PALADYN Journal of Behavioral Robotics

VERSITA

Research Article · DOI: 10.2478/s13230-013-0103-y JBR · 3(3) · 2012 · 147-155

An evaluation of environmental constraints for biologically constrained development of gaze control on an iCub robot

Patricia Shaw[1]*, **James Law**[1]†, **Mark Lee**[1]‡

1 Department of Computer Science, Aberystwyth University, Aberystwyth, Wales, UK

Received 14-12-2012

Accepted 27-03-2013

Abstract

Gaze control requires the coordination of movements of both eyes and head to fixate on a target. Using our biologically constrained architecture for gaze control we show how the relationships between the coupled sensorimotor systems can be learnt autonomously from scratch, allowing for adaptation as the system grows or changes. Infant studies suggest developmental learning strategies, which can be applied to sensorimotor learning in humanoid robots. We examine environmental constraints for the learning of eye and head coupled mappings, and give results from implementations on an iCub robot. The results show the impact of these constraints and how they can be overcome to benefit the development of fast, cumulative, on-line learning of coupled sensorimotor systems.

Keywords

developmental robotics · gaze control · sensorimotor learning · eye-head coordination · humanoid robotics

1. Introduction and Background

Developmental robotics is a field of research that focuses on ontogeny as the inspiration and primary concept for building and understanding cognitive learning systems [1]. The fundamental assumption is that autonomous cognitive robots are unlikely to be created by designing complete advanced systems; rather we must find out how to build agents that are initially *less* competent but nevertheless have the key ability to learn and grow cognitively through their own experience.

We are inspired by the enormous cognitive growth and development manifest in the human infant during the first year of life. If some of the mechanisms for sensorimotor learning, object and causality detection, imitation, etc. can be modelled from infant behaviour then it may be possible to implement these in robots that learn through experience.

In infancy, humans develop through a series of behavioural stages. These stages are well recorded in developmental psychology, and show the cycle of learning and generalisation of competencies that will support the infant during its lifetime. Behaviours rapidly emerge, consolidate, are superseded, or fused together creating new and improved competencies, during a period of intense activity and change. Although stages and their timings vary between individuals, it is widely recognised that learning progresses through a sequence influenced both by internal and external factors.

One of the most influential figures in the study of staged growth has been Jean Piaget, who placed great emphasis on the importance of early sensorimotor interaction [2]. We believe that sensorimotor interaction is also key to learning in robotics, and that algorithms for robotic learning should be grounded in the sensorimotor period. Not only is it logical to start learning at the earliest stage, as early experiences are likely to affect later learning, but a robot's 'understanding' of the world should be based on its sensor and motor experiences. The sensorimotor stages identified by Piaget are not only relevant to robotics, but it seems possible that sensorimotor coordination is a significant general principle of cognition [3]. This view that grounding and early start points are crucial for the growth of adaptive intelligence is very well argued by Smith and Gasser [4] in their "Six Lessons from Babies".

Hence, in our work we are investigating stages of development as a driver for robotic learning, with a focus on the sensorimotor stage. We take inspiration from human infant development, and the emerging control of the body over the first months of life. Although the infant may seem slow to gain control of its faculties, it is in fact developing at a rapid rate, and we believe that the identifiable stages are the manifestation of mechanisms that are key to this process. From spontaneous, uncoordinated, apparently random movements of the limbs the infant gradually gains control of the available parameters, and learns to coordinate sensory and motor signals to produce purposive acts in egocentric space [5, 6].

2. Constraints on Learning

2.1. Infant development

At birth the infant has limited torque in the neck muscles. With the effect of gravity, movements of the head are very limited. The eyes however have low gravity and inertial effects, with a high power to weight ratio. This allows the eye motor system to be very active, although vision is poor, so the neonate initially make saccades in the direction of auditory stimuli [7].

The newborn makes relatively few saccades, either to stimuli in the periphery of the visual field, or in response to acoustic stimuli. Although they can be distracted by visual stimuli, they tend to fixate on a single target. When they do saccade, they are slow to trigger, and often several saccades are required to fixate on a target [8, 9]. Saccades are

---

[*] E-mail: phs@aber.ac.uk
[†] E-mail: jxl@aber.ac.uk
[‡] E-mail: mhl@aber.ac.uk

also more often horizontal than vertical [8], and made to near objects rather than distant ones.
The head starts to make some unsupported movements during the first month. However, the muscle strength is only sufficient to hold the head for a few seconds at a time [10, 11].
At this age infants tend to fixate on object edges, rather than internal features [12], and are incapable of remembering anything about them [13]. Faster and more accurate saccades appear after 7 weeks [14]. The frequency of saccades increases, and fixations are focused within objects, rather than on their edges [12]. This coincides with an increase in visual acuity during the second month [15], and an increase in the field of view from around 20 degrees at 6 weeks to 40 degrees at 10 weeks [16]. The infant can follow moving stimuli to mid-line [17], although the infant still shows very little head movement during gaze shifts of up to 30 degrees amplitude [18]. The accuracy of saccades continues to improve and, correspondingly, the number of saccades required to fixate reduces, with near adult performance observed at 7 months [9].
By the third month, the head muscles are gaining strength and learning to contribute movements to gaze shifts. The head consistently contributes to gaze shifts greater than 30 degrees, whilst only contributing 25% of the time to 10 degree gaze shifts [18].
By the 4th month, the infant can follow objects smoothly with their eyes, and rotate their necks freely [17]. Gaze control continues to improve, until the 6th month, by which time the infant is visually insatiable, moving head and eyes to search for, and fixate on, novel stimuli [11].
This gradual development of neural structures reduces the initial complexity of learning in a noisy and dynamic environment. In effect, these limitations 'constrain' and help to focus the infants learning into manageable stages. As the infant develops, both in muscle tone and visual acuity, the gaze control is refined, finally learning how to coordinate eye and head movements to fixate on and track objects in its environment.
An extensive review of early infant development covering both the psychological and neurological literature can be found in [19].

