PRIFYSGOL
ABERYSTWYTH
UNIVERSITY

Aberystwyth University

A comparison of learning strategies for biologically constrained development of gaze control on an iCub robot

Shaw, Patricia Hazel; Law, James Alexander; Lee, Mark Howard

A comparison of learning strategies for biologically constrained development of gaze control on an iCub robot

Patricia Shaw · James Law · Mark Lee

Abstract Gaze control requires the coordination of movements of both eyes and head to fixate on a target. We present a biologically constrained architecture for gaze control and 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 two strategies (sequential and synchronous) for the learning of eye and head coupled mappings, and give results from implementations on an iCub robot. The results show that the developmental approach can give 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 con-

P. Shaw · J. Law and M. Lee
Department of Computer Science, Aberystwyth University, Aberystwyth, Wales, UK
Tel.: +44 1970 628680
Fax: +44 1970 628536
E-mail: {phs,jxl,mhl}@aber.ac.uk

cept for building and understanding cognitive learning systems (Asada et al., 2009). 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. As such, we are motivated to find methods that are able to learn in real-time without the need for any training data or off-line training phases. Achieving these goals form an important step towards truly autonomous systems that do not rely on human supervision or prior training. While this is a difficult challenge for machine learning approaches, it is achieved by humans and other animals in early infancy.

In this paper we report on experiments on sensorimotor learning for gaze control. Rapid eye saccades to fixate on stimuli of interest suffer interference from head and other body movements, hence learning to coordinate head and eye by selecting appropriate motor movements is a significant problem. In section 2 we outline the relevance of infant behaviour, then section 3 describes the role of constraints in aiding learning in robotics. Section 4 outlines the architecture developed for gaze control and explains the different strategies available for gaze learning while section 5 presents comparison results, which is followed by conclusions.

2 Infant Development and the Importance of Stages

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 consolidation 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, there is a widely recognised general sequence, and this is potentially very significant.

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 (Piaget, 1973). We believe that sensorimotor interaction is also key to learning in robotics, and that algorithms for robotic learning should be rooted 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 (Pfeifer and Scheier, 1997). This view that grounding and early start points are crucial for the growth of adaptive intelligence is very well argued by Smith and Gasser (2005) in their "Six Lessons from Babies". In Law et al. (2013), we further explore the developmental literature, from the neural architectures for gaze control, development during infancy and how constraints can be applied in shaping behaviour.

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 (Gallahue, 1982; Angulo-Kinzler et al., 2002).

3 Developmental Robotics

To create a system of staged development on a robotic platform we must first understand how development occurs in infancy. We are interested in both the manifest improvements in sensory abilities and motor control, and the underlying neurological changes that support these advancements. For this reason we have carried out an extensive review of early infant development covering both the psychological and neurological literature (Law et al., 2011). The data 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 for the iCub humanoid robot (Metta et al., 2008) and produced a comprehensive chart of the general developmental possibilities for the sensorimotor systems of the iCub (Law et al., 2010).

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 (Lee et al., 2007). There are various kinds and sources of constraints, but there are two main types and here we consider these as two alternate implementations 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 constraints limit visual exploration by the infant and reduce the complexity of the perceived environment. Since these limitations are related to the biological stage of growth, they are fairly independent of external factors and can be effectively measured in terms of a relative temporal framework.

Such constraints can be programmed from a sequence table, such as those in (Law et al., 2010), which are extracted from the infant data. Since infants develop at different rates, it is not sufficient to trigger constraints based on age alone, and account should be taken of the relative stage of biological development. In this way, type A constraints can be lifted following a generalised time line, but should be modified to reflect early or late development.

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. (2002) 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 (Law et al., 2013), has focused on using thresholds on metrics, such as novelty and habituation, to trigger their removal in a semi-structured manner (Lee et al., 2007). Our work on type B constraints (Shaw et al., 2012), explores the possibility of behavioural stages emerging internally when
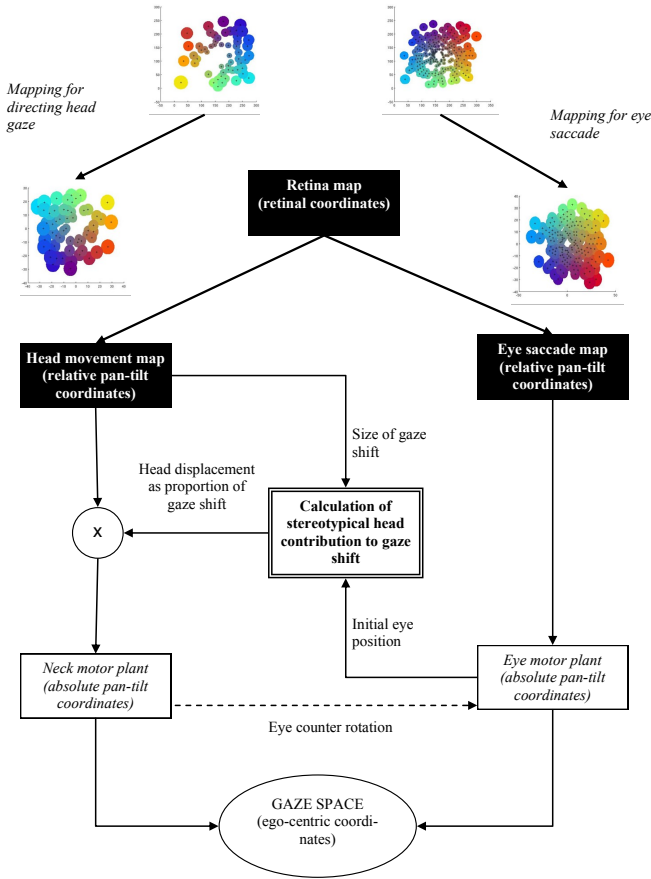
Fig. 1 Architecture for eye-head gaze control, using sensorimotor mappings to implement the neural model proposed by Freedman (2008)

sufficient structure has been created to support another stage of behaviour (Hülse and Lee, 2010). The goal must be to achieve qualitative advances in behaviour without structural change.

Here we compare the effect of both classes of constraint on learning gaze control in an eye-head system.

