

A possible representation of reward in the learning of saccades

Cornelius Weber and Jochen Triesch

Frankfurt Institute for Advanced Studies, Johann Wolfgang Goethe University
Max-von-Laue Str. 1, 60438 Frankfurt am Main, Germany

Abstract

The saccadic system is adjusted throughout lifetime. A learning scheme that also works in newborns is unlikely to rely on sophisticated object recognition, memory, and geometrical difference computation. The lack of a geometrical model is also appealing to roboticists as it renders unnecessary any hardware-specific parameterization. We review findings on the feedback signal used for saccade learning, and differentiate two possible mechanisms. The first contains a signed feedback, i.e. whether the saccade was too long or too short, and may also be influenced by magnification of the saccade target at the fovea. The second mechanism contains as feedback only the "goodness" of resolution, which is larger near the fovea; hence, there is no vectorial error signal. We demonstrate and compare both mechanisms in a model in which the first mechanism implements horizontal, the second vertical saccade adaptation. This model takes into account the separate neuroanatomical pathways for horizontal and vertical saccade control.

1. Introduction

Saccadic eye movements are among the most successfully studied sensory-motor behaviors. Yet, the saccadic patterns of newborns are poorly characterized. Two week-old infants scan geometric figures rather randomly, while 14 week-olds direct their saccades to stimulus contours more consistently (Bronson, 1990). Young infants often "approach" the new stimulus with a series of saccades rather than just one (for review see Slater, 2001). Saccades are not accurate and precise, even in adults (Figure 1).

Experiments with a systematic target shift during the saccade (intra-saccadic step, ISS paradigm) demonstrate adaptation to a new situation after a couple of trials. Noto and Robinson (2001) controlled corrective saccades and showed adaptations in saccade size with few or no correction saccades. Bahcall and Kowler (2000) found that adaptation occurred with large as well as with small target objects even

though corrective saccades were rarely performed in the case of large objects. Hence, corrective saccades do not seem to be necessary for adaptation. They propose the visual comparison model that generates an error signal from comparing the pre-saccadic image with the predicted post-saccadic image.

Existing models for the learning of saccades apply a positive reinforcement signal when the object of interest is centered on the retina (e.g. Dominey et al., 1995). Hereby they assume that the fact whether the object is centered on the retina is detected and that this information is processed as suitable input to some reward delivering neurons. However, object detection and localization is a strong assumption particularly in the young infant where invariant object recognition might not be fully functional.

The capability to detect whether an object is in the center of the retina is commonly taken for granted, but there is a need for its functional description. It is intuitive that the high resolution resulting from an object being focused in the fovea might "reward" the visual system with a better perception of the object of interest. Such a mechanism would fail if the fovea is missing, such as in age-related macular degeneration (MacAskill et al., 2002). After bilateral foveal lesions, monkeys adapt saccades to a new consistent extrafoveal preferred retinal locus which is mostly in the lower visual field. However, adaptation is much slower than in the ISS paradigm, lasting several weeks and often being incomplete (Heinen and Skavenski, 2002). Here we consider whether a "reward" that takes advantage of the foveal magnification might guide the adaptation of saccades.

We propose two different mechanisms for the feedback for saccade adaptation. For horizontal saccades, the saccade feedback can easily exploit the geometry of visual mapping: each half of the visual field (left or right) is processed by neurons in the contra-lateral hemisphere. In this implementation the saccade gain has to be increased (decreased), if after saccading the target appears in the same (other) hemisphere than before the saccade.

For vertical saccades we propose a scalar reward signal that is proportional to the increased level of neuronal activation due to foveal magnification.




Figure 1: Imprecision of saccades in human adults. The diagonal represents perfect saccades. The shaded (green) cone shaped area shows where saccades actually arrived in experiments. This area is between the mean plus/minus standard deviation. The parameters for the curves are taken from Lewis et al. (2003) who fitted experimental data. The small dark (blue) area is from (Kowler and Blaser, 1995). Their experiment yielded minimal error by a long saccade planning time and a predictable approximate distance.

Hence, saccades that bring an object closer to the fovea are more rewarded than less successful saccades, and no geometric calculations are needed. Note that the first mechanism delivers a signed error value, which in two dimensions may give the vector toward the optimal direction for adjustment. The second, even if applied in both dimensions generally delivers only a positive scalar.

2. Biological Background

Saccades are rarely precise. Figure 1 displays the area of saccade destination (between the limits of the standard deviation) dependent on the saccade eccentricity. Both, mean and standard deviation, increase linearly with eccentricity with a certain offset at zero. Often reported is a tendency for the undershooting of saccades at larger eccentricities.