2.2. Application to robotics

When creating a system of staged development on a robotic platform we are interested in both the manifest improvements in sensory abilities and motor control, and the underlying neurological changes that support these advancements. The infant development literature serves both as a foundation for developing robot behaviour, but also as a benchmark for evaluation. By taking account of the available modalities and subsystems of a given robot, it is possible to map such prototype infant data onto a developmental sequence for the robot. We have performed such a mapping using the iCub humanoid robot [20] and produced a comprehensive chart of the general developmental possibilities for the sensorimotor systems of the iCub [21].
Sensorimotor learning is conducted based on our mapping framework, and utilises the modulating influence of a dynamic constraint network to shape the developmental sequence following our approach towards constraint based learning [22]. There are various kinds and sources of constraints, but there are two main types and here we consider these as two alternate implementation strategies for robotic systems.
The first, type A, is derived from the limitations of immature neurological and physiological structures. Such limitations include the poor muscle tone that prevents the newborn from lifting its head at birth, and the lack of acuity and depth of the visual field. These are hard constraints that limit the ability of the infant and are only removed from the relevant systems when sufficient physical growth or maturity has been achieved. Although normal infants develop at different rates, their trajectories are similar, and tend to follow a common timeline, with certain stages appearing before others. Such constraints can be imposed on the system using a sequence table, such as those in [21], which are extracted from the infant data. Constraints can be released when suitable levels of competence have been achieved and we use thresholds on internal state variables to trigger their removal in a semi-structured manner [22]. Hence, a robot built using Type A constraints can be expected to follow the general infant timeline, with variances reflecting its own particular circumstances.

The second type of constraints, type B, reflect external effects that restrict or enhance development in more complex ways. Such effects may include interaction with carers, the level of stimuli in the environment, and the number of opportunities to practice. There are many experiments that have shown how the order of training on different sequences of experience can affect learning rates and the acquisition of competencies, for example the experiments by Needham et al. [23] use a 'sticky mitten' to ease the constraint on grasping thereby enabling greater interaction with objects.

Such constraints cannot be lifted according to a sequence table, because they are dependant on the experiences of the individual, and the environment it is exposed to. In this case, constraints are overcome by development of competency through learning.

We have studied both types of constraints in our work. Our work on type A constraints, has focused on using thresholds on metrics, such as novelty and habituation, to trigger their removal in a semi-structured manner [22, 24]. Our work on type B constraints explores the possibility of behavioural stages emerging internally when sufficient structure has been created to support another stage of behaviour [25].

3. Gaze Architecture

We now briefly describe our architecture for gaze control, and its biological grounding. For a full description we refer the reader to [24].

The architecture (Shown in Figure 1) begins at the point a visual stimulus has been selected for fixation, and the desired gaze displacement is known. In the human, gaze shifts are encoded topographically in the deep layers of the superior colliculus as single movements, rather than individual movements of the eye and head [26]. They are also initially encoded in retinal coordinates, before being transformed into body-centred coordinates further downstream [27]. Correspondingly, we encode the gaze shift as a single displacement in retinal coordinates on the *retinal map*.

Current thinking suggests the output of the superior colliculus is then processed on two separate pathways in the brainstem corresponding to the control of the eyes and neck [28]. We use two *mappings* (see below) to convert the gaze displacement, in retinal coordinates, to separate relative motor commands for the eye and head.

Although eye and neck movements are triggered separately, there is some cross coupling between the two [29]. Gaze shifts are not made up of a particular ratio of eye and head movements, but the amplitudes of both can be accurately predicted given knowledge of the initial position of the eye in the socket and the size of the gaze shift. We use a model by Wang and Jin [30], which calculates the gaze decomposition for saccades made in response to unexpected targets, to modulate the contribution to the gaze shift made by the head. Whilst a reduced head displacement is sent to the neck motor system, the full gaze displacement is sent to the eye motor system. This is necessary as the eye moves much faster than the head, and saccades usually reach the target well before the head has made its contribution. An eye counter-rotation mechanism is triggered during head movements to ensure the eye does not overshoot and continues to fixate the target as the head completes its movement post-saccade.

There is continuing debate about the feedback process involved during the saccade process, with current thinking pointing to gaze-error and motor-error feedback loops [29, 31, 32]. Our architecture is based on a model proposed by Freedman [32], in which the gaze displacement signal is split into eye and neck components before being subjected to the error feedback signal.

The architecture uses a series of *maps* [22] to represent sensorimotor spaces relating to the visual, proprioceptive, and motor control of the eye and neck. These maps are 2-dimensional structures representing the topological neural maps in the cortex. Each map is covered by of a set of overlapping *receptive fields*, characterising groups of neurons sensitive to stimuli in that region, which effectively partition the space into regions of equivalence (see Figure 1): stimuli occurring within a field are treated as having occurred at the same point on the map. Fields in different sensorimotor spaces are linked by *mappings*, which represent neuronal connections, allowing the location of a visual stimulus to be mapped to the motor movement required for fixation.

