent* decides whether to continue to fixate on the same location or two shift gaze. It’s state space has two dimensions: (1) the time the infant has been fixated on the current location and (2) the instantaneous reward received by the infant. The *where-agent* provides the new gaze location if the when-agent decides to shift gaze. It has one dimension only, namely, the caregiver’s gaze direction, which can be one out of the N locations, the infant ($N + 1$), or unknown ($N + 2$) in the case the infant is not looking at the caregiver. Both agents make use of a standard TD learning algorithm with tabular SARSA [23] to estimate the state-action values $Q_t(s, a)$:

$$Q_{t+1}(s, a) = Q_t(s, a) + \alpha[r_t + \gamma Q_t(s_{t+1}, a_{t+1}) - Q_t(s_t, a_t)] \quad (1)$$

The action selection is made with a standard *softmax* decision rule, where action a is chosen with probability

$$p_t(a|s) = \frac{e^{\tilde{Q}_t(s,a)/\tau}}{\sum_{a'=1}^N e^{\tilde{Q}_t(s,a')/\tau}}, \quad (2)$$

$$\text{where } \tilde{Q}_t(s, a) = \frac{Q_t(s, a)}{\max_{a'} |Q_t(s, a')|}.$$

The following parameters are important for the model:

1. α , *learning rate*: A small learning rate induces slow learning, a large learning rate oscillations.
2. γ , *discount factor*: Specifies how far in the future rewards should be taken into account. For small γ the agent is only interested in immediate rewards and does not consider long-term consequences of its actions.
3. τ , *temperature*: The lower τ , the more likely it is for the model to chose the action with the highest Q -value (exploitation). For $\tau \rightarrow \infty$, all actions will be chosen with equal probability (exploration).

In many applications these parameters are hand-tuned and fixed, but ideally they should be dynamically adapted to the environment and the agent’s performance. Schweigerhofer and Doya [22], for example, propose a meta-reinforcement learning algorithm to tune α , γ , and τ . For our purpose, the parameters were fixed in order to allow for comparisons between the various experiments.

2.4 Performance Measure

In order to measure the gaze following performance, we will only use the *gaze following index* (GFI) in this paper:

$$\text{GFI} = \frac{\# \text{ gaze shifts from CG to location looked at}}{\# \text{ gaze shifts}}$$

The gaze following index measures the frequency of gaze shifts that lead from the location of the caregiver—where the infant can determine the caregiver’s gaze location—to where the caregiver is looking. During our experiments, the learning was suspended every 1000 time steps and the gaze following performance was tested for 1000 time steps before the learning process resumed.

Note that the maximum GFI with the original caregiver as described in Section 2.2 is $\text{GFI}_{\text{max}} = \frac{N}{2(N+1)}$ as the caregiver spends on the average $\frac{1}{N+1}$ of his time

looking at the infant. For $N = 10$ regions, GFI_{\max} becomes 0.45 on the average. For $N = 2$, the maximum GFI drops to 0.33.

2.5 Finding Optimal Infant Parameters by Means of an Evolutionary Algorithm

Carlson *et al.* [6] used an experimentally determined set of parameters which consisted of the following values: $\alpha = 0.0025$, $\gamma = 0.8$, and $\tau = 0.095$. For our caregiver experiments, however, we were interested in using a nearly “optimal” infant as a vis-à-vis and we therefore decided to use an evolutionary algorithm (EA) [13] to find optimal parameters for α , γ , and τ with regards to the speed of learning and the gaze following index at the end of the learning process.

The optimization was performed by means of a standard elitist genetic algorithm with a generation gap of 0.9, a cross-over probability of 0.7, a mutation rate of 0.01, fitness-based reinsertion, and a single-point cross-over operator. The variables α , γ , and τ use a 20 bit representation for the interval $[0.0001, 1]$ on a binary genotype. The fitness function was defined as the average GFI over T simulation time steps: $\text{fitness} = \frac{1}{T} \sum_{t=1}^T \text{GFI}(t)$.

We ran the algorithm over ten trials for (1) a population size of $\text{ind} = 15$ individuals, $\text{gen} = 50$ generations, $T = 300,000$ simulation steps (5 trials), and (2) a population size of $\text{ind} = 20$ individuals, $\text{gen} = 100$ generations, $T = 50,000$ simulation steps (5 trials).