Sub-cortical and cortical control Several brain circuits are involved in the generation and learning of saccades. Theories mostly assume that there is a shift from predominantly sub-cortical control of saccades in newborns, to cortical control within the first 2-4 months (Csibra et al., 2000). In terms of Bronson’s “two visual systems model” (Bronson, 1974), a phylogenetically older, sub-cortical, “secondary visual system”, which is mature at birth and responds well to stimulus location and movement, is complemented after about 2 months of age by a cortical, “primary visual system” with high visual acuity.

Electrode stimulations to a locus on the superior colliculus (SC) elicit saccades (and matching head movements) of a characteristic amplitude and direction, independent of the current direction of

the eye (see Squire et al., 2003, for an introduction). Hence, the SC plans saccades in retino-centered representations. Neurons in the superficial layers of the SC are primarily visually responsive whereas the deeper layers – the SC is a thick three-dimensional structure rather than a sheet – host the motor cells of the SC. At birth, the retinal projection to the SC has a pattern which resembles that of the adult (Wallace et al., 1997). The SC has a much higher resolution near the fovea than in the visual field periphery, and the size of receptive fields increases with increasing distance from the fovea (Cynader and Berman, 1972).

While younger infants tend to perform saccades according to the targets’ retino-centric locations, in contrast, older infants and adults execute saccades within body-centered spatial coordinates, when effects of intervening eye movements have to be taken into account (Gilmore and Johnson, 1997). Different coordinate frames are found in the posterior parietal cortex (PPC), for example neurons encode locations retinotopically in eye-centered coordinates in the lateral intraparietal area (LIP) and in eye- and also in head-centered coordinates in the ventral intraparietal area (VIP) (Duhamel et al., 1997). Intentional saccades involve the cortical frontal eye field (FEF) (see Canfield and Kirkham, 2001, for review). FEF cells (with retinotopic code) remain active also when an experimental setup causes an intervening saccade. By the age of ~ 3 months, the FEF is involved in the prospective control of saccades, predicting the location of a moving stimulus (Canfield and Kirkham, 2001). LIP neurons perform prediction in that neurons’ receptive fields shift to compensate for the saccade immediately before movement. This also occurs in FEF, in neurons in intermediate SC layers which receive input from LIP (Walker et al., 1995), and to a lesser degree in earlier visual cortical stages (Nakamura and Colby, 2002).

Also, trans-saccadic memory may reside in higher visual areas of the cortex, as it contains abstract rather than image-like representations of the information present in each fixation (Deubel et al., 2002), and is similar to visual short-term memory in which only a handful of memory items survive the saccade.

We conjecture that visual working memory and the match of pre-saccadic expectations and the post-saccadic perception of the target need to be trained from performing saccades. Therefore, a low-level mechanism that does not rely on the functioning of these must enable initial learning of saccades. In the following, we will only address sub-cortical control.

Site of Plasticity Lesion studies demonstrate that the cerebellum is necessary for saccade adaptation, but not for saccade control (see Hopp and Fuchs, 2004, for review). Adaptation fields which

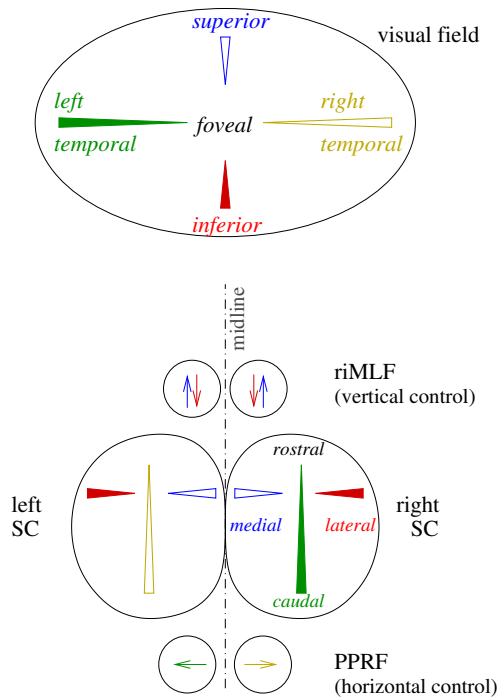


Figure 2: Brain geometries for saccades. Each SC represents the contra-lateral visual field in a topographic fashion as indicated by the color triangular patterns. Specifically, the fovea is represented in each rostral SC. Then, the PPRF control horizontal saccades, each only ipsi-lateral eye movements (indicated by the arrows within), and the riMLF control vertical saccades, where cells encoding up- and downward movements are intermingled (Sparks, 2002). Further areas exist which are not drawn.

are spatially restricted around the adaptation target hint that a structure with a spatial target representation such as the SC, and also the cortical “primary visual system” adapt. A match of the location of SC activation in middle and deeper layers with the visual object but not with the actual saccade endpoint during adaptation (Frens and Van Opstal, 1997) indicates the site of adaptation further downstream toward the motor system. Adaptable efferent connections of the SC might be worth a speculation.