4 Gaze Architecture

Using our sensorimotor mapping framework (Lee et al., 2007), we have created an architecture for gaze control which incorporates both eye and head movements. This architecture is biologically constrained, using maps to implement a model proposed by Freedman (2008). The relative contributions of eye and head movements to the gaze displacement are based on work by Wang and Jin (2001). The basic architecture of this system is shown in Figure 1. The architecture fully supports developmental learning, and we are using it to investigate the two developmental approaches mentioned above.

Maps are a two dimensional representation of a sensor or motor space that are broken down into fields,

which represent regions of equivalence. Pairs of sensor and motor maps are linked together to form sensorimotor mappings, as described in (Lee et al., 2007). Links are formed between the fields in the two maps to indicate correspondence between the sensor and motor data. For example, a mapping linking visual sensory input to eye motor commands links the corresponding motor values to move a visual target from any field in the visual input map to the centre of the visual map. The usage of links between maps is recorded, taking into account the success or failure of a link when it is reused. Further mappings can be built on top of, or learnt from existing mappings, so requirements can be included on the reliability of links.

Eye saccades are learnt via motor babbling, a strategy of potentially random motor movements, and use just the right hand route (in Fig. 1) including the retina map, the eye saccade map, and the eye motor plant. Head movements are learnt based on learnt eye movements, and use the following process: 1) the eye saccades to a target in the retina map by performing the movement indicated in the eye saccade map; 2) the head makes a random movement whilst the eye maintains fixation by making a counter rotation equivalent to the head movement; 3) a mapping for the head gaze is created based on the initial target location, total head movement, and total eye movement. To perform a gaze shift with both the eye and the head to a stimulus on the retina, the following process is used: 1) relative movements for both the eye and head saccades are selected by following the mappings from the stimulated field in the retina map; 2) the stereotypical head contribution to the gaze shift is calculated based on the size of the gaze shift and the initial position of the eye; 3) eye and head displacements are sent to the motor plants; 4) the dynamics of the system cause the eye to reach the target before, or early in, the head movement; 5) whilst the head moves 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.

The above architecture allows sensorimotor mappings to be learnt for eye movement control, head control, and eye-head interaction and compensation. Although eye-control is a pre-requisite of learning head control, there is no restriction on whether type A or type B constraints are employed. In the remaining sections we shall investigate the impact both methods have on the development of gaze control.

5 Experiments and Results

The experiments presented in this paper compare the effects of the type A and type B constraints when applied to learning gaze control online on an iCub robot. In Law et al. (2013), initial experiments were presented focusing solely on type A constraints, while in Shaw et al. (2012), experiments were presented based on type B constraints. Here we compare the effect of both classes of constraint interacting together on learning gaze control in an eye-head system.

The iCub has three degrees of freedom in the eyes (tilt, version and vergence) and three in the neck (pitch, yaw, roll), however only two degrees of each are being learnt to develop mappings based on pan and tilt. Colour blob detection is used to identify targets in the environment using a single low resolution eye camera. The central 10% of the image is defined as the foveal region and visual targets that are observed within this region are said to be 'fixated'.

The robot learns in a static environment where a selection of different coloured visual stimuli are presented. Some of these can be visually fixated using just the eyes, while others require both eye and head movement to fixate on them. When considering type A constraints, the learning is sequenced due to limitations in the neurological and physiological structures. Within the iCub, these limitations can be translated as thresholds on metrics, accuracy and coverage, which are used to trigger the release[1] of constraints. In the experiments, the thresholds are based on performance and time, with performance measured based on the number of steps taken to fixate on a target. The use of the time threshold, while artificial, is to encourage the learning to continue for longer than is required based on the performance measure. This is to encourage the mappings to become saturated and the learning to level off. In the type A experiments, once the threshold is reached the constraint restricting the learning of the head mappings is released allowing the next mapping to start developing. The learning strategy employing type A constraints will be referred to as sequential learning.

In contrast, there are no constraints applied in the type B experiments. In this case, the learning of both mappings is enabled from the start, with the aim being to highlight any natural constraints that emerge from the system restricting the development. The learning strategy employing type B constraints will be referred to as synchronous learning.

[1] Constraints can also be re-applied if metrics fall back below thresholds.

5.1 Performance

Within this first set of experiments comparing the two learning strategies, the threshold to release the constraint on head learning in the sequential strategy is based on the performance of eye saccades. The learning in both strategies is stopped once a second performance threshold is reached, based on the use of eye and head saccades for gaze direction. A comparison of the two strategies is then based on the time taken, the number of saccades completed and the coverage of the mappings learnt. When learning links in the head mapping, information from the eye mapping is used, however in this first set of experiments there are no guarantees on the accuracy of the information in the eye mapping. As a predetermined field structure is used for the mappings, if a link is added based on experience near the edge of a field, the link my not be applicable across the whole field. Therefore, it can be beneficial to verify the accuracy of a link by reusing it from different points within the field to check its reliability. In the first set of experiments, a link in the eye mapping does not need to be tested to check its accuracy before using that link to learn links in the head mapping.

In the experiments presented below, the two strategies were each run 10 times, with the learning stopping after a minimum of 10 successful combined gaze shifts and a success rate of at least 75% for the combined gaze shifts fixating on a target. The threshold on eye performance for the sequential strategy was defined based on a rolling average of the number of movements required to fixate on a given target, where the threshold was set at less than 2.0 steps per saccade over the last 10 saccades. The eye mapping was allowed to continue developing after the release of the head learning constraint, however if after releasing the constraint the eye performance degraded at any point, the constraint could be reapplied to refocus the learning on improving the eye mapping again.

