

Learning to Look: a Dynamic Neural Fields Architecture for Gaze Shift Generation

Christian Bell, Tobias Storck, and Yulia Sandamirskaya

Institut für Neuroinformatik, Ruhr-Universität Bochum
Universitätsstr. 150, 44780 Bochum, Germany
{yulia.sandamirskaya, christian.bell, tobias.storck}@ini.rub.de

Abstract. Looking is one of the most basic and fundamental goal-directed behaviors. The neural circuitry that generates gaze shifts towards target objects is adaptive and compensates for changes in the sensorimotor plant. Here, we present a neural-dynamic architecture, which enables an embodied agent to direct its gaze towards salient objects in its environment. The sensorimotor mapping, which is needed to accurately plan the gaze shifts, is initially learned and is constantly updated by a gain adaptation mechanism. We implemented the architecture in a simulated robotic agent and demonstrated autonomous map learning and adaptation in an embodied setting.

Keywords: Dynamic Neural Fields, looking, adaptation.

1 Introduction

The ability to direct gaze towards interesting objects in the surrounding environment is one of the most basic goal-directed behaviours of an embodied agent. Bringing the interesting object into the foveal (central) region of the retina not only puts the object’s image into the receptive fields of a larger number of photoreceptors, but also aligns the motor system of the agent with the outside world and allows to calibrate motor plans of other motor actions directed at objects, such as reaching, pointing, or walking towards them [12]. The neural system, responsible for the looking behaviour has been studied experimentally since the beginning of the last century. The neural circuits, involved in generating the goal-directed eye movements, have been identified [4, 11] and include the cerebellum, basal ganglia, superior colliculus, and the frontal eye field.

A prominent property of eye movements is their adaptability. Indeed, the saccadic eye movements¹ are too fast for the visual feedback to influence their accuracy. Thus, gaze shifts have to be planned based on the location of the

¹ Here, we don’t distinguish between saccades, i.e. goal-directed eye movements with restrained head, and saccadic gaze shifts, which include both eye and head movement. See [16] for a discussion of the relation of the neural mechanisms, revealed for saccadic eye movements, and the gaze shifts. In our experiments, the robotic head was moved, which corresponds to a gaze shift.

visual stimulus on the retina. However, the pathway from the retina to the representation of the motor plan and further to muscles is subject to noise and trial-to-trial as well as developmental variations. Consequently, the amplitude of the planned gaze shift towards the object, detected on the retina, needs to be constantly updated. Indeed, the experiments on saccadic adaptation [19, 17, 10] reveal pervasive adaptation capacity of the eye movement circuitry in humans and primates. In these experiments, the saccadic target is shifted during the saccade, when this shift cannot be perceived by the subject. Over several trials, the amplitude of saccades to the given target changes to compensate for the error, perceived after the manipulated saccades. Neural models of this adaptation process include cerebellum and superior colliculus [6, 15, 2, 8, 3], as well as the parietal cortex [9]. Learning mechanisms based on internal feedback and the visual error after a saccade [19] have been proposed to act in this adaptive circuitry. The models of gaze shift generation, mentioned above, focus on identifying neuronal structures involved in saccades generation and adaptation. These models typically do not demonstrate autonomous processing of visual inputs in this system and actual generation and adaptation of eye movements.

Here, we present a model for saccadic gaze shifts, which includes all stages of sensorimotor processing from acquiring visual input from a simulated camera, selecting the target, generating motor command, to actually executing the motor act with simple motor dynamics and updating the sensorimotor gains when a gaze-shift error is detected. This ‘wholistic’ and embodied approach demonstrates how the mechanistic level of a neural architecture may be bridged with the behavioural level of an embodied agent, using the framework of Dynamic Neural Fields.