These saccade mappings are learnt through the coactivation of fields, based on the following simple process:

1. A stimuli, selected for attention, excites a field on the retina, f_s, corresponding to its visual location
2. The robot makes an eye movement, m_e in an attempt to fixate on the stimulus
3. If the saccade results in fixation, the field covering the motor movement m_e is mapped to the initial stimulus location, f_s
4. Else, return to (1)

An important issue is that another biological constraint prevents learning of the head-retina mapping directly. Whenever the head is moved, a hard-wired reflex action causes the eyes to counter-rotate to maintain fixation, resulting in no change in the retinal input. Thus, in learning this mapping we must incorporate knowledge about the impact of eye movements on visual stimuli (see [24] for a detailed description of the process). The result is that an accurate mapping of head movement to retinal change requires prior learning of an accurate saccade mapping. This creates an interesting set of constraints on learning, that form part of our experiments.

With the mappings learned, as will be described in the following sections, the architecture can be used to make eye-head gaze shifts to fixate on visual targets in the following way:

1. A stimuli, selected for attention, excites a field on the retina corresponding to its visual location
2. Mappings from the stimulated field are followed to corresponding fields on the eye-saccade and head-movement maps, and in doing so are converted into relative motor movements, which are then sent to the motor systems
3. The proportion of the gaze shift to be made by the head is calculated based on the size of the gaze shift and the initial position of the eye. This then modulates the head displacement command before it reaches the motor system
4. The dynamics of the motor system result in the eye reaching the target before, or early during, the head movement
5. Whilst the head completes its movement the eye counter-rotates to maintain fixation. A limiter on the eye counter-rotation stream (not shown) prevents counter rotation until the eye has acquired the target. This enables long saccades, where the eye may not be able to reach the target until sufficient head contribution has been made

Whilst we have previously demonstrated the ability of this architecture to learn and perform gaze shifts [24], a key question is how environmental constraints impact on the learning of the mappings between the retinal and motor spaces. This is the focus of the remainder of this paper.

4. Experiments and Results

4.1. Experimental setup

The experiments presented in this paper focus on the type B constraint when applied to learning gaze control on an iCub robot, and how this interacts with the presence of a type A constraint. The type B constraint is produced from external factors in the environment, in this case the number of visual targets from which the eye control can be learnt. In the constrained scenario, a single target is presented directly in front of the robot, while in the unconstrained case, three visual targets are presented spread out across the space that can be fixated using just the eyes, whilst the head remains centred. The effect of the constrained scenario is to limit the maximum size of saccades that can be learnt whilst only moving the eyes. During these experiments a type A constraint is used to delay the onset of head learning. Once the head constraint is lifted the effect of the type B constraint disappears.

The visual targets are plain coloured cubes that can easily be identified in the images received from the iCub cameras using colour blob detection. They remain static in the environment, whilst the eyes (cameras) move to a new random starting position before each saccade. A visible object is then selected at random as the target for the saccade. The central area of the camera images are designated as the foveal region, with a diameter of approximated 10% of the image width. Objects appearing within this fovea are considered to be fixated.

Both the eyes and head have the use of two degrees of freedom during the experiments, although the head joints are locked until the type A constraint has been released. In the eyes, the degrees available are tilt and version (i.e. panning of both eyes), whilst the head uses yaw and pitch. This gives rise to a set of two dimensional mappings from the 2D image frame to the two motor dimensions. The motor space was represented in terms of the relative movements required to bring the target to fixation.

The experiments are run over a one hour period, with the timed release of the head constraint varied in ten minute intervals across different experiments. The eye learning progresses quickly, and is capable of learning a functional mapping in the first 5-10 minutes. However, the development of the head sensorimotor mapping is considerably slower, learning individual links less frequently from saccades. As a result, this stage of the learning requires more time, with initial trials showing that a one hour period of learning split between eye and head development gives rise to a mapping that is sufficiently populated that it can be successfully applied. With further development, the head mapping fills out to give a strong combined gaze control system. An analysis of the effects caused by a type A constraint on learning with no environmental constraints is presented in [24].

4.2. Map development

The map shown in Figure 2 is the retina coverage learnt from a single target, with the head constrained throughout, while the map shown in Figure 3 is the retina coverage learnt from three available targets, with the head constrained. Only the retina maps are shown, to highlight the coverage of the visual space. After 60 minutes, the maps have become saturated in the regions that are currently reachable, showing

