Narrowing of the Cone-of-Direct Gaze Through Reinforcement Learning

Sam Blakeman (sam.blakeman.15@ucl.ac.uk) & Denis Mareschal (d.mareschal@bbk.ac.uk)

Department of Psychological Sciences, Birkbeck, University of London

Abstract

The Cone of Direct Gaze (CoD) is described as the range of eye gaze deviations over which an observer reports gaze as being directed towards them. The CoD has been found to narrow with age across childhood (Mareschal et al. 2016). We investigated whether reinforcement learning, so critical in shaping eye gaze responses in infancy, was able to account for the emergence of a CoD and its narrowing in childhood. To this end, we adapted Triesch et al.'s (2006) reinforcement learning model by (1) defining a topology over object locations, and (2) introducing opponent non-linear reward profiles for looking at objects and caregivers. In Simulation 1 we show that these modifications give rise to a functional CoD in which there is reduced eye gaze following and increased fixation on the caregiver for locations with a small caregiver eye gaze eccentricity. In Simulation 2 we show that the width of this effect reduces with learning, suggesting that developmental decreases in the CoD may be driven by reinforcement learning. In Simulation 3 we explore how changes in model parameters can explain the CoD in high anxiety populations. Finally, the model provides one way of unifying the developmental gaze-following and CoD literatures, until now considered largely independent.

Keywords: Reinforcement Learning; Cone of Direct Gaze; Gaze Following; Development; Social Anxiety; Autism;

Background

The eyes are a key aspect of social intelligence. From the eyes one can infer an individual's emotions or desires, which can help guide social behavior, and also learn about the surrounding environment (Shepherd et al., 2010). Through joint attention, individuals can alert others to interesting objects in the environment by guiding their attention to that object. Eye gaze following is one form (Scaife and Bruner, 1975). In this seminal study the authors showed that infants were able to interpret the direction of another individual's eye gaze could and use that as a cue to look in the perceived direction. This allows infants to find objects of interest in the environment and learn from experienced caregivers in a non-verbal manner. Indeed, infants have been shown to be sensitive to eye gaze from a very young age, appearing to show a preference for eyes over other parts of the face (Hains & Muir, 1996).

Infants with autism have a reduced ability to follow eye gaze (Leekam et al., 1997). A recent study by Thorup et al. (2016) found that infants at high risk of developing autism rely disproportionally on directional information from the head as compared to the eyes. This reduced ability to follow eye gaze may be a contributing factor to the deficits in social cognition and communication associated with autism.

While joint attention via eye gaze following appears to be a crucial tool for the developing *infant*, the perception of eye gaze direction is not uniform across eye gaze deviations. The Cone-of-Direct gaze (CoD), is defined as the range of gaze deviations that we perceive to be looking directly at us (Gamer and Hecht., 2007). The CoD, therefore, has implications for how we perceive social situations and our interpretation of eye gaze. For example, if a gaze deviation falls inside our CoD then we may perceive it as looking directly at us and not engage in any eye gaze following behavior.

The perception of whether an individual is looking directly at you or not is also of particular interest to those investigating social anxiety disorders (Schulze et al., 2013). For example, a study by Jun et al. (2013) reported a wider CoD for high socially anxious males compared to low socially anxious males. Similarly, Gamer et al. (2011) conducted a study where participants had their CoD measured in response to a virtual head. They found that participants with social phobia had a wider CoD in the presence of a second virtual head that was directed at them. Such studies suggest that a wider CoD is associated with social anxiety and may play a role in the disorder.

Both eye gaze following and the CoD undergo changes during development. Eye gaze following emerges and improves during infant development (Brooks and Meltzoff., 2005, Deak., 2015), while the CoD becomes narrower during childhood (Mareschal et al., 2016). This equates to older children being more reliable at following eye gaze and more accurate at interpreting small eye gaze deviations as not being directed at them. It is possible that these developmental timelines for eye gaze following and the CoD are crucial for infant and child development and may be altered in clinical disorders such as autism and social anxiety. It is therefore important to understand their emergence and developmental trajectory.