The results suggest that the algorithm gets easily stuck in a local optimum (independently of the two trial setups) around $\alpha_2 = 0.5$, $\gamma_2 = 0.03$, $\tau_2 = 0.05$, whereas a global optimum seems to be located around $\alpha_1 = 0.075$, $\gamma_1 = 0.5$, $\tau_1 = 0.007$. We decided to simply retain these two showcase candidates without further investigating the EA’s underlying fitness landscape.

Figure 1 shows the evolution of the GFI as a function of four different parameter sets. *Origus* represents the original infant used by Carlson *et al.* [6], *Optimiss* and *Optimuse* are the two parameter sets found by the evolutionary algorithm, whereas *Mediocrus* is an example of a suboptimal solution for comparison (his parameters were experimentally determined).

What can we learn from these results? As one can see, *Optimuse* and *Optimiss* both perform almost equally well and are much faster learners than *Origus* and *Mediocrus*. *Optimuse* learns gaze following slightly faster than *Optimiss*, but both have approximatively the same final GFI. Due to the higher learning rate, his performance slightly oscillates. *Origus* ultimately also reaches almost the same GFI as a the two evolved infants.

Interestingly, although *Optimuse* and *Optimiss* show very similar performance, they have very different parameters. Let us try to characterize their different behaviors in order to understand why they perform almost equally well. *Optimiss* could be described as fast-learning, risky, nearsighted, and exploratory because she has a high learning rate (fast-learning) and a higher temperature (more exploratory) than *Optimuse*. She also takes more “risk” because she prefers immediate rewards (low discount factor). Her smartness, however, allows her to make up for the risk due to the exploratory short-term oriented behavior. *Optimuse*, on the other hand, might be described as tenacious, perspicacious, and exploitive because of her low learning rate and the low temperature. She is interested in getting long-term rewards and therefore closely stands by her policy without taking too much risk and without being seduced by tempting decoys. Interestingly, both of these characteristic personalities perform almost equally well, but with a very different strategy.

Why, one might ask, do *Mediocrus* and *Origus* perform less good? What do they do wrong? According to his parameters, *Mediocrus* is a quick learner, but all the same sticks to his conservative policy (exploitive). One might say that he uses a wrong combination of his capacities and therefore only reaches a GFI of about 0.3. Finally, *Origus* is a really slow learner, but makes otherwise good use of his exploitive behavior and interest in long-term rewards, which ultimately also brings him to the top.

In the remainder of this paper we shall use our two virtual toddlers *Optimuse* and *Optimiss* and their different behavioral strategies to evaluate and compare the performance of various caregivers.

2.6 Analyzing the Parameters T_{fix} , p_{shift} , and p_{valid}

In order to analyze the influence of the caregiver and the environment on the infant’s gaze following behavior we ran several simulations with different parameter sets for the caregiver as described in Section 2.2—who shall be baptized *Ancestrus*—while the infant’s parameters were that of *Optimuse* and *Optimiss* (see previous section).

Figure 2 shows *Ancestrus*’ average gaze following index as a function of each of the three parameters p_{valid} , p_{shift} , and T_{fix} . The parameter values (two were fixed during each run) were as following: $p_{\text{shift}} = 0.5$, $p_{\text{valid}} = 0.75$, and $T_{\text{fix}} = 4$.

As one can see, the predictiveness of the caregiver’s gaze p_{valid} is the sole parameter which significantly affects the infant’s GFI, i.e., for $p_{\text{valid}} < 0.5$, the GFI begins to dramatically drop for both toddlers because




Figure 1. Learning curves for different α , γ , and τ . Average GFI over 10 runs with standard error. The evolved infants *Optimiss* ($\alpha_2 = 0.5$, $\gamma_2 = 0.03$, $\tau_2 = 0.05$) and *Optimuse* ($\alpha_1 = 0.075$, $\gamma_1 = 0.5$, $\tau_1 = 0.007$) perform almost equally well and better than *Origus*. Infant *Mediocrus* is suboptimal.

the caregiver gradually becomes less predictive. On the other hand, T_{fix} and p_{shift} almost have no influence on the GFI, at least not when modified individually¹. *Optimiss* has a slightly lower GFI for all three caregiver parameter sweeps and the GFI begins to drop earlier than for *Optimuse*, which suggests that she is more sensitive to the caregiver’s behavior.