In this paper, we focus on the learning and adaptation mechanisms, which update the map of gains between the retinal representation of targets and the amplitude of the motor command, which brings the target into the fovea. The actual implementation of the model in a simulated robotic agent revealed nonlinearities in this mapping and dependence of the mapping on both the location of the target on the retina and the motor state of the system prior to the saccade. The gain map is updated based on the error after a saccade; the region of adaptation is selected autonomously based on the perceptual and motor state of the agent before the movement. Our neural-dynamic architecture may be related to the neuronal structures, involved in generation and updating of saccadic eye movements, but the focus of this work is not on neuronal modelling, but on demonstrating the adaptivity of looking behaviour and of the underlying sensorimotor circuits in an artificial system, inspired by the equivalent neuronal system.

2 Methods

2.1 Dynamic Neural Fields: choice of the mathematical framework

Dynamic Neural Fields (DNFs) are continuous in time and in the underlying behavioural space descriptions of activity of neuronal populations [1, 5]. DNFs

are the basis of the Dynamic Field Theory [14], which aims to extend the neural-dynamic modelling approach to account for cognitive behaviour in an embodied and situated agent [13]. DNFs, which form the basis for our architecture, follow the Amari equation, Eq. 1:

$$\tau \dot{u}(x, t) = -u(x, t) + h + S(x, t) + \int f(u(x', t)) w(x - x') dx', \quad (1)$$

where $\dot{u}(x, t)$ is the rate of change of the activation function, $u(x, t)$, defined over a behaviourally relevant space x (e.g., color, location, or motor command). τ is the time constant and $h < 0$ is a negative resting level, which ensures that the DNF is silent (below activation threshold) in an inactivated state. $S(x, t)$ is input to the DNF, which may come from the sensory system or other DNFs. The last term formalizes the lateral neural interactions in the DNF, which are shaped by the interaction kernel, $w(x - x')$. The lateral connectivity is homogeneous within DNFs, with nearby sites exciting each other and far-off sites inhibiting each other. $f(\cdot)$ is a sigmoidal non-linearity, which defines the output of the DNF. This non-linearity and the lateral interactions in the field result in a special form of solution of the DNF equation – a localised activity peak, which is the unit of representation of DFT and a bridge between the continuous in time and in space dynamics of the activation function (which may be directly linked to the sensory input coming from a physical sensor) and discrete, categorical states of the cognitive system. DNFs were previously used to detect, select, and stabilise the representation of target objects, as well as to account for timing of gaze-shift generation in early models for looking behaviour [7, 18, 20].

2.2 The DNF architecture for looking

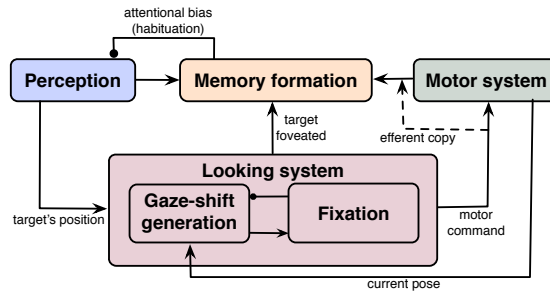


Fig. 1: The DNF architecture for looking. Shaded regions mark four components of the architecture. Arrows show direction of activation flow between different subsystems, lines with circles denote inhibitory couplings.

The complete DNF architecture for looking has four components, shown in Fig. 1. Adaptation of the saccadic gains happens in the gaze-shift generation system. The fixation system tracks the target during fixation and drives the memory formation processes when the target is foveated. The memory formation

system combines the features of the currently attended object with the motor state of the camera head when the object is fixated, creating a body-centred representation of the visual scene.

Fig. 2 shows the gaze-shift generating circuitry, which is the focus of this work. A *perceptual DNF* is defined over the retinotopic space and the visual feature (color, the third dimension, not shown in the figure). Activity peaks over salient objects in the visual field are induced in this DNF by the visual input.