Reinforcement learning (Sutton and Barto, 1998) has received much interest in recent years and there is now good evidence of the neurocomputational basis of reinforcement learning (Schultz et al., 1997). It has been proposed as a possible mechanism for the emergence and improvement of eye gaze following in infants. Triesch et al. (2006), describe a reinforcement learning model in which rewards obtained by following a caregiver's gaze led to the reinforcement of eye gaze following behavior. According to this account, infants associates the rewarding object that the caregiver is looking at with the act of following the caregiver's gaze, thereby building a predisposition to follow gaze as a consequence of experience rather than an innate behavior.

While reinforcement learning may account for the emergence of eye gaze following, the mechanism behind the emergence of a CoD is yet to be elucidated. To what extent is the CoD and its development the result of an innate prior

and to what extent is it the result of learning and the external environment? To investigate such a question, we explored whether the reinforcement learning framework could also account for the CoD and its changes through development. If reinforcement learning where to play a role in the emergence of a CoD then it may provide a link between eye gaze following and the CoD. It would also highlight reinforcement learning as a promising target for the therapeutic investigation of disorders such as autism and social anxiety.

Triesch et al.'s (2016) Model

Triesch et al.'s (2006) model serves as a spring board for this study. The model consists of an infant, a caregiver and an object (Figure 1). Both the infant and caregiver remain in fixed positions while the object is able to move around N discrete locations. Two parameters, T_min and p_shift are responsible for the movement of the object around these locations. T_min specifies the minimum amount of time an object must spend in a location, while p_shift specifies the probability of shifting to a new location per time step after T_min. This shifting of the object also determines the shifting of the caregiver's gaze.
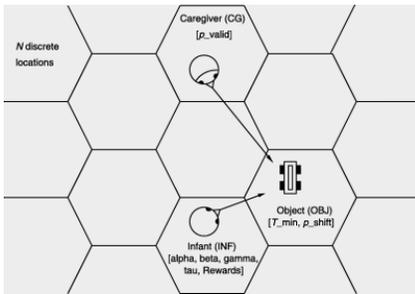

Figure 1. *Diagram of the gaze following reinforcement learning model proposed by Triesch at al. (2006).*

When the object moves to a new location the caregivers gaze is shifted to a new location. The caregiver can look at N+1 potential locations; the N locations the object can reside in plus the location of the infant. The parameter p_valid determines the probability that the new location of the caregiver's gaze is the same as the new location of the object. This probability effectively models two scenarios. The first scenario is when the caregiver's gaze may not be a 100% predictive of where the object is and the second scenario is when the infant's interpretation of the caregiver's gaze may not be a 100% accurate. The p_valid parameter accounts for both of these scenarios because both an inaccurate caregiver gaze or a poor interpretation of the caregiver gaze will lead to the infant following the caregivers gaze to an incorrect location.

The infant's behavior is modelled using a reinforcement learning framework whereby it is essentially driven to maximize rewards in the environment. The infant is broken up into two agents, a 'when' agent and a 'where' agent. The when agent is responsible for deciding whether it is time to shift gaze on a time step and the where agent is responsible

for deciding where to shift the gaze. These decisions are driven by the rewards encountered in the environment by the infant. In this environment the infant has four possible views, each of which having an associated reward:

1. An empty location ($R_{nothing}$)
2. A location containing the object (R_{object})
3. A profile view of the caregiver ($R_{profile}$)
4. A frontal view of the caregiver as they look directly at the infant ($R_{frontal}$)

For each of these views the infant receives the associated reward (R_x) multiplied by a habituation value. This habituation value exponentially decreases as the infant fixates on a location. The degree of this decrease is controlled by the habituation parameter beta (β). Equally, the reward value for locations that have been habituated to, but the infant is no longer looking at, recover at the same rate. Habituation is important in a reinforcement learning framework such as this because otherwise the infant could just fixate on a single reward (e.g. the caregiver), and never have the motivation to shift gaze.

Taking the reward structure and habituation into account, the state-space of the *when* agent becomes two dimensional. The first dimension is how long has the infant been looking at the same location and the second dimension is the reward received by the infant. Representing the state-space with these two dimensions allows the *when* agent to decide whether to carry on looking at the same location or look somewhere else. If the decision is made to look somewhere else, then the *where* agent then specifies the location of the new gaze. The state-space of the *where* agent varies along a single dimension, which represents the gaze of the caregiver. This corresponds to N+2 states. There are N number of states for when the caregiver is looking at each of the N object locations. It is these states that the infant uses to interpret where the caregiver is looking. Another state is for when the caregiver is looking directly at the infant and a final state is for when the gaze of the caregiver is unknown to the infant. While the action-space of the *when* agent is simply stay or move, the action-space of the *where* agent is of size N+1. The *where* agent can decide to shift gaze to one of the N object locations or to look directly at the caregiver.