Figure 3 shows four plots for characteristic parameter sets. With *Ancestrus* as a vis-à-vis, *Optimuse* quickly learns to follow the caregiver’s gaze. Caregiver *Ancestrus I* performs less good with $p_{\text{valid}} = 0.2$ because he less often looks at the object. *Ancestrus II* has the same p_{valid} as *Ancestrus* but a relocation probability of $p_{\text{shift}} = 1$ and an object fixation time of $T_{\text{fix}} = 1$, which prevents the infant from learning gaze following because the object is relocated at every time step and the caregiver is “nervously” shifting gaze. Finally, although *Ancestrus III* has a fixation time of $T_{\text{fix}} = 1$ only and a low p_{shift} , the infant learns gaze following fairly well due to $p_{\text{valid}} = 0.3$ (compare also with the average GFI of Figure 2).

Simulations for *Optimiss* show a different picture for *Ancestrus I* and *Ancestrus III*: they have a much lower GFI (average of about 0.1 and 0.15 respectively) than for *Optimuse*, which confirms the foregoing finding that

¹Note that the minimal T_{fix} is 1, i.e., the object moves at each time step if $p_{\text{shift}} = 1$.

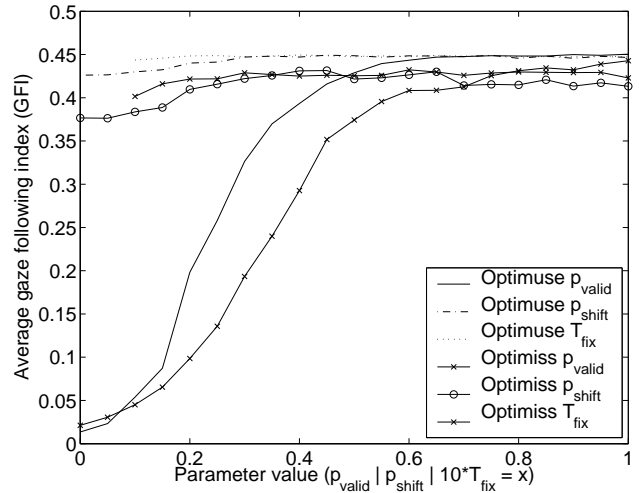


Figure 2. Average GFI as a function of p_{shift} , p_{valid} , and T_{fix} over $T = 100,000$ time steps for *Optimuse* and *Optimiss*. The fixed values were: $p_{\text{shift}} = 0.5$, $p_{\text{valid}} = 0.75$, and $T_{\text{fix}} = 4$. Average GFI over 4 runs without standard error for better legibility.

she is more sensitive to the caregiver’s behavior and well adapted to *Ancestrus* only.

It can be concluded—what seems intuitively obvious for this setting and these toddlers—that the more predictive and structured the environment is (i.e., the higher p_{valid}), the better and faster the infant learns gaze following. Thereby, the predictiveness of the caregiver plays a crucial role, whereas the object’s behavior is less important.

In the next section we will analyze four new caregivers and compare them with our two toddlers in combination with two developmental disorders.

3 New Caregiver Models

The caregiver used so far did not adapt to the infant’s behavior, which certainly represents a gross oversimplification of a real mother-infant interaction. One might hypothesize that an adaptive caregiver would allow the infant to learn gaze following faster and more reliably. In order to test this hypothesis, we implemented several caregiver models and environments and evaluated the gaze following performance of normally developing and autistic infants as well as of children with the Williams syndrome. The caregivers used are as following:

- *Ancestrus*: Original caregiver as described in Section 2.2.

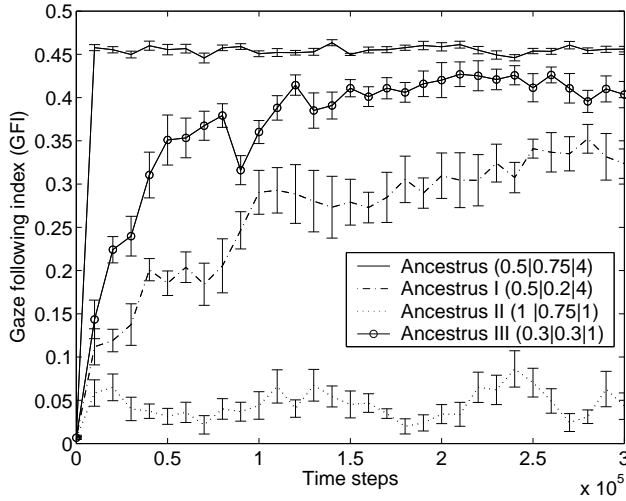


Figure 3. *Optimuse's* learning curves for four different caregiver parameter sets p_{shift} , p_{valid} , and T_{fix} . Average GFI over 10 runs with standard error, $\alpha = 0.075$, $\gamma = 0.5$, $\tau = 0.007$.