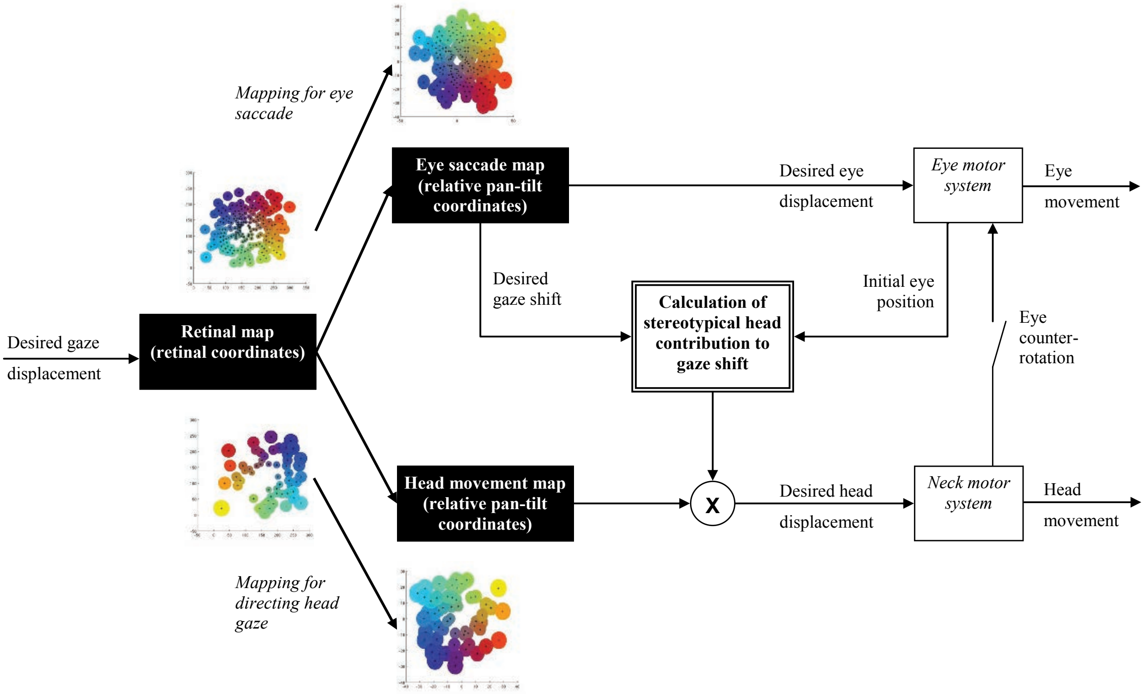

Figure 1. Our gaze control architecture (described fully in [24]). Eye and head control is learnt through a pair of topographic mappings linking stimuli in retinal coordinates to the corresponding motor movements required for fixation. Stereotypical head contributions are calculated from the initial eye position and size of gaze shift using formulae from Wang and Jin [30]. An eye counter-rotation mechanism adjusts for movements made by the neck to ensure the eye maintains fixation during head movements.

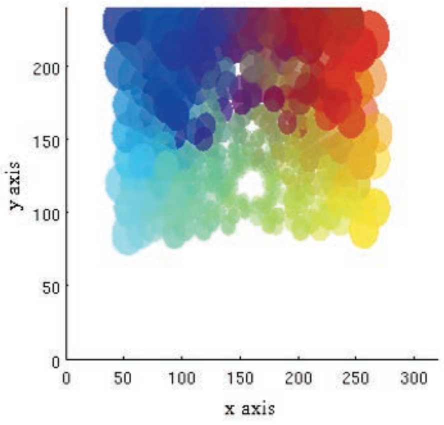

Figure 2. Retina eye map from learning with a single target for 60 minutes without any head movement

the full scope of what can be learnt from the visual targets available. The scale on the two maps are the same, so the difference between them illustrates the additional coverage obtained when three targets are available and spread out across the visible space. The reason for the restriction, caused by only having a single target available to learn from, is that the maximum horizontal saccade that can be made is only half the range of the motors, whilst in the case where three targets are present, almost the full horizontal range of the motors can be learnt. Similarly, most of the fields fall in the top half of the retina map due to a vertical offset of the target. Clearly, different arrangements and numbers of targets would give rise to different initial mappings, when learning with only the eye movements enabled.

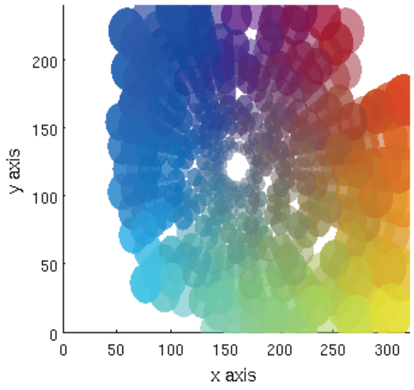

Figure 3. Retina eye map from learning with three targets for 60 minutes without any head movement

The mapping resulting from learning with a single target is concentrated within a small area, whilst the learning from three targets initially develops a wider coverage of fields spread out over a larger area. Both mappings have their advantages; the constrained learning developing a mapping with a detailed and focused coverage across a smaller area, whilst the less constrained learning develops a mapping with a wider coverage. This wider coverage can help when the head learning is enabled as the extra coverage increases the probability of finding an existing eye link to learn the retina displacement. However, having a more constrained mapping means the links will have been tested and refined more frequently.

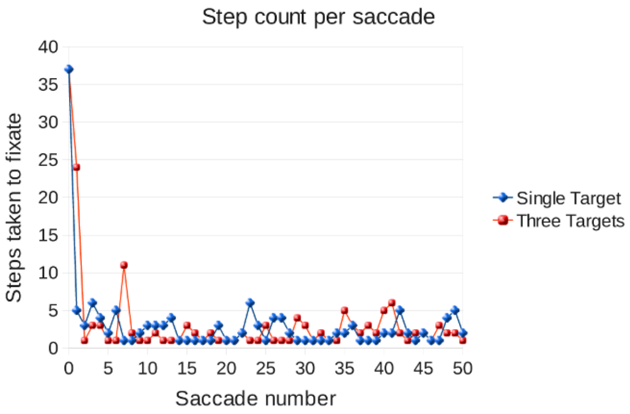

Figure 4. Graph showing the number of movements taken to fixate on a target. Only the first 50 saccades are shown, with no head movement involved. Best viewed in colour.

Figure 4 shows the number of steps taken to fixate on the target. In both the single target and three target experiments, the first saccade consists purely of random motor movements until the target is fixated. This can take a large number of steps, however once the target has been fixated, each step taken can be used to learn a link in the mapping by chaining together the motor movements made and the position of the target at each step. In this case, almost 40 steps were required on both initial saccades, however after this point, the number of steps reduces significantly as existing links are more likely to be encountered and reused. When a single target was present and the whole 60 minutes was spent using only eye movements, a total of 672 saccades were completed. The average number of steps for the first 100 saccades was 2.37, whilst the last 100 saccades had a step count average of 1.53. Even towards the end of the hour there is still learning occurring in the system, either where gaps are found in the mapping, or links are being refined. An example of the early motor movements can be seen in Figures 5 and 6, showing the first saccade randomly exploring the space, whilst the others quickly encounter existing fields that are able to bring them either directly to, or close to, the target. Note the motor movements have been adjusted relative to the target, so the target appears at (0,0) in this relative motor space.