Both the *when* and *where* agents learn using temporal difference (TD) learning and the SARSA algorithm (Rummery and Niranjan, 1994; Equation. 1).

$$\delta_t = r_t + \gamma Q_t(s_{t+1}, a_{t+1}) - Q_t(s_t, a_t)$$
$$Q_{t+1}(s_t, a_t) = Q_t(s_t, a_t) + \alpha \delta_t$$

$s_t = state\ at\ time\ t$
$a_t = action\ at\ time\ t$
$r_t = reward\ at\ time\ t$
$\delta_t = temporal\ difference\ error\ at\ time\ t$
$Q_t(s_t, a_t) = state - action\ value\ at\ time\ t$
$\alpha = learning\ rate$
$\gamma = discount\ factor\ for\ future\ rewards$

In this framework the temporal difference error is calculated and the parameter gamma (γ) is used as the discount factor

2

for future rewards. This temporal difference error is then used to update the appropriate state-action value and the size of this update is controlled by the learning parameter alpha (α). A softmax function was used to mapped the state-action values to actions. This allowed for a balance between exploration and exploitation, determined by the parameter tau was responsible for this balance. A larger value of tau results in more exploration and less exploitation of the state-action values.

Simulation 1 – Emergence of the Cone of Direct Gaze

To investigate whether reinforcement learning could also account for the emergence of a CoD, the reward structure of Triesch et al.'s model was modified to include a spatial topology in the reward space

Methods

A CoD is inherently a spatial phenomenon that is assessed by having an individual look further and further away from a participant until the participant judges the gaze to no longer be directed at them. For this reason, the N object locations in Triesch et al.'s model were first given a spatial location identity based on the gaze deviation required from the caregiver to look at them (Figure 2). Values stepped four degrees at a time and the N locations were arranged in a linear manner. This alteration to Triesch et al.'s model allowed for the analysis of gaze following behaviour based on the spatial location of the caregiver's gaze.
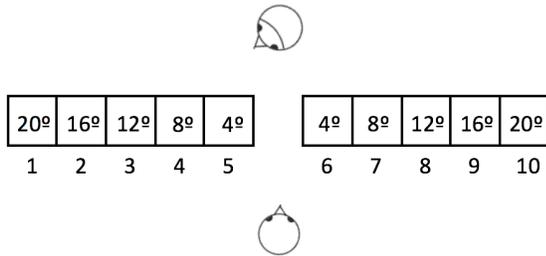

Figure 2. *Layout of the object locations in the model.*

Next, the model's reward structure was modified. Specifically, the caregiver profile reward ($R_{profile}$) and object reward (R_{object}) were changed in an opposing manner using Gaussian functions (Figure 3).

Various studies have shown that infants prefer direct gaze over averted gaze (e.g., Farroni et al., 2002), which in the model equates to $R_{frontal} > R_{profile}$. In terms of a CoD, small eye gaze deviations are likely to be interpreted as being direct and so $R_{profile}$ should have a higher reward value for small eye gaze deviations compared to large eye gaze deviations. A Gaussian function was therefore applied to $R_{profile}$ so that it increased in value as the caregiver looked at locations which required smaller eye gaze deviations.

The opposite transformation was applied to R_{object} so that it decreased in value as the caregiver looked at locations which required smaller eye gaze deviations. This aimed to represent the fact that objects outside of the infant's current

visual field are more likely to be unexpected (so informative) and therefore more rewarding than objects that currently reside in the visual field. These two modifications to the reward structure had the net effect of increasing the caregiver's relative reward at smaller eye gaze deviations and increasing the object's relative reward at large eye gaze deviations.
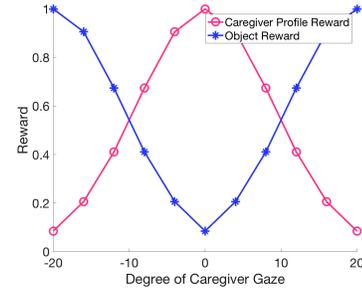