Figure 2 shows an example of the typical data obtained from a sequential and synchronous test. At the start, both learning strategies are making large numbers of steps to fixate on the target as they randomly babble around. The fixation on the first target can randomly be achieved quite quickly, as in the sequential approach here, which initially fixated in just 9 steps, while the synchronous approach took 25 steps. This is not related to the specific approach and in other runs the first saccade for both approaches took much longer, with an average across all initial saccades of 51 steps. At the end of this first saccade, the first set of links is learnt in the mapping, with the number of links added related to the number of steps taken. The more ran-
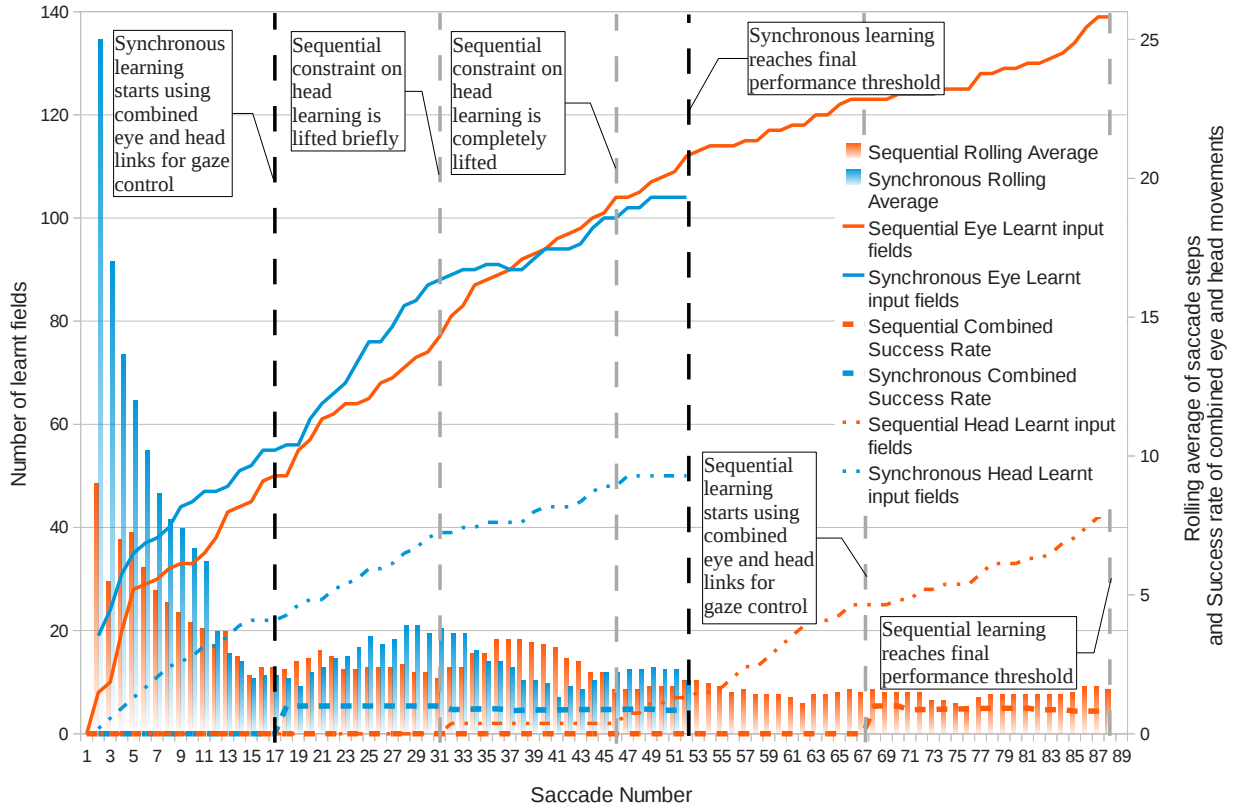
Fig. 2 Example learning rates from the sequential and synchronous learning with threshold constraints based on performance. The maximum number of fields available to learn is 586 and the success of combined head movements is measured between 0 and 1

dom steps in the first saccade the greater the coverage of the maps and hence links from this first saccade. This leads to a rapid decrease in the number of steps in subsequent saccades to find an existing link that can be followed to fixate on the target as is seen in the synchronous step counts. The size of the rolling average block is ten saccades, so after the tenth saccade there is a significant drop as the large size of the first saccade is no longer included in the rolling average. Conversely, if the first saccade randomly manages to fixate rapidly on the target, as is seen in the sequential approach in this example, the subsequent early saccades still require a lot of random exploration in order to discover existing links or fixate on the target again. In this case, the reduction in the number of steps per saccades is much slower, with a possible increase in the steps required over the first few saccades.

In terms of the overall number of saccades required between the two approaches, the synchronous learning strategy consistently requires less saccades to reach the desired level in performance. Comparing the breakdown of the saccades in the sequential approach, shown in Table 1, the number of saccades involving the head is less than the total number of saccades required in the syn-

Table 1 Data for time taken, number of saccades and number of learnt fields for sequential and synchronous learning

		Sequential		Synchronous	
		mean	s.d.	mean	s.d.
Time (seconds)	Eye Only	111.2	22.3	—	—
	Eye & Head	487.3	72.3	566.6	74.4
	Total	598.5	79.5	566.6	74.4
Saccades	Eye Only	39.1	14.3	—	—
	Eye & Head	39.2	6.1	49.4	5.8
	Total	78.3	13.7	49.4	5.8
Learnt Fields (%)	Eye	25.6	2.6	20.5	2.6
	Head	7.5	1.4	9.4	1.2
	Total	16.7	1.6	15.0	1.7

chronous approach. This shows that some initial learning of the eye mapping is always required before it is possible to start learning links in the head mapping. In the example shown in Figure 2, the first head links in the synchronous learning strategy begin learning during the second saccade. Within a learning episode where the initial learning is slower, the delay in learning the head links is much longer, as can be seen in the sec-

ond set of experiments below where the requirements for learning the head mappings are increased to slow the onset of this learning.

Despite the synchronous approach using considerably less saccades to achieve the same level of performance, the amount of time it takes is only slightly less than that required by the sequential learning strategy (Table 1). Comparing the breakdown of the time taken during the two phases of the sequential learning, it is clear that the learning of the eye on its own, without moving the head is very quick, taking on average just 111 seconds, compared to the total duration of nearly 600 seconds. The average length of time spent learning individual eye saccades is 2.8 seconds per saccade, while it takes 12.4 seconds per eye-head saccade during learning. The majority of the time is spent in the phase with the combined eye and head learning, with the slower head movements and the increased possibility of the target being outside of visual range due to the combined gaze direction of the eye and head. These effects will slow down the rate of learning resulting in the overall time saved by the synchronous learning strategy being minimal.
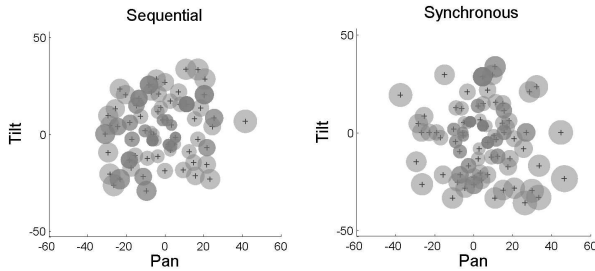
Fig. 3 Illustration of the distribution of the first 100 learnt eye motor fields for sequential and synchronous learning