Horizontal and vertical control Figure 2 visualizes how visual field position and saccade direction are represented in the SC. A target seen in the left (right) visual field is represented in the right (left) SC and will activate neurons in the left (right) paramedian pontine reticular formation (PPRF) which in turn will elicit an ipsi-lateral horizontal saccade. Positions along a vertical line in the visual field are represented within an SC along its medio-lateral axis, and the rostral interstitial nucleus of the medial longitudinal fasciculus (riMLF) contains upward and downward saccade-encoding neurons in an intermin-

gled fashion (Sparks, 2002). Further areas and connections exist dealing with issues like intra-saccadic feedback control and the coordination of vertical and horizontal components of oblique saccades.

While saccade vectors are primarily spatially represented on the SC (place code), the PPRF and the MLF encode saccades by the duration and frequency of bursts. This requires a spatio-temporal transformation which is performed – for the horizontal component – by a graded strength of anatomical connections: neurons of the caudal SC have more synaptic boutons in the PPRF than neurons from the rostral SC (Moschovakis et al., 1998).¹ There is no corresponding literature so far for the vertical component.

Hence, different pathways process horizontal and vertical saccades. Processing- and learning principles may also differ, even though there is little direct evidence. For example, Scudder et al. (1998) found that vertical gain changes during adaptation were only between 69 and 89% of the changes in horizontal gain under same conditions. Parkinson patients display primarily vertical saccadic hypometria (Rottach et al., 1996). Further, after horizontal, but not after vertical saccade adaptation, saccade accuracy is decreased (Watanabe et al., 2003) and a facilitating effect on subsequent re-adaptation is found (Kojima et al., 2005).

Possible reward signals Action-relevant feedback can influence even lower visual areas such as primary visual cortex, V1. Wegener et al. (2004) showed that not only the frequency of occurrence of oriented bars shapes V1 cell responses, but that behavioral relevance leads to the learning of more specific orientation responses in V1 cells.

More specifically, the basal ganglia are involved in the generation of voluntary and reward-oriented saccades, and in the suppression of reflex saccades (Hikosaka et al., 2006). Basal ganglia damage during birth predicts major visual abnormalities evident in the following months, such as abnormal fixation shift and visual evoked responses (Mercuri et al., 1997). Direct influences of reward on responses of neurons in lower visual cortex have only recently been discovered (Shuler and Bear, 2006). Chemical reward signals may act directly in visual cortex. Opiate receptors in higher visual cortical areas suggest that large activation may be rewarded (“endomorphin hypothesis”, Vessel et al., 2003). They also report larger activity in higher visual areas for images which subjects preferred over others.

¹Multiple-site stimulation of the SC with synchronous pulses was found to bias saccade vectors toward averaging, whereas asynchronous stimulation biases toward vector summation (Brecht et al., 2004). The latter, paradoxical observation can be explained by a summation-saturation model for the conversion from place codes to rate codes (Groh, 2001), if asynchronous stimulation is regarded as low-amplitude input.

Revising the view that saccade properties are only determined by the site of SC stimulation, Stanford et al. (1996) find that movement amplitude increases with train duration, and movement velocity increases with higher frequency of stimulation. Edelman and Goldberg (2003) find enhanced discharge in the SC for movements directed to high contrast stimuli. This motivates the idea that saccades aim at raising the post-saccadic activation level, e.g. in the SC. For vertical saccade adaptation, we will describe a direct implementation of such a mechanism.

Such a mechanism, however, cannot explain the following variant of the ISS paradigm, which was performed with horizontal saccades. Robinson et al. (2003) waited until a saccade was almost complete before displacing the target to a position with a constant offset from the saccade endpoint. Saccades adapted into the direction of the target offset, even though by experimental design this gain change produced no signal of learning success. As this experiment was performed only with horizontal saccades and offsets, and as the vertical saccade circuit is distinguished and separate from the horizontal circuit, we may propose the scalar feedback mechanism for the vertical (components of) saccades.

Existing Models A review of models of saccade generation is given in Girard and Berthoz (2005).

For saccade adaptation, many models use an error vector in 2-dimensional visual space between the target object and the actual location of the fovea after performing the saccade (e.g. Weber et al., 2005), possibly expressed as a vectorial gaze motor error considering the difference between the desired and actual length of the oculomotor muscles (e.g. Gancarz and Grossberg, 1999). To explain how the distance between target object and fovea can be measured, a model can make the assumption that the object persists in short term memory, for example by having been focused previously (Rao and Ballard, 1994).