Figure 3. *The modified reward structure for the object and the profile view of the caregiver.*

To evaluate the effect of these modifications 500 simulations were run for 100,000 learning iterations to establish gaze following. After learning, each simulation was run for 10,000 iterations without any learning to gather stable gaze following measurements, which were then averaged across simulations. The model parameters were as follows: Number of locations (N)= 10; Degree step per location (D)= 4; Reward for looking at empty location ($R_{nothing}$) = 1; Reward for looking at the frontal view of the caregiver ($R_{profile}$)=1; Peak Reward for looking at the object (R_{object})= 1; Sigma of the Gaussian applied to the object reward (S_{object})= 9; Sigma of the Gaussian applied to the caregiver ($S_{profile}$)=9; Habituation rate (β) =0.5; Learning Rate (α) =0.0025; Discount Factor (γ) =0.8; Exploration vs. exploitation(τ)=0.095; Minimum fixation time (T_{min})=4; Probability of shifting (p_{shift})=0.5; Predictiveness of caregiver gaze (p_{valid})=0.75. Unless otherwise stated, these values were used in all simulations

Results

Two measurements were used to assess the effect of the modified reward structure. The first measurement was the mean time spent by the infant fixating on the caregiver. This represented how long the infant looked at the caregiver before shifting gaze and served as an indirect measure of the probability of shifting gaze. The second measurement was the total number of gaze follows made by the infant. This was a direct measure of eye gaze following behaviour. Both of these measurements were examined as a function of object location.

After implementing the reward structure in Figure 3, the mean time spent fixating on the caregiver was larger when the caregiver was looking at locations that required small eye gaze deviations (Figure 4, left panel). This contrasted with the model's performance when endowed with a flat reward profile. In addition, the total number of gaze follows

3

was smaller when the caregiver was looking at locations that required small eye gaze deviations (Figure 4, right panel). These findings are consistent with the concept of a CoD. By fixating on the caregiver for longer during small eye gaze deviations the infant acts as if the caregiver is looking directly at them and is unable to follow their gaze to another location. Similarly, the increased number of eye gaze follows for large eye gaze deviations indicates that the infant is correctly classifying them as indirect and can therefore follow them to the object. These findings suggest that a CoD can emerge under a reinforcement learning framework where the caregiver and object rewards act in an opposing manner.
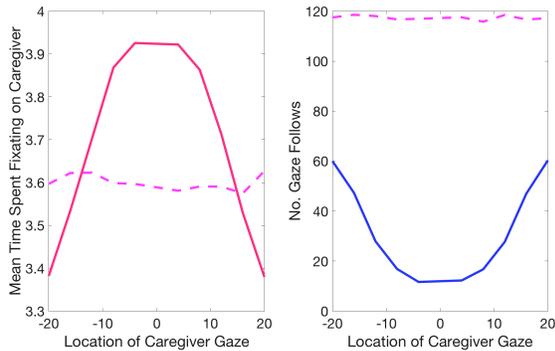

Figure 4. *(Left panel) Mean time spent by the infant fixating on the caregiver as a function of the caregiver's gaze. The dotted magenta line represents the results when the object reward and caregiver profile reward are not modified, as in the original Triesch et al. (2006) model. (Right pane) Total Number of gaze follows made by the infant as a function of the caregiver's gaze. The dotted magenta line represents the results when the object reward and caregiver profile reward are not modified.*

Simulation 2 – Developmental Trajectory of the Cone of Direct Gaze

After confirming that reinforcement learning could lead to the emergence of a CoD, we investigated the effect of reinforcement learning on the CoD over time to see if it could also explain known developmental changes.

Methods

In order to get a measure of the width of the induced CoD, the mean fixation duration and the number of gaze follows were overlaid and their intersects calculated. To achieve this, it was first necessary to rescale the feature so that both measurements were operating on the same scale (Equation 2). Each value had the minimum value subtracted and this was then divided by the range of the values. This produced a final value that ranged between 0 and 1. Gaussian curves were then fit to both feature scaled measures, with the mean fixation time requiring a single term and the number of gaze follows requiring two terms. Finally, the two intersection points of the fitted Gaussian curves were calculated and the width between the two points

was taken as a proxy for the width of the CoD in the model (Figure 5).