The maximum number of fields available in any of the eye and head sensorimotor maps is 586, so the number of learnt retina fields shown in Figure 2 gives an indication of the percentage of those fields being linked to eye motor fields. At the point when the synchronous learning finishes, the number of links learnt by the two strategies is very similar, however the distribution of these fields is different. Figure 3 shows the distribution of the first 100 fields on the retina maps learnt during the example learning sessions shown in Figure 2. It is clear just by looking at the distribution that the fields are much more clustered in the sequential learning strategy, whilst the coverage generated during the synchronous learning strategy overlays a broader area with fields further apart. More fields are added for the periphery motor movements than in the sequential mapping where the fields are focused in the central region. Due to the tight clustering in the map from the sequential strategy there is more overlap, reducing the total coverage of the map. This is confirmed when comparing the percentage area of the map that is covered as fields are learnt, as shown in Figure 4, where for example the mean coverage at 100 fields is 23.5% of the area of the map in the sequential learning strategy and 27% in the synchronous learning, with a variance of 1% in the sequential learning and 3% in the synchronous. Despite the apparently low percentage of cover, the performance obtained can still produce good results in terms of number of steps to fixate and the success of combined eye-head gaze shifts to fixate on targets.
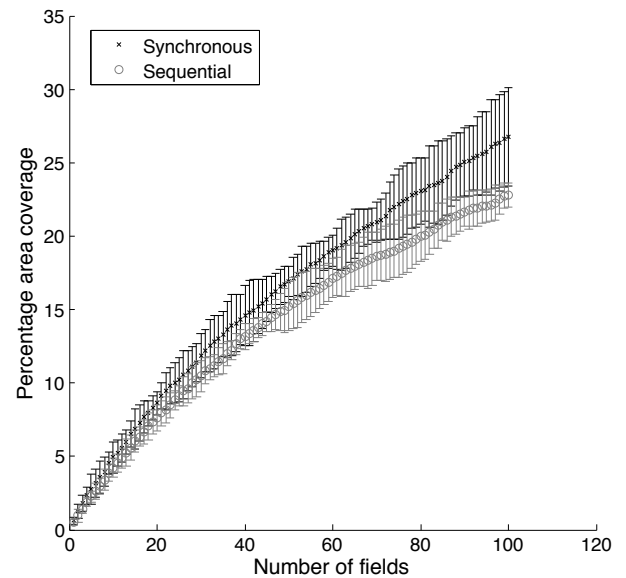
Fig. 4 Graph showing the mean percentage with variance of learnt field coverage for the first 100 fields over the eye motor maps from learning sequentially and synchronously

One interesting feature to notice from the graph in Figure 2 is the increase in the step counter when head movements are introduced during the sequential learning strategy. The introduction of the head movements highlights gaps in the mapping produced whilst only moving the eye. As a result, the head constraint is briefly reapplied whilst this gap is learnt. This increase could be related to the increase in the available gaze space, with fields in the periphery now being stimulated that were not required when simply using the eye on its own. This is backed up by looking at the coverage of the fields that have been learnt, as shown in Figure 3. The head constraint in the example is first released when there are approximately 80 fields in the mapping, however the performance of the eye saccades decreases so this constraint is reapplied until there are approx-

imately 105 fields learnt. During the initial phase of the sequential learning strategy, the head remains still and only a subset of the visual targets can be reached using just the eye motors, and the full motor range is unlikely to be exploited. When the head starts moving, targets are more likely to appear in the periphery of the retina map, requiring a greater range of motor movements to reach them, stretching the coverage of the maps. In the synchronous learning strategy, both the head and eye are moving form the start, allowing more periphery stimulation to be captured.

The increased coverage provided by the synchronous learning allows the same level of performance to be achieved with a smaller number of fields, due to the increased distribution of the fields. This is highlighted in Table 1 where the total number of fields learnt during synchronous learning is less than the total number generated during sequential learning.

5.2 Timed experiments

In this second set of experiments, the focus is aimed at allowing the learning to continue to a point where it is reaching saturation and hence levelling off. This allows evaluation of the learning process, looking for refinements in performance over time. The requirements for developing the head mapping have also been increased, ensuring that any eye link which is used to calculate a head link has now been reused successfully to fixate on a target. Within the sequential learning strategy, once the threshold was reached to enable the head learning, the eye learning was disabled. However, any eye links that failed were 'unlearnt' to allow pruning of the mapping to just the core links that were necessary. The results compare the mappings generated in terms of performance, map coverage, and fields required.

A comparison is performed between the learning output of both sequential and synchronous approaches over a period of six hours, recording results after each hour. In comparison, a basic and sufficient threshold for performance can be reached after just 10 minutes, as seen above. However, by allowing the learning to continue, a higher level of performance is attainable, with a more developed and refined mapping.

In the case of the sequential learning, the 'eye only' phase of the learning was set learning for an hour, whilst the head remained still. This eye mapping was then used for learning the head mapping while the eye learning was disabled. In both the sequential and synchronous experiments, the learning was performed in one-hour segments, with the level of development being compared at the end of each segment.

Figure 5 shows the data obtained over the 6 hours for the sequential and synchronous learning strategies, with the final eye mappings shown in Figure 6. Figure 5(a) focuses on the detailed variation of the eye saccades, clipping the initial saccades to allow the main data set to be more visible. In this illustration, the first saccade in the synchronous learning strategy took 98 steps to fixate, whilst the first saccade in the sequential learning strategy just took 22 steps, both rapidly learning mappings that allowed the target to be reliably fixated in a small number of steps. In this figure, the number of saccades used to calculate the steps per saccade rolling average has also been increased to enable the gradual improvement over time to be clearer.

While both approaches were left learning for the same duration, the actual number of saccades completed was different. In the case of the sequential learning approach, during the first hour only the eye was moving. As the eye moves much faster than the head, it is possible to achieve more saccades in the same duration, so the difference in the number of saccades overall is the effect of the slower head movements in the first hour of synchronous learning. The initial 'eye only' phase of the learning involved approximately 850 saccades in the first hour, however the combined eye and head learning, in both the sequential and synchronous learning approaches, completed an average of 430 saccades per hour.