Weaker assumptions are made in a model for gaze-following (Jasso et al., 2006) where a binary reward signal is given when the target object is in (or near) the fovea. Correct and incorrect trials (determined as “left” or “right”) are distinguished also in a conditional visual discrimination model (Dominey et al., 1995). It associates activation on cue-related IT cells with activation on striatal cells which initiate a saccade to the corresponding target for the current cue. Initially random saccades, when successful, strengthen the correct associations. Still, these models do not explain the origin and implementation of the feedback signal.

The model of Bandera et al. (1996) applies reinforcement learning to recognize objects with a minimal number of saccades. The model “interrogates” the fixation points by its high-resolution fovea; rather

than aiming at increasing neural activations, saccades are learned to target those locations where the most informative features are expected, dependent on the model database and the already seen features.

On the “BabyBot” robot platform that uses log-polar vision, Orabona et al. (2005) propose an object-based model of visual attention. The center of mass of the most salient blob is selected for a saccade elicited to accumulate information about an object. Saccading in the “BabyBot” is done using a closed-loop control system which is adapted based on a vectorial retinal position error (Metta et al., 2000).

A reverse technique of technical interest for robotic saccade adaptation is implemented in the algorithm of Rodemann et al. (2004). It identifies the patch in the image center after a saccade and searches for the position of the patch in the pre-saccadic image so to learn this position’s relation to a motor command.

Itti and Baldi (2005) predict human saccades when watching complex video clips. They define a Bayesian “surprise” as the Kullback-Leibler divergence between a concurrent Prior and the Posterior from the data, and predict saccades as maximizing surprise. Prediction significantly benefits from modeling increased resolution at the fovea (Itti, 2005). Similar information theoretic measures have been used for sensory-motor mappings with action perception loop. Intelligent adaptive curiosity, for example, pushes a robot toward situations in which it maximizes its learning progress (Oudeyer et al., 2005). Similar to our proposed subcortical mechanism for vertical saccades, these principles may tend to maximize activation levels in certain brain areas.

3. Model and Results

Our model addresses what the feedback signal for saccade learning consists of. While we can assume a functional, cortical working memory in models of older infant saccade learning, there are restrictions for newborns which are likely to rely on sub-cortical mechanisms. For saccade adaptation, a geometrical vectorial error signal between a saccade endpoint and the actual target is also a strong assumption. When not knowing about one’s own fovea, one can only ask, “What kind of eye movements benefit the view?”.

A consideration motivating our proposed mechanism for vertical saccade adaptation is the following. Let us assume that a salient object (in terms of e.g. contrast or motion signals which are accessible to an infant) activates some peripheral retinal cells, and hence also some peripheral cells downstream where saccades are elicited, such as the SC deeper layers. Further, random saccades may be generated, occasionally foveating (or near-foveating) the object by chance. We consider the simple idea that an increased resolution of that salient object being perceived in the fovea can generate a learning signal (re-

ward) by activating an increased number of neurons in, e.g., the SC. This is used to train the peripheral neurons in the motor layers of the SC that have been activated by the object prior to a vertical saccade.

Log-Polar Vision A log-polar mapping between the retina and lower visual cortical areas (Schwartz, 1977) is increasingly being used for artificial vision systems (e.g. Barnes and Sandini, 2000). Accordingly, a point described on the retina in polar coordinates (r, θ) is mapped to SC coordinates $(\xi = q_{rc} \log(sr + \epsilon), \eta = q_{ml}\theta)$. Parameters are ϵ , s , q_{rc} and q_{ml} . We set $\epsilon > 0$ to avoid the singularity of the logarithm at zero (another possible remedy would be to set $\xi = \text{const} \cdot r$ near zero and shift the logarithm to fit). Parameters were fitted to experiments: First, we set $\epsilon = 1$ to yield only positive values. Next, we set $s = 0.25[\frac{1}{deg}]$, so that the region within 10° around the fovea will occupy 40% of the SC (according to Figure 3 in Wallace et al., 1997). Finally, we scaled the curve to cover a SC of rostro-caudal extension of 5mm (see <http://brainmaps.org/>) by setting $q_{rc} = 1.58[mm]$. A hemifield of 180° covers 5mm in the medio-lateral direction of the SC by setting $q_{ml} = \frac{5}{180}[\frac{mm}{deg}]$.