Equation 2
$$x' = \frac{x - \min(x)}{\max(x) - \min(x)}$$

To observe the change in this width over time, 500 simulations were run for 1,000,000 learning iterations. At 100,000 learning iteration intervals, learning was halted and 10,000 iterations were run to gather stable gaze following measurements. These results were averaged across all simulations for each break in the learning process.
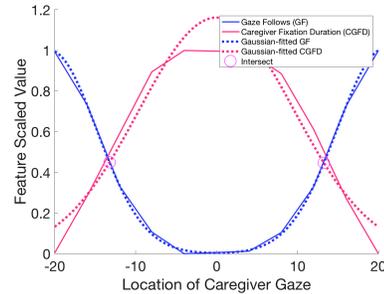

Figure 5. *Quantification of the width of the CoD effect. After feature scaling each measure, Gaussians were fit to both of them. The intersect points of these Gaussians were then calculated and the horizontal distance between the intersection points was taken as a proxy for the width of the CoD effect.*

Results

The CoD width was found to decrease as the number of learning iterations increased (Figure 6). This is consistent with the finding that the CoD decreases during child development (Mareschal et al., 2016) and suggests that reinforcement learning may be one explanation for these changes.
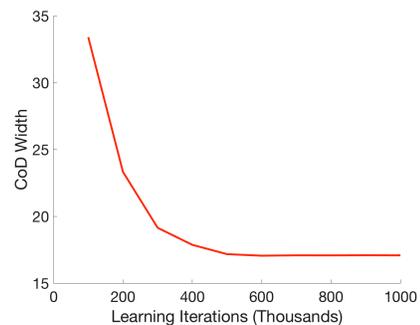

Figure 6. *Change in the width of the CoD effect as a function of the number of learning iterations.*

Simulation 3 – High Anxiety Populations

In this simulation, we explore different parameter values in an attempt to capture known differences in the CoD for individuals with social anxiety.

4

Methods

The peak reward value for both the frontal and caregiver rewards ($R_{frontal}$, $R_{profile}$) were systematically decreased. For each set of $R_{frontal}$ and $R_{profile}$ values 500 simulations were run for 1,000,000 learning iterations. At 200,000 learning iteration intervals, learning was halted and 10,000 iterations were run to gather stable measurements. These results were averaged across all simulations for each break in the learning process. The same method was used for decreasing values of β (habituation rate) but results were taken at 100,00 learning iteration intervals. The β values allowed for a finer temporal resolution than adjusting the $R_{frontal}$ and $R_{profile}$ values because when $R_{frontal}$ and $R_{profile}$ were set to a value of 0.5 the CoD effect broke down at 100,000 iterations.

Results

Socially anxious individuals have a wider CoD than control individuals (Jun et al., 2013; Gamer et al., 2011). One possible explanation for this is that they are avoiding eye contact (Schneier et al., 2011; Schulze et al., 2013) and therefore may have a reduced 'caregiver' reward compared to non-socially anxious individuals. To test this hypothesis, we reduced the rewards associated with the caregiver ($R_{frontal}$, $R_{profile}$) and looked at the effect on the width of the CoD. Simulations were run for longer than previous simulations because the wider CoD for highly anxious individuals should be present during adulthood (Jun et al., 2013; Gamer et al., 2011). Reducing $R_{frontal}$ and $R_{profile}$ did not result in a wider CoD (Figure 7). On the contrary, the simulations suggest that after 600000 iterations, there is a trend towards a narrower CoD when $R_{frontal}$ and $R_{profile}$ were reduced.
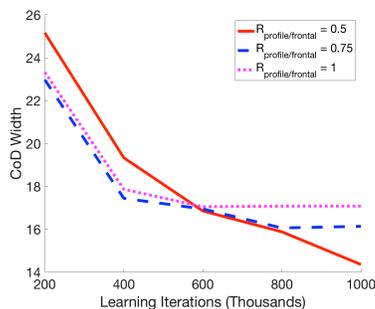

Figure 7. *Change in the width of the CoD effect for different values of $R_{profile}$ and $R_{frontal}$ as a function of the number of learning iterations.*

