A Reinforcement Learning Model Explains the Stage-wise Development of Gaze Following

http://mesa.ucsd.edu

Hector Jasso[1], Jochen Triesch[2], Christof Teuscher[2]

1: Dept. of Computer Science and Engineering, UC San Diego, USA hjasso@cs.ucsd.edu

2: Dept. of Cognitive Science, UC San Diego, USA triesch@cogsci.ucsd.edu, christof@teuscher.ch

Summary

➢We present a **model of the development of gaze following** consisting of a **reinforcement learning** algorithm [1] maximizing visual reward.

➢The model **captures the progression that human infants go through:** gaze following is first learned for targets positioned in front of the infant, and only later for targets behind.

➢It also **learns to overcome the "Butterworth error"** [2] [3], where younger infants sometimes follow gaze to the correct side of the room but stop at distractors in the visual path.

Problem statement

•**Gaze following** is the ability to redirect one's visual attention towards an object that someone else is looking at (see Figure 1, right).

•**Objective:** To build a **parsimonious model** of the emergence of gaze following in human infants.

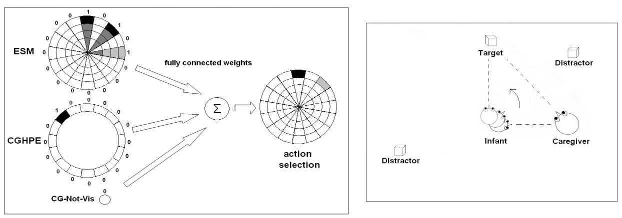

Figure 1. Left: Details of the model (see below). The activations shown correspond to the room configuration on the right figure, when the infant is looking at the caregiver. **Right**: Infant following gaze in the presence of distractors.

A reinforcement learning model of gaze following

•**Model: Linear, gradient-descent Sarsa(_)** reinforcement learning.

•**Input features:** From Expected-Saliency-Map (**ESM**): presence of objects/caregiver at different headings (see Figure 1, left). From Caregiver-Head-Pose-Estimate (**CGHPE**): memory of caregiver's head pose. From Caregiver-Not-Visible (**CG-not-Vis**): whether the caregiver is within the infant's field of view.

•**Actions:** Direct visual attention to a specific location (heading/depth).

•**Rewards:** Visual saliency of the location where attention is being directed to. "Neck cost" subtracted: small for looking straight ahead, large for looking backwards.

Training

•**Room setup**: At the beginning of each training trial, target and distractor objects are positioned **randomly** in the room within a fixed perimeter from the infant (see Figure 2).

•**Parameters:** 500 trials per epoch, 20 time steps per trial. Saliencies: target = 50; distractors = 50; caregiver = 5. _ = 0.001, _ = 0.2, _ = 0.05.

Testing

Experiment 1

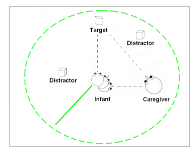

Figure 2. Room setup. Target and distractors positioned randomly in the room, within a fixed perimeter from the infant

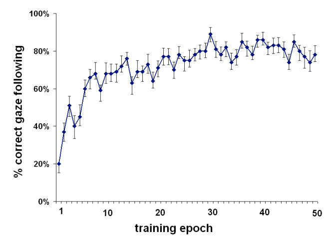

Figure 3. Infant learns to follow gaze to target within the first 50 epochs.

Experiment 2

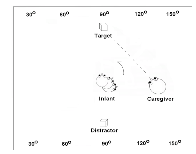

Figure 5. Replication of experiment in [3]. Target along the wall at 30º, 60º, 90º, 120º, and 150º, and a distractor object on the opposite side of the room.

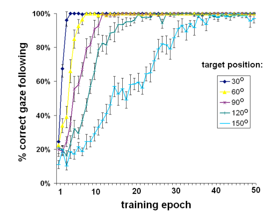

Figure 6. The infant learns to follow the target with a front-to-back progression, as observed in [3].

Experiment 3

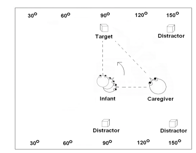

Figure 7. Replication of experiment in [3]. Target along the wall at 90º, 120º, and 150º, second in the line of vision, with a distractor in the same side of the room, at 30º, 60º, and 90º, respectively, and two distractors positioned symmetrically at the other side of the room.

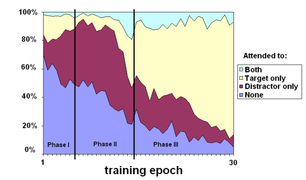

Figure 8. The infant learns to overcome the Butterworth error, as observed in [3]. Phase I: Gaze following is not learned yet. Phase II: Infant stops at distractor object in the gaze path ("Butterworth error). Phase II: Infant follows gaze to correct target.

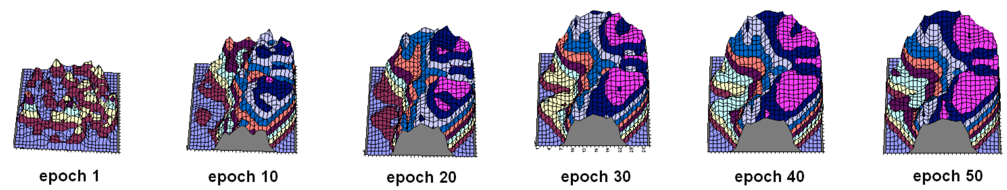

Figure 4. Percentage of trial successes for different target locations (grid alignment corresponds to room setup shown in Figure 2). Gaze following is first learned for objects positioned in front of the infant.

Conclusions and further work

•Gives a **parsimonious account** of gaze following in a spatial setting.
•Based on **reinforcement learning**, which has been proposed by the neuroscience community as a model of learning in the brain.
•Results **match the front-to-back progression** found in infants, as well as the **Butterworth error** and its solution at later stages.
•Currently being adapted to work in a **virtual reality platform** [4] (see Figure 9) to enable more sophisticated experiments.

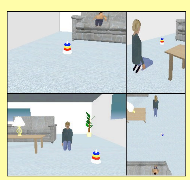

Figure 9. Modeling gaze following in a virtual reality platform.

References

[1] Sutton, R.S., and Barto, A.G. (1998). *Reinforcement Learning: An Introduction.* MIT Press.

[2] Lau, B., and Triesch, J. *Learning Gaze Following in Space: a Computational Model.* 3rd International Conference for Development and Learning, ICDL'04, La Jolla, California, USA, October 20-22, 2004.

[3] Butterworth, G., & Jarrett, N. (1991). *What minds have in common is space: Spatial mechanisms serving joint visual attention in infancy.* British Journal of Developmental Psychology, 9, 55-72.

[4] Jasso, H., and Triesch, J. *A Virtual Reality Platform for Studying Cognitive Development.* 3rd International Conference for Development and Learning, ICDL'04, La Jolla, California, USA, October 20-22, 2004.

[5] Fasel, I., Deak, G. O., Triesch, J. and Movellan, J. R. *Combining embodied models and empirical research for understanding the development of shared attention.* 2nd International Conference on Development and Learning, ICDL'02, Cambridge, MA, USA, June 12-15, 2002.

[6] MESA project, http://mesa.ucsd.edu

Acknowledgements

This work is part of the MESA project at UCSD [5] [6]. We acknowledge funding from the UC Davis MIND Institute and the National Alliance for Autism Research.