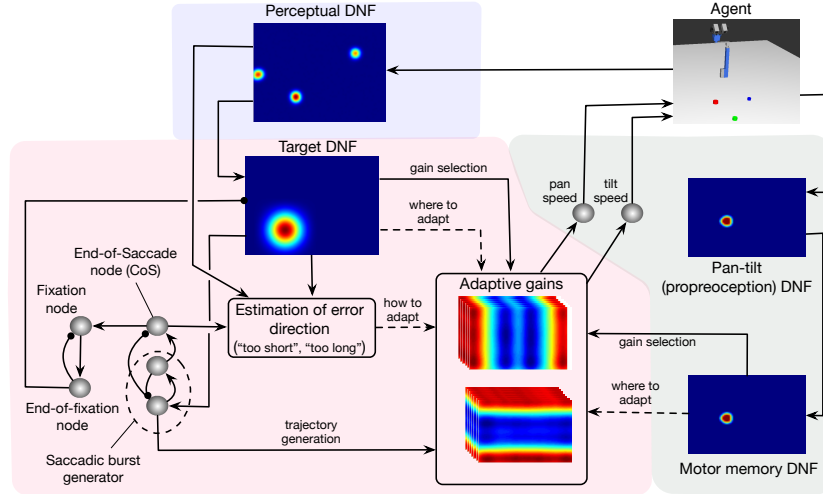


Fig. 2: The gaze-shift generation system. Shaded regions and arrows have the same meaning as in Fig. 1. Dashed lines show couplings involved in adaptation.

The perceptual DNF provides input to the *target DNF*, also defined over the retinotopic space. A single peak evolves in the target DNF over the location of the most salient object. In our architecture, the activity peak in the target DNF is stabilised by very strong lateral interactions, such that this peak is self-sustained. Once initiated, the peak conserves its location even if the initial input from the perceptual system moves (e.g., due to the initiated eye movement). This is a critical property of the architecture: it creates a stabilised, self-sustained representation of the initial location of the target in retinal coordinates, which is critical to enable learning of the gain to the motor system, associated with this retinal location.

An activity peak in the target selection DNF drives a *saccadic burst generator*, which consists of two connected nodes: the excitatory node eventually drives the motor system of the agent and activates the inhibitory node, which, in its turn, inhibits the excitatory node. This system behaves as an oscillator, which generates a stereotypical (i.e. the same for saccades of all directions and amplitudes) trajectory for the motor signal and is similar to neuronal burst generators, involved in saccades generation.

Before arriving in the motor system, the oscillatory signal is scaled with an *adaptive gain*, which depends on the location of the target object on the retina

(represented by the activity peak in the target selection DNF) and the gaze direction of the agent before the saccade (represented in the *motor memory DNF*, driven by the proprioceptive inputs). These dependences are marked with “gain selection” labels in Fig. 2. The latter input is required in our robotic agent, because the mapping from the retinal locations of targets to the respective motor commands (in terms of the amplitude of pan and tilt joints’ movements) is non-linear because of the geometry of the robot. In the human looking system, non-linearities come from the non-linearities of the neuromuscular plant. The motor signal eventually drives the motor system, setting velocities of the two motor joints, which control the camera head in our experiments.

After a gaze shift, an *end-of-saccade node* (condition-of-satisfaction) is activated by the saccadic burst generator, which has finished its single oscillation. The end-of-saccade node activates the fixation system (see Fig. 1), which stabilises the target in the central portion of the visual field and enables memory formation for the gaze-centred representation of the foveated object. A detailed description of the fixation system, as well as memory formation are outside the scope of this paper.

The end-of-saccade node also excites an *error-estimation circuit*, which compares the location relative to the fovea of the activity peaks in the perceptual DNF (after the saccade) and in the target DNF (holding the memory for the target location before the saccade) and represents the decision whether the saccade was too short or too long along the horizontal and the vertical directions in the image plane. This error-sign signal defines whether the gain, specified by the target location and the initial motor state (“where to adapt” labels in Fig. 2), will be increased or decreased (“how to adapt” label in the figure). Next, we describe this adaptation mechanism in more detail.