Another common theory relating to social anxiety is the hyper-vigilance-avoidance hypothesis (Horley et al., 2004) (Wieser et al., 2009). This hypothesis states that socially anxious individuals are hyper-vigilant towards anxiety provoking stimuli and tend to engage in avoidance by looking away. To investigate whether such a hypothesis could account for the wider CoD in socially anxious individuals we increased the value of the habituation parameter β. The goal of this modification was to reduce the

infant's gaze fixation time on the caregiver, as would be expected from avoidance. Increasing the value of β resulted in a progressively narrower CoD effect (Figure 8). This effect was evident after around 200,000 learning iterations.
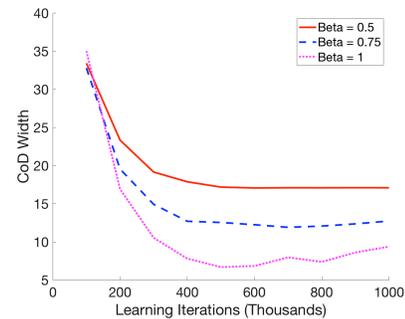

Figure 8. *Change in the width of the CoD effect for different values of β as a function of the number of learning iterations. Larger values of β resulted in a narrower CoD effect.*

Discussion

We have demonstrated that a reinforcement learning account of eye gaze following behavior can be extended to account for the emergence and development of a CoD. In addition, the model also captured the developmental narrowing of the CoD (Mareschal et al., 2016). While a preference for direct gaze may be present from birth for example (Farroni et al., 2002), the fact that the CoD appears to narrow under the influence of reinforcement learning, as seen in developing children, suggests that at least some aspects of the CoD are experience dependent.

The fact that the CoD appears to be influenced by experience and learning poses interesting questions for researchers investigating clinical populations of socially anxious individuals. Importantly it suggests that a critical developmental period may exist that could act as a therapeutic window to reduce the occurrence of behaviors such as social anxiety. We used the model to investigate which aspects of the reinforcement learning framework could influence the developmental trajectory of the CoD. One theory for why socially anxious people may have a wider cone than control individuals is because of their aversion to direct eye contact (Schneier et al., 2011; Schulze et al., 2013). To probe this further, we reduced the rewards associated with the caregiver ($R_{frontal}$ and $R_{profile}$) and looked at the effect on the width of the CoD. Reducing $R_{frontal}$ and $R_{profile}$ resulted in a trend towards a narrower CoD, the opposite of what is seen in highly anxious individuals.

As an alternative to reducing $R_{frontal}$ and $R_{profile}$, we also investigated the effect of increasing the value of the habituation parameter β. This was done in an attempt to capture the hyper-vigilance-avoidance hypothesis, which states that socially anxious individuals are quicker to engage and then avoid anxiety provoking stimuli. An increase in the value of β however, did not lead to a wider CoD. That said, care must be taken when making conclusions from this result. We used mean the infant's fixation time on the

5

caregiver as a functional measure of COD. This measure was used because under a CoD an infant should fixate for longer at small eye gaze deviations because it judges the gaze to be direct. However, this poses a problem when we want to model hyper-vigilance and avoidance by increasing β. An individual exhibiting hyper-vigilance and avoidance will have a reduced mean fixation time on the 'caregiver' even if they perceive the gaze to be directed at them because they would rather shift their gaze to nothing than hold direct gaze. So while increasing the value of β may capture this behaviour, the measure of mean fixation duration will be lower for these cases causing the width of the CoD effect to be smaller even if the CoD is actually wider. Therefore, in order to accurately assess the effect of hyper-vigilance and avoidance on the width of the CoD effect, a different measure that captures when the infant perceives the gaze as being direct is needed.

Our findings (and Triesch et al's) have potential implications for several disorders other than social anxiety. The fact that both eye gaze following and the CoD appear linked by reinforcement learning could provide novel opportunities to investigate disorders that produce both characteristic eye gaze following and CoD behavior. One example of this is autism spectrum disorder. Triesch et al. (2006) put forth multiple candidates under the reinforcement learning framework that could produce the reduced eye gaze following described in individuals with autism. These candidates included a reduced learning rate, reduced caregiver reward and increased shifting latency. The fact that reinforcement learning can account for both eye gaze following and the CoD allows us to explore these candidates further. For example, studies have suggested that individuals with autism have a narrower CoD (Matsuyoshi et al., 2014) and so these candidates should be able to account for that. Indeed, in this study we demonstrated that reducing the caregiver rewards resulted in a slightly narrower cone in the later stages of a simulation. This finding lends weight to reduced caregiver rewards being present in autism. A similar approach should be taken for the other candidates to see if they too can account for both the eye gaze following and CoD differences described in autism, thereby either confirming or rejecting their validity.