Theoretical Considerations We are considering a system which adjusts the precision of its saccades by measuring activation increase in the SC. In order to do so, this difference must be reliable and sensitive to deviations from a perfect saccade. Figure 3 shows that the area that a disk seen in the visual field will occupy on the SC rises significantly when the disk is centered on the fovea. A learning system has to deal with the following quantitative trends: (i) the differences in SC occupation area are extreme for a given disk when seen at e.g. 0° and 30° eccentricity, while (ii) the differences in SC occupation area depend lesser on disk sizes; and (iii) at small deviations of the target from the center of the fovea, a decrease in SC occupation area is strong for small disks, but for large disks the decrease becomes strong only at a larger deviations.

Model Architecture The model architecture is shown in Fig. 4. Stimuli on the retina project to the SC by a log-polar mapping without modeling spatially extended receptive fields. The representation of visual stimuli is expanded if seen near the fovea.

Then, horizontal and vertical control circuits are distinguished. Trained weights link SC to one horizontal burst generator (BG) neuron for each, left and right, direction.

For the vertical direction, there is a motor neuron for every vertical saccade distance (which may reside in the deeper layers of the SC). This is necessary, because we “reward” such actions that lead to good sac-




Figure 3: The area which a disk shown in the visual field occupies on the SC dependent on the distance (in degrees) to the fovea. Assumed is the log-polar mapping between the visual field and the SC. Shown for three disk sizes, of radius 10° , 5° and 2° . The ordinate denotes fraction of SC covered; for example, a disk of 10° radius shown in the center of the retina projects to an area covering 40% of the SC.

cades by strengthening the weights that led to this action. If there was only one rate- (burst-) coding neuron, then slightly overshooting saccades which might still enhance resolution would lead to even stronger connections and larger overshoots.² Therefore, we have implemented the vertical SC motor neurons carrying a spatial code, and a transformation to a temporal code at the vertical BG neuron(s) would need to be done in another, non-adaptive step.

Learning Horizontal Saccades The adaptation of the horizontal components of saccades is driven by the laterality of post-saccadic SC activation. When a blob of activation is shown on the retina representing the target, the log-polar mapping produces a pre-saccadic activation a_j^{SC} of an SC unit j . Then $m_c^h = \sum_j w_j^h a_j^{SC}$ is the activation of the contra-lateral horizontal saccadic BG, with the weights w_j^h crossing the midline. After the saccade, the target will either be represented primarily on the same SC (too small saccade) or on the contra-lateral SC (too large saccade). A useful measure to obtain the sign is to take the difference between average ipsilateral and contra-lateral post-saccadic SC activation: $D = \bar{a}^i - \bar{a}^c$. This term will be negative for overshooting saccades, and can be readily used as a modulator in an otherwise Hebbian-like weight learning rule: $\Delta w_j^h \approx m_c^h a_j^{sc} D$.

Since the difference D can be large near the fovea but is tiny at the periphery, this can in principle explain the foveal lesion experiment of Heinen and Skavenski (2002) in which adaptation took weeks, rather than minutes as with intact foveae.

²An architecture with just one neuron may be realizable with some mechanism exploiting the tendency for undershoot of saccades.

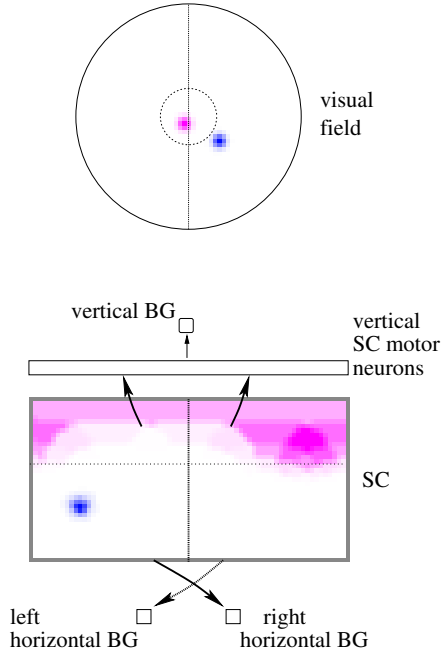


Figure 4: Model architecture. There is a log-polar mapping between the retina (the visual field is shown) and the superior colliculus (SC). Two colored Gaussian patches on the retina and their projections on the SC are shown, highlighting the large magnification of the foveal patch (pink). Learned connections are shown by thick arrows; below the SC are horizontal burst generator (BG) neurons; above is the circuit for the vertical saccade control.