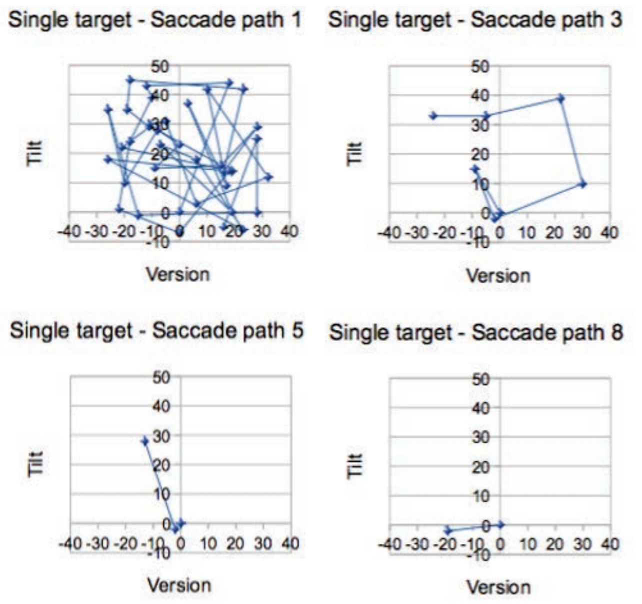

Figure 5. Motor movements for fixating on a single target. Movement tracks are taken from saccades 1, 3, 5 and 8 showing the rapid early development. Movements are relative to the target location.

In the case of the learning based on three targets, the area that can be learnt with eyes alone is greater. This means that the initial exploration can discover more gaps in the maps, requiring more learning early on. Over a 60 minute period of eye only movements, 803 saccades were completed to three targets. The average number of steps from the first 100 saccades was 2.45 showing slightly more steps taken to fixate on the initial targets, however the average over the last 100 saccades was just 1.08 showing a high percentage of single step saccades towards the end.

While the environment constrains the scope of learning using just the eyes, the release of an internal constraint can counteract the environmental constraint. In this case, when the head starts to move, the range of positions in which the target can appear on the retina increases. This allows the eye saccade system to then continue learning the areas of the map that were previously unreachable. Figure 7 shows this effect when the constraint on head movement was released after 10 minutes of eye learning, followed by a further 50 minute period of learning with both systems enabled. Similarly, in Figure 8, the retina maps learnt from three targets are shown. Interestingly, the final coverage from learning with a single target includes areas that were not learnt with three targets in the same time frame, even though the single target formed one of the three targets used in the second scenario. This could be because additional eye motor babbling occurs when the head constraint is released, increasing the exploration in these previously inaccessible areas, whilst the mapping learnt from three targets already has a good coverage meaning there is always a link nearby that could be used to bring the target close to the centre for a two step saccade, reducing the amount of random babbling and opportunities for learning new links.

The graphs shown in Figures 9 and 10 show the percentage coverage of the retina map as the number of links learnt increases. Within each scenario, the eye only learning (solid lines) follows the same curve, with a final coverage of just under 50% reached from learning with a single target, compared to almost 80% coverage obtained when three targets are available. When the head constraint is released (dashed lines), this has very little impact on the rate of learning or coverage learnt from three targets, however when only a single target is available it makes a noticeable impact. The percentage of coverage increases dramatically after the release of the head constraint, ultimately learning a greater coverage than that gained from three targets.

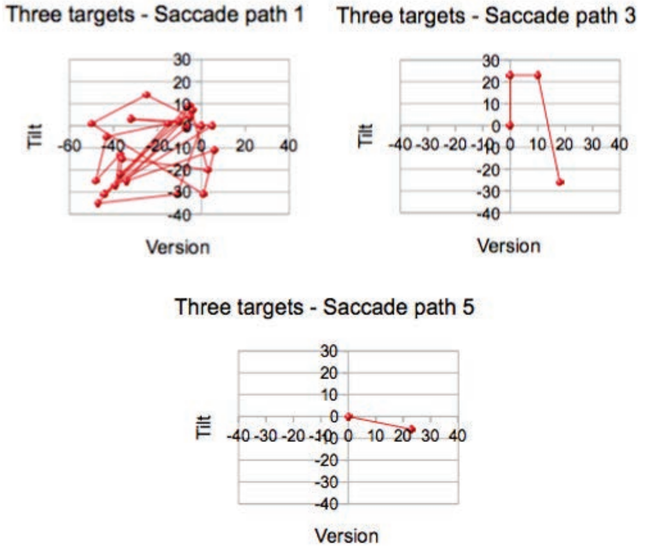

Figure 6. Motor movements for fixating on three targets. Movement tracks are taken from saccades 1, 3 and 5 showing the rapid early development. Movements are relative to the target location.

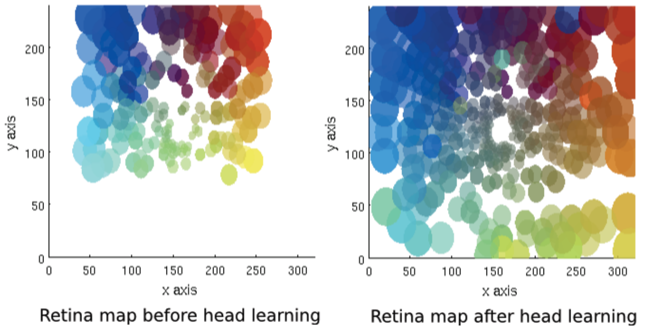

Retina map before head learning

Retina map after head learning

Figure 7. Retina eye map from learning from a single target for 10 minutes with just the eye, then a further 50 minutes with the head moving.