In this set of experiments, the learning of the links in the head mapping was much more restricted. Despite this, the first link in the head mapping was learnt after just 10 saccades in the synchronous strategy, and later used successfully, before more head links were regularly being learnt after 85 saccades. The increase in the number of combined gaze shifts being used was much more gradual in the synchronous approach, as is seen in Figure 5(b), where as once they start to be used in the sequential strategy they are employed consistently. There is only a small number of fields required in the head mapping before this is able to occur, with some learning continuing afterwards. At the end of the 6 hours of learning, both learning strategies have approximately the same number of fields learnt in the head mapping, with both showing indications that the rate of new fields being learnt is levelling off. This is partly due to the majority of saccades now combining existing links to fixate on targets, however this shows sufficient coverage has been obtained to make the saccades. The final head mappings for the two approaches are shown in Figure 7.

The mappings generated from the learning in the two strategies are shown in Figure 6, with Figure 6(a) showing the state of the sequentially learnt eye map-

(a) Gaze performance   (b) Fields learnt
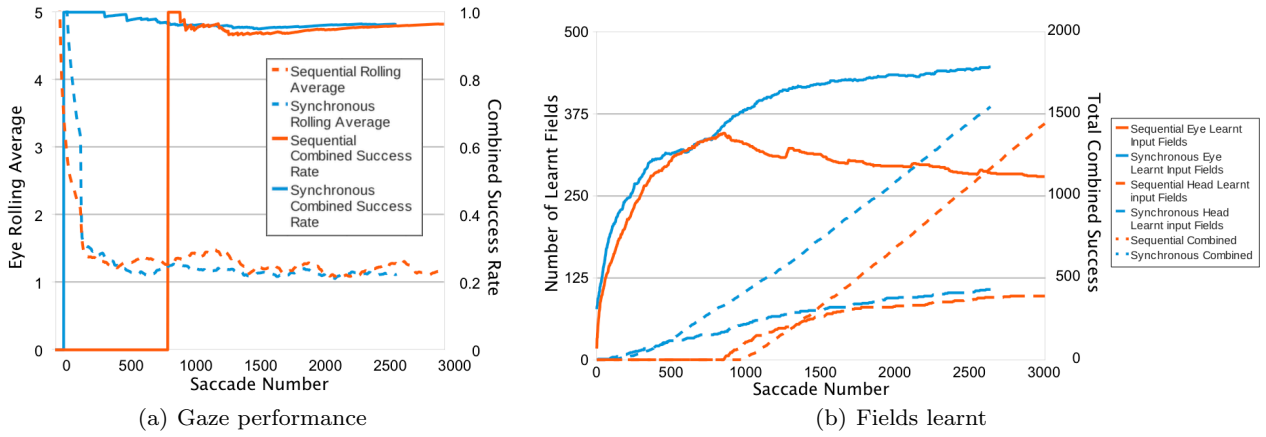
Fig. 5 Set 2 learning rates from sequential and synchronous learning showing gaze improvement and learning rates

ping at the end of the first hour of learning. As the head is stationary during this first phase, the learning is focused on the central area of the maps, generating a dense coverage by the fields. As the head starts moving, the eye is then required to move further to fixate on the same targets that now appear further in the periphery of the vision, to compensate for the offset of the head. At this phase, further learning is then performed to build up mappings around the periphery. Despite the general learning of new eye fields being suppressed, the updating of existing fields is still active. Links that fail cause the fields to be marks as 'unlearnt' to indicate that alternative links need to be tested next time that field is considered. Gradually, the number of fields learnt for the retina and eye motor maps levels off in the sequential mapping to show the minimum number of fields required to be able to perform accurate eye saccades.

In comparison, the learning of the eye mapping is allowed to continue throughout the whole duration of the synchronous experiment, as shown in Figure 5(b). Gradually, the rate of learning begins to level off, and if allowed to continue on for longer would reach a stable state where no more fields were being learnt. The final mapping that is produced is shown in Figure 6(c) where there is greater coverage of the periphery fields. This follows through to the head mappings, shown in Figure 7, where links can only be learnt if there are existing links in the eye mapping. As a result, the sequential mapping has less coverage around the periphery when compared to the synchronously learnt head mapping.
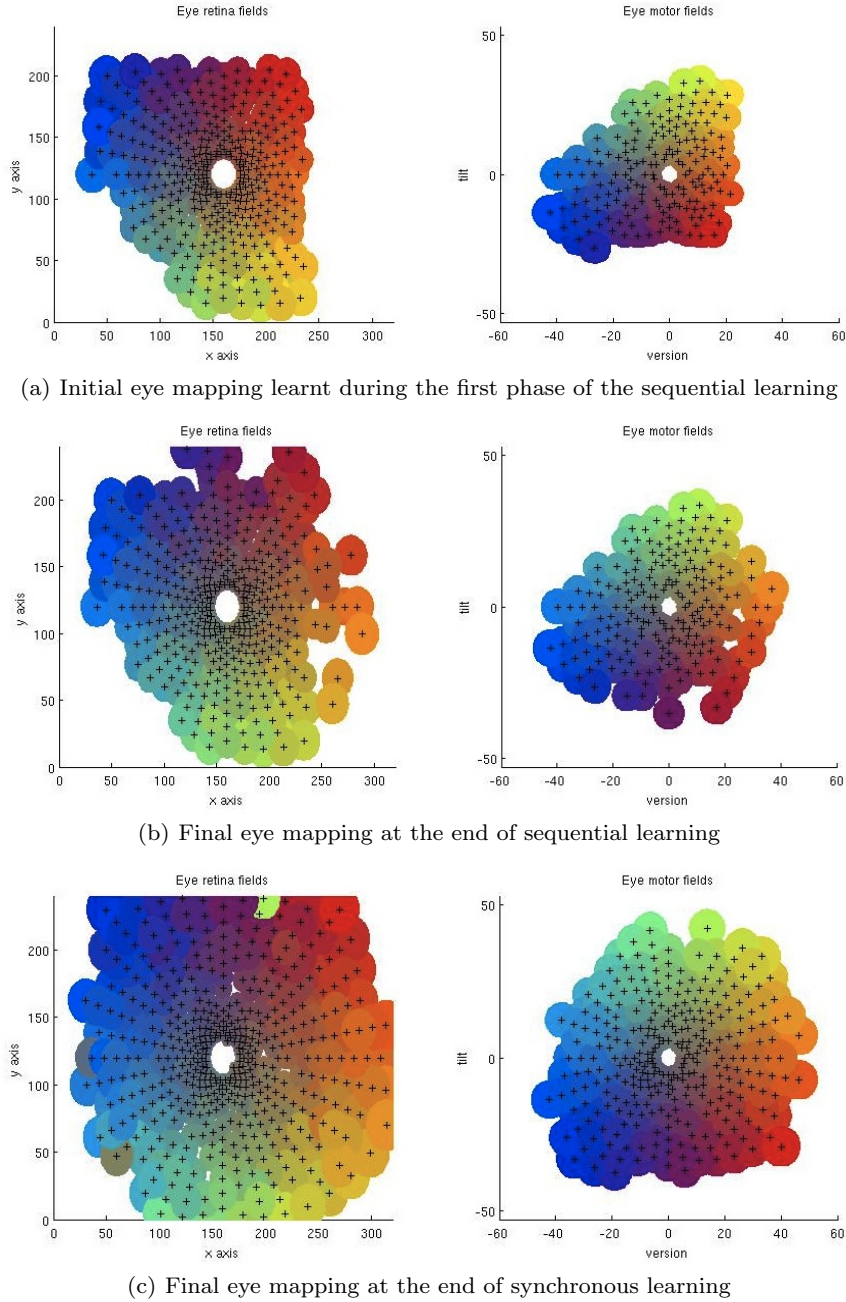
Looking at the resultant maps at the end of the learning for the two approaches shows that the overall coverage from the synchronous learning is visibly greater than the coverage in the sequential case. This is due to the eye mapping continuing to develop for