2.3 Learning and adaptation of the sensorimotor gain maps

Here, we describe the learning process, which initially learns and constantly updates the gain maps if a saccadic error is detected. Initially, the gains are set to small random values and are updated according to the following learning rule:

$$\tau_l \dot{G}^{h,v}(x, y, m, t) = \epsilon^{h,v}(t) f(u_{EoS}(t)) f(u_m(m, t)) f(u(x, y, t)). \quad (2)$$

Here, $G^{h,v}(x, y, m, t)$ are two sets of gain maps (for the horizontal and vertical components of movement). Each of the m gain maps in the two sets is defined over the dimensions of the target DNF, $u(x, y, t)$. The sets span m different initial motor states (the tilt joint angle in our setup). The gains change in the map(s), which are selected by the output of the motor DNF, $f(u_m(m, t))$, at the locations, which are set by the activity peak in the target DNF. $f(u_{EoS}(t))$ is the output of the end-of-saccade node, which is required to be positive (saccade finished) for learning to become active. $\epsilon^{h,v}(t)$ is the error in each of the movement components, τ_l is the learning rate. Using this rule, the whole gain map may be learned in a learning session, where the whole visual-motor space

is sampled. The map is updated locally over a few gaze-shifts if an unexpected change in the sensorimotor plant happens.

3 Results

In our learning experiments we used a simulated robot CoRa, which may be seen in the upper right corner of Fig. 2. The implementation allows to transfer the architecture onto a real robot, but this was not the focus of our work here. In this paper, we report two sets of experiments.

3.1 Learned gain maps

In the first experiment, we investigated how the whole gain-map may be learned by the agent based on a coarse prestructure that associates horizontal movements on the retina with the pan joint of the camera head, and vertical movements with the tilt joint, which is not correct for all initial joint configurations. The objects were placed systematically (on a virtual 16x15 grid in the image, projected on the table) in front of the simulated robot to sample the whole visual space. The initial camera head pose was also varied systematically to sample the whole space of initial motor states (we used only five different initial poses).

Fig. 3 shows the learned gain maps for five slices along the tilt motor dimension. It took 6000 saccades to learn these maps (i.e. approx. five saccades per location and initial tilt). In our experiments, the pan angle of the initial pose of the robot did not affect the gain maps, because of the geometry of the robot. The dependence on the initial tilt angle was successfully learned by the agent and calibrated the sensorimotor mapping, needed to direct actions at the visually perceived objects.





Fig. 3: Gain maps: each map is defined over retinotopic (image-based) coordinates. Maps are arranged according to the initial state (tilt angle) of the camera head. Note that the gain maps change with the initial tilt angle.

3.2 Adaptation experiments

In the second set of experiments, we simulated a gaze adaptation session, in which the target object was displaced during the saccade. This shift is not perceived by the system, since the activity peak in the target DNF is self-sustained

and thus ‘decoupled’ from the perceptual system. Fig. 4 shows results for an experiment with the initial pose fixed at $[0, 0]$ (the robot looks straight ahead). The robot performs horizontal saccades, during which the target is shifted against the saccade direction.

Fig. 4a shows the time-course of the gaze error. The gaze error is estimated only in the periods when the end-of-saccade node is active and is zero otherwise (see the plot). First six saccades demonstrate a low error, since the system has already learned the gain map. Starting with the seventh saccade, the target is shifted (during each gaze shift) against the saccade direction horizontally. Over a few saccades, the error is decreased again; the agent performs shorter saccades to the manipulated target. Fig. 4b shows a slice of the adapted gain map, which corresponds to the initial pose of the robot in the adaptation session. Note the slightly decreased amplitude of the gains around the adapted location ($[25, 55]$ in image coordinates), where the target was perceived.



(a) Gaze shift error. Blue: horizontal, (b) The adapted horizontal gain map. red: vertical.

Fig. 4: Results of an adaptation experiment.