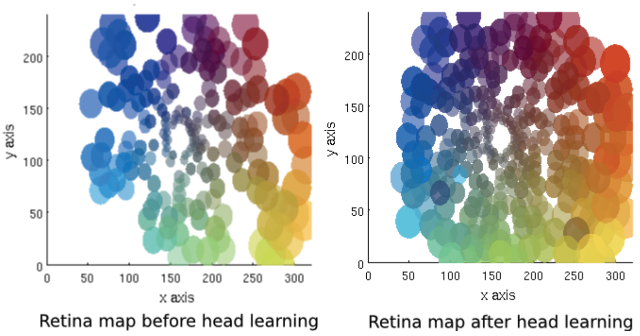

Retina map before head learning

Retina map after head learning

Figure 8. Retina eye map from learning from three targets for 10 minutes with just the eye, then a further 50 minutes with the head moving.

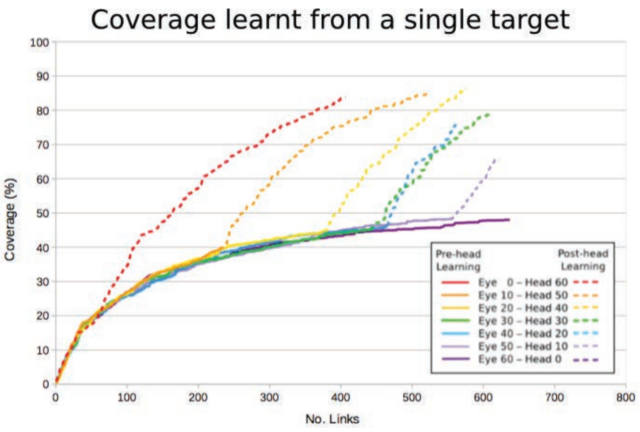

Figure 9. Graph showing the percentage of map covered as more links are learnt, based on learning with a single target for 60 minutes with the head constraint being released after 10 minute intervals across different trials (solid lines represent eye-only phase). Best viewed in colour.

Although these graphs do not show the timing at which the links are learnt, they do give an indication of the rate of learning. Initially when no links are present in the mappings, the opportunity to learn new links is very high. As time passes and the maps become more populated, so the number of steps taken to fixate on the target is reduced, limiting the number of links that can be learnt per saccade. As a result the rate of learning decreases. Where three targets are present, the learning is able to cover up to 80% of the mapping, however with a single target and only eye movements being made, less than 50% of the mapping is reachable. Once the head constraint is lifted, further regions of the maps are effectively opened up, giving rise to an increase in the rate of learning as the large gaps in the map require more steps to fixate on the target once again.

The amount of coverage in the eye mapping when the head constraint is released affects the performance of the head map development. Low coverage of the retina map slows down the learning of the head mapping, as additional time after the release of the head constraint is required to improve the eye mapping. However, in both scenarios, when the head mapping was released early, after 10 minutes, the retina and head maps had slightly greater coverage by the end of the 60 minutes. In the case of the head gaze maps (see Figure 11), this was 15.7% coverage learnt from a single target compared to 13% learnt with three targets. Although these numbers seem low, the links that are learnt can be used straight away. However, the frequency in which existing links are encountered increases as the coverage increases.

4.3. Map comparison

The mappings that were generated were tested to assess their coverage of the visual area. A set of 100 initial eye and head motor configurations were selected, with a single target in front of the iCub. The mappings were then queried to find links from where the target appeared in the image to fixate on the target using a single step saccade. The first 50 positions had the head centered and only the eyes were allowed to move to fixate on the target, whilst the second half also changed the initial position of the head and if a matching head link existed then a combined eye and head saccade was used to fixate on the target.

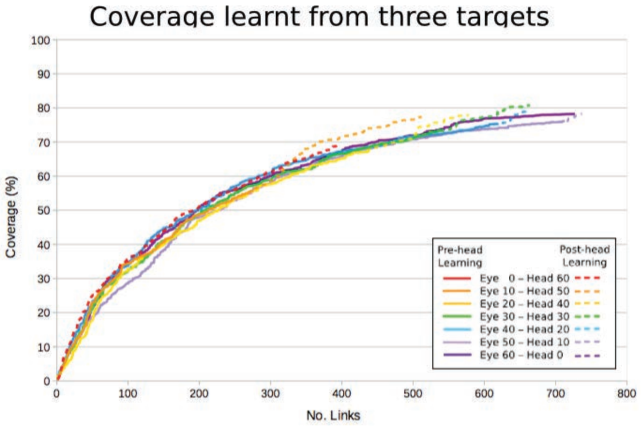

Figure 10. Graph showing the percentage of map covered as more links are learnt, based on learning with three targets for 60 minutes with the head constraint being released after 10 minute intervals across different trials (solid lines represent eye-only phase). Best viewed in colour.