Table 2 Overview of success rates for sequential and synchronous learning

		Sequential learning	Synchronous learning
Saccades	Eye only	941	-
	Head only	2066	-
	Total	3007	2643
Avg. eye steps to saccade		1.292	1.244
Combined success rate (not including eye-only 'combined' movements)		95.331	96.093
Overall combined success rate		96.441	96.354

the whole duration in the synchronous approach, whilst only a small amount of further development occurred in the sequentially learnt eye mapping when the original mapping proved insufficient to perform the task. In addition, the focus of the learning in the sequential approach is on the central area, where as incorporating the head movement as shown above increases the coverage in the peripheral regions.

As shown in Table 2, the accuracy of the eye saccades and the combined gaze shifts is high. For the eye saccades, the majority of the saccades are fixated in a single step, with a small number requiring multiple steps. The averages shown are the averages across all the saccades, however the average during the final hour is 1.20 and 1.13 for the sequential and synchronous respectively. In the case of the combined gaze shifts, when calculating the proportion of the head movement to use, it is possible that this proportion will be zero. This is because for small gaze shifts within a 'comfort' zone on the eye motors, the head is not required. Subsequently, the performance of the combined gaze shifts is broken down to highlight those gaze shifts which do involve a head movement. During the final hour, in the sequen-

(a) Initial eye mapping learnt during the first phase of the sequential learning

(b) Final eye mapping at the end of sequential learning

(c) Final eye mapping at the end of synchronous learning

Fig. 6 Eye sensorimotor maps showing the links between the retina maps and the eye motor maps. Linked fields are indicated by colour

tial case there are 211 combined gaze shifts and 119 eye only 'combined' gaze shifts out of 440 saccades, compared to 205 combined gaze shifts and 133 eye only gaze shifts, out of 475 saccades in the synchronous strategy. Out of these, just 4 fail (2 combined gaze shifts) in the sequential learning, while 11 (4 combined) fail in the synchronous learning, giving a 99.05% success rate by the end of the sequential learning and 98.04% success rate for the synchronous learning. By the end of the sequential learning, 75.0% of the saccades are reusing

existing links to fixate on the target, while 76.8% of the saccades in the synchronous approach are reusing links. The slightly larger percentage for the synchronous case is due to the greater coverage, suggesting that in both cases if the eye and head mappings were allowed to continue developing, these percentages would likely increase.

While these mappings present the complete set of fields and links learnt, they do not give the full story. The occasional inconsistent links can be seen in both

(a) Final head mapping at the end of sequential learning

(b) Final head mapping at the end of synchronous learning
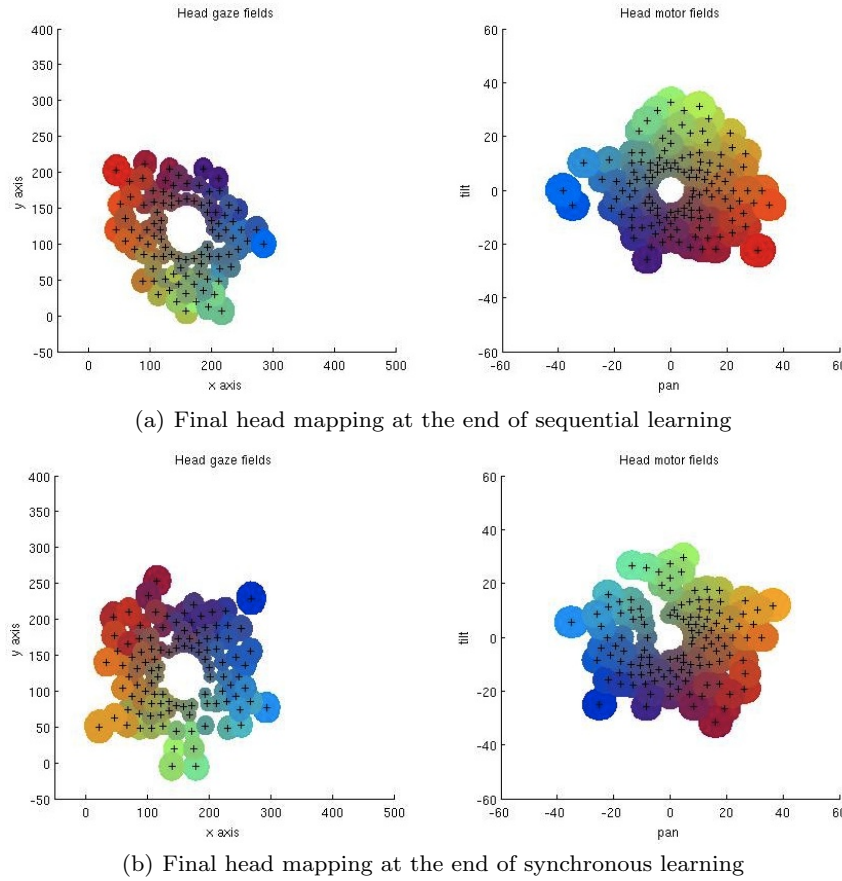
Fig. 7 Head sensorimotor maps linking the retina maps to the head motor maps. Linked fields are indicated by colour

mappings, although more so in the synchronous eye mapping, however these links do not last long when they are identified so do not have a significant impact on performance. In order to really see what the mapping looks like and how the performance levels are obtained it is necessary to look at the amount of usage for the various fields. Figure 8 shows the usage of the fields in the synchronously generated mappings, with a similar result for the sequential mappings. The more transparent a field is, the less it has been used. Any inconsistent links that appear in the full mappings are not visible when viewing the mappings based on usage. While they may not be deleted, they are identified as inconsistent so when selecting links to follow, these links are not used leading to the high degree of performance.