Learning Vertical Saccades The learning rule which we propose for the vertical components of saccades modulates learning in proportion to the foveal magnification increase of the seen target object. From the pre-saccadic SC unit activations, first the vertical motor neurons’ inner activations a_i^v (membrane potential) are computed as: $a_i^v = \sum_j w_{ij}^v a_j^{SC}$. The most active unit (winner) is determined, but in order to allow for random exploration, a unit in the spatial vicinity is determined according to the probability of a Gaussian ($\sigma = 2$) centered around the winner. Another Gaussian ($\sigma = 4$) is set around this unit to set the activations $m_i^v(\bar{a}^v)$, and the vertical saccade displacement encoded in this population is performed: each motor unit i has a vertical saccadic displacement m_i^v assigned, the length m^v of a saccade is based on the pooled contribution from the motor neurons. The difference between post- and pre-saccadic total SC activation $T = \bar{a}^{post} - \bar{a}^{pre}$ is measured and modulates the Hebbian learning: $\Delta w_{ij}^v \approx m_i^v a_j^{SC} T$. Outgoing weight normalization $\sum_i (w_{ij}^v)^2 = 1$ for SC neurons j enforces competition.

Simulation and Parameters The modeled circular retina was cut out from a square of 101×101 units. The SC has 51 units along the rostro-caudal,

101 units along the medio-lateral direction, which crosses the midline. It has modifiable connections to two horizontal BG neurons and to an array of 101 vertical SC motor units. The latter code for retinal vertical saccade displacements of -54 to 54 pixels and have (virtual) fixed connections, used for pooling, to a vertical BG neuron to execute that command. We have used a 180° wide visual field, even though so large saccades cannot be done. Each of the 101 units therefore covers a field of about 1.8° ; similar is the precision of the vertical SC motor units.

The visual stimuli are Gaussian activation patches ($\sigma=2$ in pixel units) shown to the retina (visual field in Fig. 4). Such activations might be a result of a saliency computation, which is not a topic of this paper. One patch is presented before each saccade and again after a saccade at the corresponding displacement. The position of each pre-saccadic patch is random with equal probability for every position on the retina. Corrective saccades are not done; after each saccade, a new random patch is created. All Gaussians (σ ’s) were roughly five-fold wider at the beginning and narrowing to their final values during the first 5% of training steps in order to aid learning.

Model Results Trained weights from SC to the horizontal BG forming a rostro-caudally increasing gradient (Moschovakis et al., 1998) are shown in Fig. 5 a). Sample weights to three of the vertical SC motor units are shown in Fig. 5 b). Each such unit receives input from a region in SC for which the same vertical saccade displacement has to be made in order to foveate the visual target. Figure 5 c) shows an error map for retinal target locations. 150,000 training steps have been made and 250,000 error averaging steps.

4. Discussion

We have shown two alternatives of visual feedback for saccade adjustment, and implemented both in a combined model. The first mechanism detects whether a saccade is too large or too short, and retrieves a signed error signal. It extends current models by proposing that the activation level difference between contralateral and ipsilateral SC modulates learning. The second mechanism cannot detect the sign of the error, but only how far the saccade is off target.

Assumption on Visual Stimuli The stimuli must activate a localized peripheral location before the saccade, and must significantly activate the fovea at a higher resolution after the saccade. Helpful for such a salient response might be spatial summation and surround inhibition of SC neurons (Wallace et al., 1997), or newborns’ and adults’ preference to moving stimuli. This could also alleviate




Figure 5: Learned weights and the resulting saccade errors. **a)** shows the weights of the left- and right horizontal BG neurons. Each neuron has strong weights in the caudal (bottom) contra-lateral SC. **b)** shows weights of three of the 101 vertical SC motor neurons. Middle: unit which elicits zero vertical displacement. Left and right, units eliciting vertical saccade displacements of approximately -27° and 27° , respectively. **c)**, **d)** and **e)** show horizontal, vertical and absolute saccadic error, respectively, for all visual target points in the visual field. Deviations of saccade endpoints from the center are coded dark. Scale bars denote errors in pixel units (1 pixel $\approx 1.8^\circ$).

the interference of multiple stimuli. It is also conceivable that the object at the periphery pre-cues some features (e.g. a certain color) and that such features will elicit a greater response at the fovea. Such pre- and post-saccadic image comparisons are implemented in some models (e.g. Rao and Ballard, 1994), but may rather be cortical.

Suggestions for Research Since the second proposed adaptation mechanism is in conflict with the experiment of Robinson et al. (2003), it may be implemented in the brain only for vertical (components of) saccades. It might be falsified when performing the Robinson experiment also for vertical saccades.

The prime alternative would be a mechanism similar to that for horizontal adaptation. Central to horizontal gain control is the rostro-caudal gradient of SC-to-PPRF connections (Moschovakis et al., 1998). To demonstrate similarity for vertical gain control, one has to find the following gradients in the SC-to-rIMLF connections: from center SC to *medial* SC (open blue triangle in Fig. 2), connections would increase for *upward* rIMLF saccade burst generator (BG) cells, while from center SC to *lateral* SC, connections would increase to *downward* BG cells.