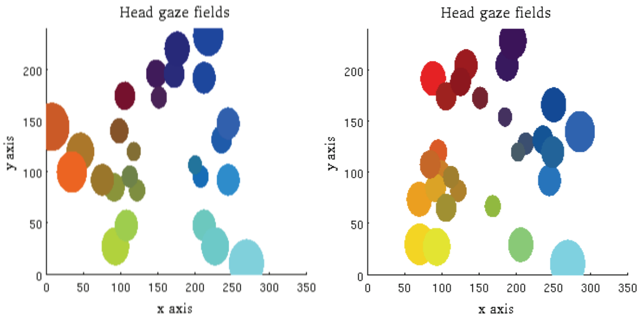

Single target, 50 minutes

Three targets, 50 minutes

Figure 11. Head gaze maps after learning for 10 minutes of eye only learning, followed by 50 minutes of learning with both eye and head systems enabled. Left: single target scenario, Right: three targets scenario

Table 1. Number of targets fixated based on mappings learnt with one and three targets using only the eyes

Trial Eye-Head	1 Target Map Targets fixated (%)	3 Target Map Targets fixated (%)
0-60	17 (34)	12 (24)
10-50	21 (42)	34 (68)
20-40	29 (58)	26 (52)
30-30	25 (50)	38 (76)
40-20	33 (66)	33 (66)
50-10	30 (60)	38 (76)
60- 0	32 (64)	37 (74)

Table 2. Number of targets fixated based on mappings learnt with one and three targets using the eyes and the head

Trial	1 Target Map		3 Target Map	
Eye-Head	Targets fixated (%)	Combined saccades (%)	Targets fixated (%)	Combined saccades (%)
0-60	15 (30)	0 (0)	15 (30)	2 (4)
10-50	21 (42)	5 (10)	18 (36)	2 (4)
20-40	16 (32)	1 (2)	25 (50)	4 (8)
30-30	21 (42)	2 (4)	30 (60)	3 (6)
40-20	22 (44)	2 (4)	24 (48)	2 (4)
50-10	24 (48)	2 (4)	30 (60)	1 (2)
60- 0	20 (40)	0 (0)	32 (64)	0 (0)

Tables 1 and 2 show the number of saccades that resulted in fixation based on the available links during the two phases of testing.

On average, the mappings learnt with three targets present were able to complete more saccades in both phases of the testing. An interesting point is in the combined eye and head saccades, during the second phase of the testing. The number of successful combined saccades, in the mappings learnt using three targets, peaks where the head constraint was released after 20 minutes. As the links are more spread out when learning from three targets, it takes longer to gain a good coverage and reuse existing links to check their accuracy, both of which are needed to learn the head mapping. After 20 minutes of eye only learning, the eye mapping has a good coverage with reliability, and still have time to learn a good head mapping, leading to this trial giving the best overall results for the eye and head mapping.

Meanwhile in the single target scenario, most of the combined saccades come from one early mapping, where the head constraint was released after just 10 minutes. This illustrates the additional time required to fill out the coverage of the eye-mapping after the head constraint is released, and still have time to be able to learn a good coverage in the head mapping.

5. Conclusion

We have compared the learning of the gaze control when restricted by a type B constraint. When learning with only a single target visible, there was a clear limitation on the area of the eye mapping that could be learnt, a restriction that was then lifted by the movement of the head, allowing the eye mapping to continue developing whilst starting to learn the head mapping. This effect slowed the rate of learning in both mappings, resulting in the learning taking longer to reach a similar level.

The difference in field distribution, and hence coverage, can be both an advantage and a disadvantage. In one case, a densely populated mapping over the focal region is generated with a more sparse coverage of the surrounding regions where less detail is required. On the other hand, the less constrained learning initially generates a more even mapping that covers a larger area without specialising at any particular point. In terms of biological development, this is an example of a critical phase in learning, illustrating the relationship between the sensorimotor mappings of the eye and head when developing gaze control.

The literature on infant development has inspired the two categories of constraints, A and B, that can be identified as influencing the developmental trajectory. The timing of the release of head learning is used to illustrate type A constraints where a clearly defined sequence is given for the order in which the development proceeds. Meanwhile, the availability of objects within the environment is used to simulate type B constraints.

Regardless of the structure of the environment, both cases produce mappings that are capable of combining links from the eye and head mappings to direct the gaze at desired targets. In the case of the constrained environment, the release of the head constraint is required to learn the full eye mapping, while in a less constrained setting, the same mapping can be learnt with just the eye movements. If the release of the head constraint is delayed, this will significantly delay the time taken to gain the full coverage of the eye mapping. However, whilst the head learning is constrained, a good coverage can be developed around the foveal region where a detailed mapping is required, allowing a sparser peripheral mapping, where less detail is required, to be learnt later.

Overall, the limitations on the learning can be internal to the system, as considered in [24], or they can be external to the system as shown here. When only the eye is moving in a constrained environment, the focus is on a single object directly in front of the robot. From this point, only small motor movements are ever required to build up a mapping that is sufficient to cope with this environment. Without modifying the physical environment, the introduction of the head frees up a much greater range of motions, allowing for larger saccades to be learnt to the single object and potentially making additional targets, that were previously out of visual range, reachable.

Acknowledgments

This research has received funds from the European Community 7th Framework Programme (FP7/2007-2013), "Challenge 2 - Cognitive Systems, Interaction, Robotics", grant agreement No. ICT-IP-231722, project "IM-CLeVeR - Intrinsically Motivated Cumulative Learning Versatile Robots".

References

[1] M. Asada, K. Hosoda, Y. Kuniyoshi, H. Ishiguro, T. Inui, Y. Yoshikawa, M. Ogino, and C. Yoshida, "Cognitive developmental robotics: A survey", *IEEE Transactions on Autonomous Mental Development*, vol. 1, no. 1, pp. 12–34, 2009.

[2] J. Piaget, *The Child's Conception of the World*, Paladin, London, 1973.

[3] R. Pfeifer and C. Scheier, "Sensory-motor coordination: the metaphor and beyond", *Robotics and Autonomous Systems*, vol. 20, no. 2, pp. 157–178, 1997.

[4] L. Smith and M. Gasser, "The development of embodied cognition: Six lessons from babies", *Artificial Life*, vol. 11, no. 1-2, pp. 13–29, 2005.