In Figure 5(a), the number of failed combined eye and head movements can be seen in both the sequential and synchronous learning. There is an initial period where the early combined movements are successful, followed by a period during which there is an increase in the number of attempts failing. After this period of failings, the number of failures reduces dramatically. This is strengthened by the usage data that rapidly grows in

terms of the range of usage across the links as shown in Tables 3 and 4.

During the process of learning links, multiple links can be created from the same field. This occurs when the first link is found not to work from a subset of the field, so a new link is added. However, this new link may only work for a different subset of the field. Repeated links between two maps can be made involving a small number of fields until gradually the link with the least amount of error, i.e. the greatest accurate coverage for the field, will emerge as the strongest link that is then repeatedly selected. Figure 9 shows the frequency of the usage of each of the links. Looking at the ranges of the links strengths, it is clear to see that this process is repeated many times in the synchronously learnt eye mapping, where links on some fields are generating strongly negative field strengths, while in the sequentially learnt eye mappings, very few attempts are required to find a strong link. The usage of a small number of links is very high, indicating the existence of key links that are repeatedly used successfully, while the majority of the links are not used at all or very few times. As with the usage maps shown in

(a) Synchronous eye usage mapping

(b) Synchronous head usage mapping
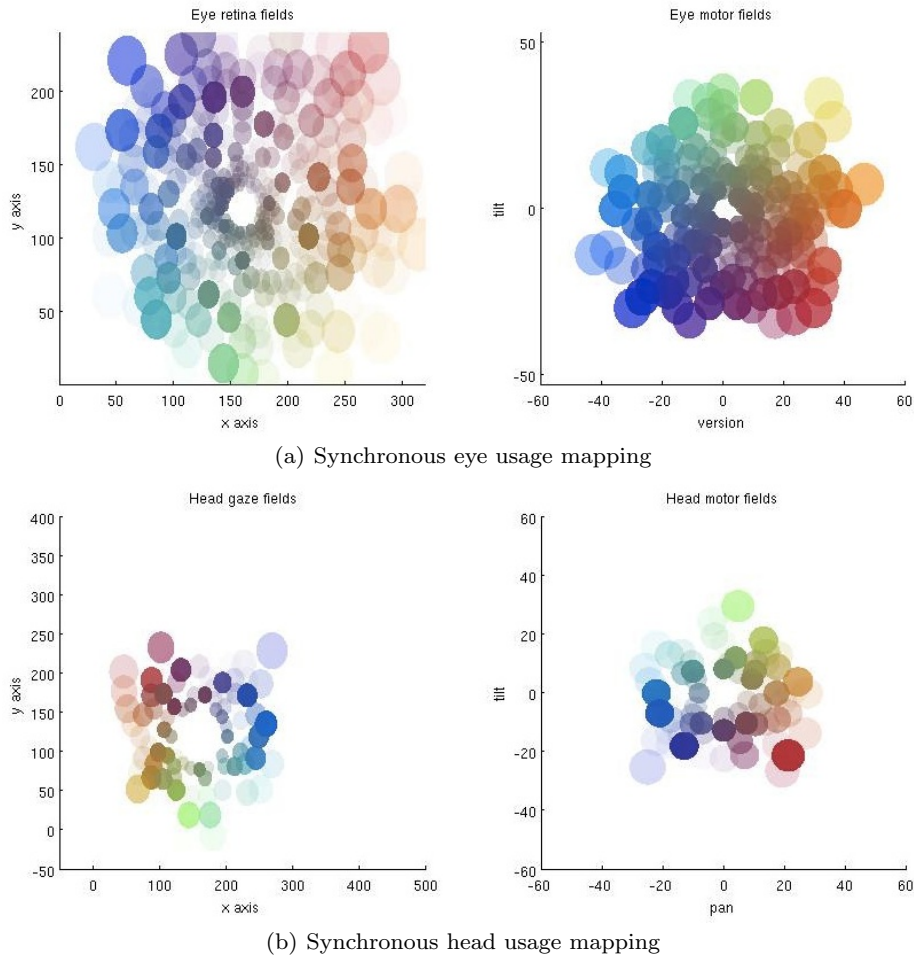
Fig. 8 Final usage mappings from the synchronous learning. Opacity is used to indicate usage with fields that are more opaque having greater usage

Table 3 Break down of the fields and links learnt at each stage during sequential learning

Sequential	1hr	2hrs	3hrs	4hrs	5hrs	6hrs
Learnt eye input fields	345	308	299	294	282	279
Learnt eye links	546	557	575	586	592	601
Eye link usage range	$0-27$	$-2-28$	$-2-31$	$-2-35$	$-2-42$	$-2-56$
Learnt head input fields	0	48	74	83	91	97
Learnt head links	0	64	105	133	158	191
Head link usage range	$-$	$-1-20$	$-1-27$	$-1-44$	$-1-64$	$-1-80$

Figure 8, it is clear that only a subset of the fields are required to obtain a good performance.

6 Conclusion

We have compared the performance of a sequential learning strategy to that of a synchronous learning strategy in relation to the learning of sensorimotor gaze control on an iCub robot. While both approaches achieve the same final goal in terms of performance, the results presented have illustrated some of the differences that

arise due to the variation in the learning strategy. These variations appear in terms of the distribution of fields learnt, leading to differences in the number of fields learnt and the number of steps taken.