- *Randomus*: Random caregiver and object, i.e., at each time step the object and the caregiver's gaze are individually moved to a random location. The caregiver therefore looks at the infant with probability $\frac{1}{N+1}$ and even less rarely looks at the infant.
- *Careus*: This caregiver waits until he can establish mutual gaze contact with the infant. He then moves the object to a random location and directs his gaze to the very same location where he waits until the infant looks at the object. The caregiver then returns his gaze to the infant and waits again for mutual gaze contact.
- *Avoidus*: Same as *Ancestrus*, but he never looks at the infant.
- *Boreus*: The object and the caregiver's gaze both move together, stepwise, and indefinitely from location 1 to N and back, one step each $T_{\text{fix}} = 4$ time steps.

Figure 4 shows *Optimuse's* gaze following performance vis-à-vis of the five above described caregivers. One can see that the random caregiver *Randomus* does not provide a sufficiently structured environment to the infant, whereas *Ancestrus* only allows the infant to slowly acquire gaze following. *Careus* performs best as he adapts to the infant and “guides” his gaze. *Avoidus* still allows the infant to learn gaze following very well, although he never establishes mutual gaze

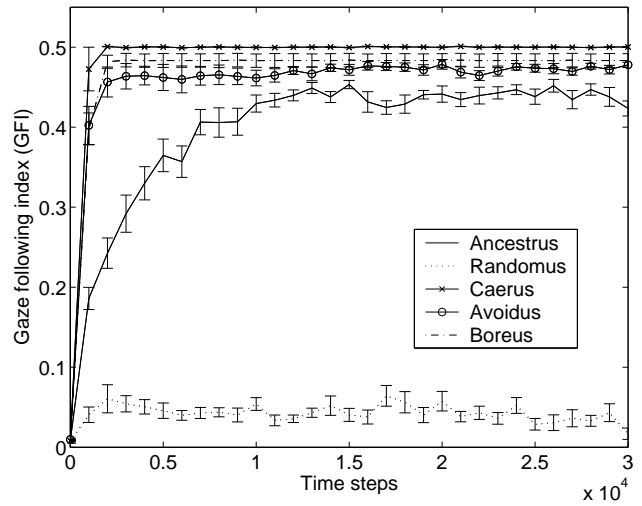


Figure 4. *Optimuse's* GFI for the new caregivers. Average over 10 runs with standard error, $\alpha = 0.075$, $\gamma = 0.5$, $\tau = 0.007$.

avg GFI	<i>Optimuse</i>			<i>Optimiss</i>		
	N	A	W	N	A	W
<i>Ancestrus</i>	0.44	0.004	0.05	0.41	0	0.38
<i>Randomus</i>	0.04	0	0.05	0.04	0	0.04
<i>Careus</i>	0.5	0	0.5	0.5	0	0.5
<i>Avoidus</i>	0.48	0.27	0.42	0.43	0.05	0.5
<i>Boreus</i>	0.47	0.45	0.24	0.49	0	0.5

Table 1. Summary of the average GFI for the two toddlers facing the new caregivers. Average values over $T = 300,000$ time steps and over two runs. N=Normal, A=Autist, W=Williams syndrome.

contact. Similarly, *Boreus* provides a highly structured environment because of the deterministic object trajectory and his fully predictive gaze direction. Columns “N” (normal) of Table 1 summarize the simulation results for this experiment for both infants. As one can see, the results are very similar for *Optimiss*.