Acknowledgments

We thank anonymous reviewers for valuable feedback. JT and CW are financed by the Hertie Foundation.

References

- Bahcall, D. and Kowler, E. (2000). The control of saccadic adaptation: implications for the scanning of natural visual scenes. *Vision Res.*, 40:277996.
- Bandera, C., Vico, F., Bravo, J., Harmon, M., and Baird, L. (1996). Residual q-learning applied to visual attention. In *International Conference on Machine Learning*, pages 20–27.
- Barnes, N. and Sandini, G. (2000). Direction control for an active docking behaviour based on the rotational component of log-polar optic flow. In *Proc. Europ. Conf. on Computer Vision, Vol. 2*, pages 167–81.
- Brecht, M., Singer, W., and Engel, A. (2004). Amplitude and direction of saccadic eye movements depend on the synchronicity of collicular population activity. *J Neurophysiol*, 92:424–32.
- Bronson, G. (1974). The postnatal growth of visual capacity. *Child Development*, 45:87390.
- Bronson, G. (1990). Changes in infants' visual scanning across the 2- to 14-week age period. *J Exp Child Psychol.*, 49(1):101–25.
- Canfield, R. and Kirkham, N. (2001). Infant cortical development and the prospective control of saccadic eye movements. *Infancy*, 2(2):197–211.
- Csibra, G., Tucker, L., Volein, A., and Johnson, M. (2000). Cortical development and saccade planning: the ontogeny of the spike potential. *Neuroreport*, 11(5):1069–73.
- Cynader, M. and Berman, N. (1972). Receptive-field organization of monkey superior colliculus. *J Neurophysiol.*, 35(2):187–201.
- Deubel, H., Schneider, W., and Bridgeman, B. (2002). Transsaccadic memory of position and form. *Progress in Brain Research*, 140:165–80.
- Dominey, P., Arbib, M., and Joseph, J. (1995). A model of corticostriatal plasticity learning oculomotor associations and sequences. *J Cog Neurosci*, 7(3):311–36.
- Duhamel, J., Bremmer, F., Benhamed, S., and Graf, W. (1997). Spatial invariance of visual receptive fields in parietal cortex neurons. *Nature*, 389:845–8.
- Edelman, J. and Goldberg, M. (2003). Saccade-related activity in the primate superior colliculus depends on the presence of local landmarks at the saccade endpoint. *J Neurophysiol*, 90:1728–36.
- Frens, M. and Van Opstal, A. (1997). Monkey superior colliculus activity during short-term saccadic adaptation. *Brain Research Bulletin*, 43(5):473–84.
- Gancarz, G. and Grossberg, G. (1999). A neural model of the saccadic eye movement control explains task-specific adaptation. *Vis Res*, 39:3123–43.
- Gilmore, R. and Johnson, M. (1997). Body-centered representations for visually-guided action emerge during early infancy. *Cognition*, 65(1):B1–9.
- Girard, B. and Berthoz, A. (2005). From brainstem to cortex: computational models of saccade generation circuitry. *Prog Neurobiol.*, 77(4):215–51.
- Groh, J. (2001). Converting neural signals from place codes to rate codes. *Biol. Cybern.*, 85:159–65.