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 lighter coverage of the surrounding regions where less detail is required. On the other hand, the synchronous learning shows a more even mapping that covers a larger area without specialising

Table 4 Break down of the fields and links learnt at each stage during synchronous learning

Synchronous	1hr	2hrs	3hrs	4hrs	5hrs	6hrs
Learnt eye input fields	301	335	405	426	432	446
Learnt eye links	398	526	671	766	839	928
Eye link usage range	-5 − 9	-9 − 19	-11 − 20	-13 − 26	-15 − 30	-18 − 37
Learnt head input fields	18	38	66	80	96	108
Learnt head links	22	50	95	128	164	194
Head link usage range	0 − 11	0 − 32	-1 − 39	-1 − 45	-1 − 51	-1 − 56
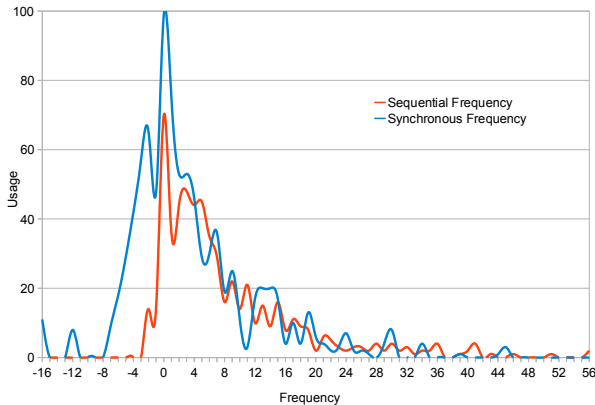
Fig. 9 Final link usage frequency plots for the learnt eye mappings

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. When only the eye is moving, the focus is on objects in a small area 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 and makes additional targets, that were previously out of visual range, reachable.

The two learning strategies are inspired by literature on infant development where two categories of constraints, A and B, can be identified as influencing the developmental trajectory. The sequential learning approach is used to simulate the type A constraints where a clearly defined sequence is given for the order in which the development proceeds. Meanwhile, the synchronous learning approach is used to simulate an aspect of type B constraints. The type B constraints refer to external impacts on the development, such as from the environment or carers within the environment, however they can also reflect internal support structures that are developed through experience and can be pursued through novelty and discovery within the environment. In the experiments presented here, the

environment does not change and the iCub robot is left on its own to learn. In this sense, the aspect of type B constraints that is evaluated is that of the internal support structures. It is anticipated that if every aspect of the sensorimotor system was allowed to start learning at the same time, a series of self imposed, emergent constraints would appear, whereby certain components would need to be learnt before other components could learn anything useful. The exact ordering could vary depending on motivations or nuances within the environment, leading to a more dynamic approach to learning, rather than the carefully sequenced learning of the type A constraints. While the learning of only two components has been considered here, a clear delay in the development of the head mappings has been seen. If this learning was extended to include torso and reaching, it is likely that in some cases the reaching may start learning before the torso, while in other cases it would be vice versa. The two types of constraints are likely to work in tandem, for example initial constraints from muscle tone will limit learning to just the eye, however as muscles develop, other components may be learnt more synchronously. In the experiments presented here, we only consider the learning of the eye and head for gaze control, however we use this phase of the learning to compare a sequential approach to learning to a synchronous approach, which could be applied to learning in general.

Regardless of whether the learning is performed sequentially or synchronously, both approaches produce mappings that are capable of combining links from the eye and head mappings to direct the gaze at desired targets. The rate of learning between the two approaches is very similar, with a similar number of fields learnt after a matching number of saccades. The performance based on link usage developed in the sequential approach suggests that the links learnt early on are more reliable than those learnt in the synchronous approach, and the initial high density of accurate links gives a strong basis for further mappings to be learnt. This early focus is also important to provide good coverage around the foveal region where the detailed visual processing will take place. However, it is also clear that

some synchronous learning is required to enable learning of saccades to the periphery.

References

Angulo-Kinzler, R. M., Ulrich, B., and Thelen, E. (2002). Three-month-old infants can select specific leg motor solutions. *Motor Control*, 6(1):52–68.

Asada, M., Hosoda, K., Kuniyoshi, Y., Ishiguro, H., Inui, T., Yoshikawa, Y., Ogino, M., and Yoshida, C. (2009). Cognitive developmental robotics: A survey. *IEEE Transactions on Autonomous Mental Development*, 1(1):12–34.

Freedman, E. (2008). Coordination of the eyes and head during visual orienting. *Experimental Brain Research*, 190(4):369–387.

Gallahue, D. L. (1982). *Understanding Motor Development in Children*. John Wiley, NY.

Hülse, M. and Lee, M. (2010). Adaptation of coupled sensorimotor mappings: An investigation towards developmental learning of humanoids. In Doncieux, S., editor, *SAB 2010, LNAI 6226*, pages 468–477. Springer.

Law, J., Lee, M., and Huelse, M. (2010). Infant development sequences for shaping sensorimotor learning in humanoid robots. In *Proc. 10th Int. Conf. on Epigenetic Robotics*, pages 65–72.

Law, J., Lee, M., Hülse, M., and Tomassetti, A. (2011). The infant development timeline and its application to robot shaping. *Adaptive Behaviour*, 19:335–358.

Law, J., Shaw, P., and Lee, M. (2013). A biologically constrained architecture for developmental learning of eye-head gaze control on a humanoid robot. *Autonomous Robots*, 35(1):77–92.

Lee, M. H., Meng, Q., and Chao, F. (2007). Staged competence learning in developmental robotics. *Adaptive Behaviour*, 15(3):241–255.

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

Needham, A., Barrett, T., and Peterman, K. (2002). 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*, 25(3):279–295.

Pfeifer, R. and Scheier, C. (1997). Sensory-motor coordination: the metaphor and beyond. *Robotics and Autonomous Systems*, 20(2):157–178.

Piaget, J. (1973). *The Child's Conception of the World*. Paladin, London.

Shaw, P., Law, J., and Lee, M. (2012). An evaluation of environmental contraints for biologically constrained development of gaze control on an icub robot. *PALADYN Journal of Behavioral Robotics*, 3(3):147–155.

Smith, L. and Gasser, M. (2005). The development of embodied cognition: Six lessons from babies. *Artificial Life*, 11(1-2):13–29.

Wang, X. and Jin, J. (2001). A quantitive analysis for decomposing visual signal of the gaze displacement. In Feng, D. D., Jin, J., Eades, P., and Yan, H., editors, *Proceedings of the Pan-Sydney area workshop on Visual information processing - Volume 11 (VIP '01)*, volume 11, pages 153–159, Darlinghurst, Australia. Australian Computer Society, Inc.