3.1 Developmental Disorders

Some autistic children show little or no eye contact and tend to avoid looking at faces whereas children with the Williams syndrome have an abnormally high preference for faces. Note that there seems to exist controversial evidence whether autistics perceive direct gaze as aversive or not. Carlson *et al.* [6] have demonstrated that simple changes in the infant's re-

ward structure can lead to behaviors reminiscent of autism and the Williams syndrome. In this paper we used the following reward structures:

- **Autist:** $R_{\text{frontal}} = -1, R_{\text{profile}} = 0$;
- **Williams syndrome:** $R_{\text{frontal}} = 2, R_{\text{profile}} = 2$.

Columns “A” (autist) and “W” (Williams syndrome) of Table 1 summarize the simulation results for this experiment for both infants. In addition, Figure 5 illustrates the GFI of *Optimuse* as an autist and as a toddler with Williams syndrome. She faces *Ancestrus*, *Avoidus* for the autist, and *Careus* for the Williams toddler. As one can see, gaze following does not emerge with *Ancestrus* and the autist, whereas the Williams toddler only learns it badly. *Avoidus*, however, helps the autist to successfully learn gaze following because he avoids mutual gaze contact (i.e., avoids negative rewards for the infant), whereas *Careus* succeeds in “guiding away” the Williams toddler from staring at his face to the object.

From Table 1 we can also see that the random caregiver *Randomus* is unsuccessful in all situations. Looking at *Optimuse*, we find that *Boreus* performs even better than *Avoidus* for the autist because he too never looks at the infant, but provides an even more deterministic behavior. For the same infant with Williams syndrome, *Avoidus* performs also very well because he avoids mutual gaze.

The situation is a little different for *Optimiss*. In her autistic version, she prevents any gaze following to evolve whereas she does a much better job as a “Williams infant”: all caregivers, except *Randomus*, succeed in “teaching” her gaze following. This is somehow surprising as she was the one who was more sensitive to the environment. However, this might exactly be the explanation: she is more likely to be disturbed and drifted away, which is beneficial for a Williams toddler, i.e., to shift attention away from the highly attracting face.

4 Conclusions

We presented a computational gaze following framework as first introduced by Carlson *et al.* [6] and optimized the infant’s parameters by means of an evolutionary algorithm. The outcome were two infants which performed almost equally well, but used different strategies. We then analyzed the original and several new caregiver models and showed that the caregiver behavior plays a crucial role in the development of gaze following, especially for virtual infants with “developmental disorders”. A second finding is that no “univer-

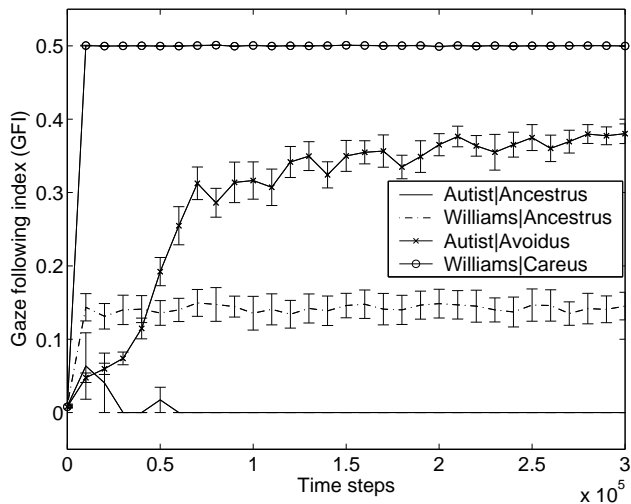


Figure 5. GFI of disordered *Optimuse* who faces caregivers *Ancestrus*, *Avoidus*, and *Careus*. Average over 10 runs with standard error, $\alpha = 0.075, \gamma = 0.5, \tau = 0.007$.

sally” optimal caregiver exists: every virtual toddler—whether with a developmental disorder or not—has its own needs and requires a particular caregiver. This immediately suggests that the ideal caregiver should itself be a learning agent that constantly adapts to the developing infant.

This finding might seem intuitively obvious and simply goes into the direction of developing more realistic computational models of the emergence of gaze following, and eventually shared attention. Nevertheless, the current simplicity of the model can also be considered as a strength since it brings the computational essence of the underlying mechanisms into focus.

Acknowledgments

This work described here is part of the UC San Diego MESA project (<http://mesa.ucsd.edu>) and is supported by the National Alliance for Autism Research (NAAR). Christof Teuscher was supported by the Swiss National Science Foundation under grant PBEL2-104420. The authors are grateful to Gedeon Deák, Hector Jasso, Boris Lau, and Ursina Teuscher for their helpful comments and discussions.