- Heinen, S. and Skavenski, A. (1992). Adaptation of saccades and fixation to bilateral foveal lesions in adult monkey. *Vision Res*, 32(2):365–73.
- Hikosaka, O., Nakamura, K., and Nakahara, H. (2006). Basal ganglia orient eyes to reward. *J Neurophysiol*, 95(2):567–84.
- Hopp, J. and Fuchs, A. (2004). The characteristics and neuronal substrate of saccadic eye movement plasticity. *Progress in Neurobiology*, 72(1):27–53.
- Itti, L. (2005). Quantitative modeling of perceptual salience at human eye position. *Visual Cognition*.
- Itti, L. and Baldi, P. (2005). A principled approach to detecting surprising events in video. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 631–7.
- Jasso, H., Triesch, J., Teuscher, C., and Deák, G. (2006). A reinforcement learning model explains the development of gaze following. In *International Conference on Cognitive Modeling (ICCM)*.
- Kojima, Y., Iwamoto, Y., and Yoshida, K. (2005). Effect of saccadic amplitude adaptation on subsequent adaptation of saccades in different directions. *Neurosci Res*, 53(4):404–12.
- Kowler, E. and Blaser, E. (1995). The accuracy and precision of saccades to small and large targets. *Vision Res.*, 35(12):1741–54.
- Lewis, A., Garcia, R., and Zhaoping, L. (2003). The distribution of visual objects on the retina: connecting eye movements and cone distributions. *J Vision*, 3(11):893–905.
- MacAskill, M., Anderson, T., and Jones, R. (2002). Saccadic adaptation in neurological disorders. *Progress in Brain Research*, 140:419–33.
- Mercuri, E., Atkinson, J., Braddick, O., Anker, S., Cowan, F., Rutherford, M., Pennock, J., and Dubowitz, L. (1997). Basal ganglia damage and impaired visual function in the newborn infant. *Archives of Disease in Childhood*, 77:F1114.
- Metta, G., Panerai, F., Manzotti, R., and Sandini, G. (2000). Babybot: an artificial developing robotic agent. In *SAB*.
- Moschovakis, A., Kitama, T., Dalezios, Y., Petit, J., Brandi, A., and Grantyn, A. (1998). An anatomical substrate for the spatiotemporal transformation. *J Neurosci*, 18(23):10219–29.
- Nakamura, K. and Colby, C. (2002). Updating of the visual representation in monkey striate and extrastriate cortex during saccades. *PNAS*, 99(6):402631.
- Noto, C. and Robinson, F. (2001). Visual error is the stimulus for saccade gain adaptation. *Cog Brain Res*, 12:301–5.
- Orabona, F., Metta, G., and Sandini, G. (2005). Object-based visual attention: a model for a behaving robot. In *3rd International Workshop on Attention and Performance in Computational Vision*.
- Oudeyer, P., F., K., Hafner, V., and Whyte, A. (2005). The playground experiment: Task-independent development of a curious robot. In *AAAI Spring Symposium Workshop on Developmental Robotics*.
- Rao, R. and Ballard, D. (1994). Learning saccadic eye movements using multiscale spatial filters. In Tesauro, G., Touretzky, D., and Leen, T., (Eds.), *Proceedings of NIPS 7*, pages 893–900.
- Robinson, F., Noto, C., and Bevans, S. (2003). Effect of visual error size on saccade adaptation in monkey. *J Neurophysiol*, 90:1235–44.
- Rodemann, T., Joubin, F., and Körner, E. (2004). Saccade adaptation on a 2 dof camera head. In *Self-Organization and Adaptive Behaviour*, pages 94–103.
- Rottach, K., Riley, D., DiScenna, A., Zivotofsky, A., and Leigh, R. (1996). Dynamic properties of horizontal and vertical eye movements in parkinsonian syndromes. *Annals of Neurology*, 39(3):368–377.
- Schwartz, E. (1977). Spatial mapping in the primate sensory projection: Analytic structure and relevance to perception. *Biol Cybern*, 25(4):181–94.
- Scudder, C., Batourina, E., and Tunder, G. (1998). Comparison of two methods of producing adaptation of saccade size and implications for the site of plasticity. *J Neurophysiol*, 79(2):704–15.
- Shuler, M. and Bear, M. (2006). Reward timing in the primary visual cortex. *Science*, 311:1606–9.
- Slater, A. (2001). *Blackwell Handbook of Infant Development*, chapter Visual Perception, pages 5–34.
- Sparks, D. (2002). The brainstem control of saccadic eye movements. *Nat Rev Neurosci*, 3:952–64.
- Squire, L., Roberts, J., Spitzer, N., Zigmond, M., S.K., M., and Bloom, F., (Eds.) (2003). *Fundamental Neuroscience*. Academic Press, 2 edition.
- Stanford, T., Freedman, E., and Sparks, D. (1996). Site and parameters of microstimulation: evidence for independent effects on the properties of saccades evoked from the primate superior colliculus. *J Neurophysiol*, 76(5):3360–81.
- Vessel, E. A., Biederman, I., and Cohen, M. S. (2003). How opiate activity may determine spontaneous visual selection. *J Vision*, 3(9):6a.
- Walker, M., Fitzgibbon, E., and Goldberg, M. (1995). Neurons in the monkey superior colliculus predict the visual result of impending saccadic eye movements. *J Neurophysiol.*, 73(5):1988–2003.
- Wallace, M., McHaffie, J., and Stein, B. (1997). Visual response properties and visuotopic representation in the newborn monkey superior colliculus. *J Neurophysiol*, 78:2732–41.
- Watanabe, S., Ogino, S., Nakamura, T., and Koizuka, I. (2003). Saccadic adaptation in the horizontal and vertical directions in normal subjects. *Auris Nasus Larynx*, 30:S41–5.
- Weber, C., Karantzis, K., and Wermter, S. (2005). Grasping with flexible viewing-direction with a learned coordinate transformation network. In *Proc. Humanoids*.
- Wegener, D., Freiwald, W., and Kreiter, A. (2004). The influence of sustained selective attention on stimulus selectivity in macaque visual area MT. *J Neurosci*, 24(27):6106–14.