[5] D. L. Gallahue, *Understanding Motor Development in Children*, John Wiley, NY, 1982.

[6] R. M. Angulo-Kinzler, B. Ulrich, and E. Thelen, "Three-month-old infants can select specific leg motor solutions.", *Motor Control*, vol. 6, no. 1, pp. 52–68, 2002.

[7] Darwin Muir and Jeffery Field, "Newborn infants orient to sounds", *Child Development*, vol. 50, pp. 431–436, 1979.

[8] J. F. Rosenblith, *In the Beginning: Development from Conception to Age Two*, Sage Publications, Newbury Park, CA, second edition, 1992.

[9] C. M. Harris, M. Jacobs, F. Shawkat, and D. Taylor, "The development of saccadic accuracy in the first seven months", *Clinical Vision Sciences*, vol. 8, no. 1, pp. 85 – 96, 1993.

[10] R. Griffiths, *The abilities of babies*, University of London Press, London, 1954.

[11] M. D. Sheridan, *From Birth to Five Years*, NFER Publishing, Windsor, USA, 1973.

[12] D. Maurer and C. Maurer, *The World of the Newborn*, Basic Books, New York, 1988.

[13] A. E. Milewski, "Infants' discrimination of internal and external pattern elements", *Experimental Child Psychology*, vol. 22, pp. 229–246, 1976.

[14] R. N. Aslin, "Visual and auditory development in infancy", in *Handbook of Infancy*, J. D. Osofsky, Ed. John Wiley, New York, second edition, 1987.

[15] J. Oates, C. Wood, and A. Grayson, *Psychological Development and Early Childhood*, Blackwell Publishing, Malden, MA, 2005.

[16] E. Tronick, "Stimulus control and the growth of the infant's effective visual field", *Perception & Psychophysics*, vol. 11, no. 5, pp. 373–376, 1972.

[17] M. R. Fiorentino, *A Basis for Sensorimotor Development, Normal and Abnormal: The Influence of Primitive, Postural Reflexes on the Development and Distribution of Tone*, Charles C Thomas, Springfield, Ill, 1981.

[18] F. Goodkin, "The development of mature patterns of head {eye coordination in the human infant", *Early Human Development*, vol. 4, pp. 373–386, 1980.

[19] J. Law, M. Lee, M. Hülse, and A. Tomassetti, "The infant development timeline and its application to robot shaping", *Adaptive Behaviour*, vol. 19, pp. 335–358, 2011.

[20] G. Metta, G. Sandini, D. Vernon, L. Natale, and F Nori, "The icub humanoid robot: an open platform for research in embodied cognition", in *Proceedings of the 8th Workshop on Performance Metrics for Intelligent Systems*, Washington DC, USA, 2008, ACM, pp. 50–56.

[21] J. Law, M. Lee, and M. Huelse, "Infant development sequences for shaping sensorimotor learning in humanoid robots", in *Proc. 10th Int. Conf. on Epigenetic Robotics*, 2010, pp. 65–72.

[22] M. H. Lee, Q. Meng, and F. Chao, "Staged competence learning in developmental robotics", *Adaptive Behaviour*, vol. 15, no. 3, pp. 241–255, 2007.

[23] A. Needham, T. Barrett, and K. Peterman, "A pick-me-up for infants' exploratory skills: Early simulated experiences reaching for objects using 'sticky mittens' enhances young infants' object exploration skills", *Infant Behavior and Development*, vol. 25, no. 3, pp. 279–295, 2002.

[24] J. Law, P. Shaw, and M. Lee, "A biologically constrained architecture for developmental learning of eye-head gaze control on a humanoid robot", *Autonomous Robots*, 2013, Accepted.

[25] M. Hülse and M. Lee, "Adaptation of coupled sensorimotor mappings: An investigation towards developmental learning of humanoids", in *SAB 2010, LNAI 6226*, S. Doncieux, Ed. 2010, pp. 468–477, Springer.

[26] R. A. Marino, C. K. Rodgers, R. Levy, and D. P. Munoz, "Spatial relationships of visuomotor transformations in the superior colliculus map", *Journal of Neurophysiology*, vol. 100, no. 5, pp. 2564–2576, 2008.

[27] E. M. Klier, H. Wang, and J. D. Crawford, "The superior colliculus encodes gaze commands in retinal coordinates.", *Nat Neurosci*, vol. 4, no. 6, pp. 627–32, 2001.

[28] N. J. Gandhi and H. A. Katnani, "Motor functions of the superior colliculus", *Annual Review of Neuroscience*, vol. 34, no. 1, pp. 205–231, 2011.

[29] H. H. L. M. Goossens and A. J. van Opstal, "Human eye-head

coordination in two dimensions under different sensorimotor conditions", *Experimental Brain Research*, vol. 114, pp. 542–560, 1997.

[30] X. Wang and J. Jin, "A quantitive analysis for decomposing visual signal of the gaze displacement", in *Proceedings of the Pan-Sydney area workshop on Visual information processing – Volume 11 (VIP '01)*, David Dagan Feng, Jesse Jin, Peter Eades, and Hong Yan, Eds., Darlinghurst, Australia, 2001, vol. 11, pp. 153–159, Australian Computer Society, Inc.

[31] D. Guitton and M. Volle, "Gaze control in humans: eye-head coordination during orienting movements to targets within and beyond the oculomotor range", *Journal of Neurophysiology*, vol. 58, no. 3, pp. 427–459, 1987.

[32] E. Freedman, "Coordination of the eyes and head during visual orienting", *Experimental Brain Research*, vol. 190, no. 4, pp. 369–387, 2008.