Acknowledgments

References

Brooks R. & Meltzoff A.N. (2005) The Development of Gaze Following and its Relation to Language. *Dev. Sci,* 8:535-543.

Deak G.O. (2015) When and Where do Infants Follow Gaze? 5th *International Conference on Development and Learning and on Epigenetic Robotics*, Providence, RI USA.

Farroni T., Csibra G., Simion F. & Johnson M.H. (2002) Eye Contact Detection in Humans From Birth. *PNAS, 99*: 9602-5.

Gamer M. & Hecht H. (2007) Are you Looking at me? Measuring the Cone of Gaze. *JEP: HPP, 33*:705-15.

Gamer M., & Hecht H., Seipp N. & Hiller W. (2011) Who is Looking at me? The Cone of Direct Gaze Widens in Social Phobia. *Cog. & Emotion, 25*:756-64.

Hains S.M.J. & Muir D.W. (1996) Infant Sensitivity to Adult Eye Direction. *Child Development, 67*:1940-51.

Horley K., Williams L.M., Gonsalvez C. & Gordon E. (2004) Face to Face: Visual Scanpath Evidence for Abnormal Processing of Facial Expression in Social Phobia. *Psychiatry Res, 127*:43-53.

Jun Y.Y., Mareschal I., Clifford C.W. & Dadds M.R. (2013) Cone of Direct Gaze as a Marker of Social Anxiety in Males. *Psychiatry Res, 210*:193-8.

Leekam S.R., Baron-Cohen S., Perrett D.I., Milders M. & Brown S. (1997) Eye-Direction Detection: A Dissociation Between Geometric and Joint Attention Skills in Autism. *BJDP, 15*:77-95.

Mareschal I., Otsuka Y., Clifford C.W., Mareschal D. (2016) "Are you Looking at me?" How Children's Gaze Judgments Improve with Age. *Dev. Psy, 52*:695:703.

Matsuyoshi D., Kuraguchi K., Tanaka Y., Uchida S., Ashida H. & Watanabe K. (2014) Individual Differences in Autistic Traits Predict the Perception of Direct Gaze for Males, but not for Females. *Molecular Autism, 5*:12.

Rummery G.A. & Niranjan M. (1994) On-Line Q-Learning Using Connectionist Systems. *CUED/F-INFENG/IR 166*. Cambridge University, UK.

Scaife M. & Bruner J.S. (1975) The Capacity for Joint Visual Attention in the Infant. *Nature, 253*:265-66.

Schneier F.R., Rodebaugh T.L., Blanco C., Lewin H. & Liebowitz M.R. (2011) Fear and Avoidance of Eye Contact in Social Anxiety Disorder. *Compr Psychiatry, 52*:81-7.

Schulze L., Renneberg B. & Lobmaier J.S. (2013) Gaze Perception in Social Anxiety and Social Anxiety Disorder. *Front. in Human Neuroscience, 7*:872.

Schultz W., Dayan P. & Montague P.R. (1997) A Neural Substrate of Prediction and Reward. *Science, 275*:1593-9.

Shepherd S.V. (2010) Following Gaze: Gaze-Following behaviour as a Window into Social Cognition. *Front. in Integrative Neuroscience, 4*:5.

Sutton R.S.& Barto A.G. (1998) *Reinforcement Learning: An Introduction. Cambridg*e, MA: MIT Press.

Thorup E., Nystrom P., Gredeback G., Bolte S., Falck-Ytter T. & the EASE Team (2016) Altered Gaze Following During Live Interaction in Infants at Risk for Autism: an Eye Tracking Study. *Molecular Autism, 7*:12.

Triesch J., Christof T., Deak G.O. & Carlson E. (2006) Gaze Following: Why (not) Learn it? *Dev. Sci, 9*:125-47.

Wieser M.J., Pauli P., Weyers P., Alpers G.W. & Muhlberger A. (2009) Fear of Negative Evaluation and the Hypervigilance-Avoidance Hypothesis: An Eye-Tracking Study. *Journal of Neural Transmission, 116*:717-23.

6