References

- [1] S. Baron-Cohen. Development of a theory of mind: Where would we be without the intentional

- stance? In M. Rutter and D. Hay, editors, *Development Through Life: A Handbook for Clinicians*, pages 303–308. Blackwell Scientific, Oxford, UK, 1994.
- [2] S. Baron-Cohen. The eye direction detector (EDD) and the shared attention mechanism (SAM: Two cases for evolutionary psychology). In C. Moore and P. J. Dunham, editors, *Joint Attention: Its Origins and Role in Development*, chapter 6, pages 41–59. Erlbaum, Hillsdale, NJ, 1995.
- [3] C. Breazeal and B. Scassellati. Infant-like social interactions between a robot and a human caregiver. *Adaptive Behavior*, 8(1):49–74, 2000.
- [4] J. Bruner. *Child’s Talk*. Norton, New York, 1983.
- [5] G. Butterworth and N. Jarrett. What minds have in common is space: Spatial mechanisms serving joint visual attention in infancy. *British Journal of Developmental Psychology*, 9:55–72, 1991.
- [6] E. Carlson and J. Triesch. A computational model of the emergence of gaze following. In H. Bowman and C. Labiouse, editors, *Proceedings of the 8th Neural Computation and Psychology Workshop*, Progress in Neural Processing. World Scientific, 2003.
- [7] M. R. A. Chance. Attention structure as the basis for primate rank orders. *Man*, 2(503–518), 1967.
- [8] M. Cole and S. Cole. *The Development of Children*. Freeman, New York, 3 edition, 1996.
- [9] G. O. Deák, R. Flom, and A. D. Pick. Perceptual and motivational factor affecting joint visual attention in 12- and 18-month-olds. *Developmental Psychology*, 36(511–523), 2000.
- [10] N. J. Emery. The eyes have it: the neuroethology, function and evolution of social gaze. *Neuroscience and Biobehavioral Reviews*, 24:581–604, 2000.
- [11] I. Fasel, G. O. Deák, Jochen Triesch, and J. Movellan. Combining embodied models and empirical research for understanding the development of shared attention. In *Proceedings of the 2nd International Conference on Development and Learning (ICDL02)*, pages 21–27, Los Alamitos, CA, 2002. IEEE Computer Society Press.
- [12] J. M. Findlay and R. Walker. A model of saccade generation based on parallel processing and competitive inhibition. *Behavioral and Brain Sciences*, 22:661–674, 1999.
- [13] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Redwood City, CA, 1989.
- [14] K. Kaye. *The Mental and Social Life of Babies*. University of Chicago Press, Chicago, 1982.
- [15] B. Lau and J. Triesch. Learning gaze following in space: A computational model. Third International Conference on Development and Learning (ICDL’04), La Jolla, California, October 20–22, 2004.
- [16] G. Matsuda and T. Omori. Learning of joint visual attention by reinforcement learning. In E. M. Altmann and A. Cleeremans, editors, *Proceedings of the Fourth International Conference on Cognitive Modeling*, pages 157–162, Mahwah, NJ, US, 2001. Lawrence Erlbaum Associates.
- [17] C. Moore. Theories of mind in infancy. *British Journal of Developmental Psychology*, 14(19–40), 1996.
- [18] C. Moore and P. J. Dunham, editors. *Joint Attention: Its Origins and Role in Development*. Erlbaum, Hillsdale, NJ, 1995.
- [19] Y. Nagai, K. Hosoda, A. Morita, and M. Asada. A constructive model for the development of joint attention. *Connection Science*, 15(4):211–229, December 2003.
- [20] M. Scaife and J. S. Bruner. The capacity of joint visual attention in the infant. *Nature*, 253:265–266, 1975.
- [21] B. Scassellati. Theory of mind for a humanoid robot. *Autonomous Robots*, 12:13–24, 2002.
- [22] N. Schweighofer and K. Doya. Meta-learning in reinforcement learning. *Neural Networks*, 16:5–9, 2003.
- [23] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. A Bradford Book, MIT Press, Cambridge, MA, 1998.
- [24] M. Tomasello and J. Todd. Joint attention and early lexical acquisition style. *First Language*, 4:197–212, 1983.