4 Discussion

In this paper, we presented the neural-dynamic architecture for generation of saccadic gaze shifts and their adaptation, which may be coupled to sensory inputs and drive a physical motor system. On the one hand, the architecture offers a neural-dynamic framework, in which various aspects of saccades and gaze shift generation, as well as their adaptation, may be studied. On the other hand, the system allows to study sensorimotor learning in an adaptive artificial cognitive agent.

References

1. S Amari. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27:77–87, 1977.

The project was funded by the DFG SPP ‘Autonomous learning’ within the Priority program 1527.

2. P Dean, J E Mayhew, and P Langdon. Learning and maintaining saccadic accuracy: a model of brainstem-cerebellar interactions. *Journal of cognitive neuroscience*, 6(2):117–38, January 1994.
3. G Gancarz and S Grossberg. A neural model of saccadic eye movement control explains task-specific adaptation. *Vision research*, 39(18):3123–43, September 1999.
4. B Girard and A Berthoz. From brainstem to cortex: computational models of saccade generation circuitry. *Progress in Neurobiology*, 77(4):215–251, 2005.
5. S Grossberg. Nonlinear neural networks: Principles, mechanisms, and architectures. *Neural Networks*, 1:17–61, 1988.
6. J J Hopp and A F Fuchs. The characteristics and neuronal substrate of saccadic eye movement plasticity. *Progress in neurobiology*, 72(1):27–53, January 2004.
7. K Kopecz and G Schöner. Saccadic motor planning by integrating visual information and pre-information on neural dynamic fields. *Biological cybernetics*, 60:49–60, 1995.
8. L M Optican and Ch Quaia. Distributed model of collicular and cerebellar function during saccades. *Annals of the New York Academy of Sciences*, 956(1):164–177, 2002.
9. M Panouillères, O Habchi, P Gerardin, R Salemme, C Urquizar, A Farne, and D Pélisson. A Role for the Parietal Cortex in Sensorimotor Adaptation of Saccades. *Cerebral Cortex*, 2012.
10. D Pélisson, N Alahyane, M Panouillères, and C Tilikete. Sensorimotor adaptation of saccadic eye movements. *Neurosci Biobehav Rev*, 34(8):1103–1120, 2010.
11. C Quaia, P Lefèvre, and L M Optican. Model of the control of saccades by superior colliculus and cerebellum. *Journal of Neurophysiology*, 82(2):999–1018, 1999.
12. J Reuschel, F Rösler, D Y P Henriques, and Katja Fiehler. Spatial updating depends on gaze direction even after loss of vision. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 32(7):2422–9, February 2012.
13. Y Sandamirskaya, Stephan K.U. Zibner, S Schneegans, and G Schöner. Using Dynamic Field Theory to extend the embodiment stance toward higher cognition. *New Ideas in Psychology*, 31(3):322–339, December 2013.
14. G Schöner. Dynamical Systems Approaches to Cognition. *Dynamical Systems*, 2008.
15. Nicolas Schweighofer, Michael A Arbib, and Peter F Dominey. A model of the cerebellum in adaptive control of saccadic gain. *Biological cybernetics*, 75(1):19–28, 1996.
16. D L Sparks. Conceptual issues related to the role of the superior colliculus in the control of gaze. *Current opinion in neurobiology*, 9(6):698–707, 1999.
17. R Srimal, J Diedrichsen, E B Ryklin, and C E Curtis. Obligatory adaptation of saccade gains. *Journal of neurophysiology*, 99(3):1554–8, March 2008.
18. T P Trappenberg, M C Dorris, D P Munoz, and R M Klein. A model of saccade initiation based on the competitive integration of exogenous and endogenous signals in the superior colliculus. *Journal of cognitive neuroscience*, 13(2):256–71, February 2001.
19. J Wallman and A F Fuchs. Saccadic gain modification: visual error drives motor adaptation. *Journal of Neurophysiology*, 80(5):2405–2416, 1998.
20. C Wilmzig, S Schneider, and G Schöner. The time course of saccadic decision making: dynamic field theory. *Neural networks : the official journal of the International Neural Network Society*, 19(8):1059–74, October 2006.