Constraint Facilitated Learning in Artificial Systems

Frank Chien and Wystan N. Palmer
Swarthmore College Department of Computer Science
May, 2010

The goal of our study was to determine whether the implementation of sensory and motor constraints on a developing robot would significantly improve learning progress. Developing robots were initiated in a rectangular hallway facing a wall whose color was dependent on the robot's movement and position. Throughout the developmental process, the robot had to predict future sensorimotor vectors based on past experiences and report the error in this prediction. Experiments were run on both constrained and unconstrained robots. A comparison was made between the prediction successes in the final stages of development in both the constrained and unconstrained robots. We found no significant evidence showing that the addition of constraints led to lower prediction error. We also found no significant evidence of decrease in prediction error in the unconstrained agents. The data collected led us to believe that neither the constrained nor unconstrained experimental groups were successful in learning to predict their environment and that the designed world was too difficult for the robot to learn with the given parameters and developmental architecture.

1. Introduction

Human infants are born with poor mastery over their muscular-skeletal locomotive functions. Further, physical maturation and the emergence of fine coordinated muscular control follows a developmental pathway shared among all infants. For example, rolling precedes crawling, which in turn precedes standing and walking. In developmental biology and robotics literature, the freezing and release of degrees of freedom in movement has been speculated as a phenomenon that contributes to this developmental pathway. It has been suggested that freezing and freeing degrees of freedom facilitates development by reducing a large range of potential movements that may overwhelm an infant. A baby learning to coordinate its body has heavy coupling in many joints, which effectively constrains the degree of freedom. These couplings are then relaxed as the infant masters its rudimentary locomotive action sets and explores more complex movements which require increased joint freedom. (Lungarella et. al, 2003).

Developmental robotics experiments have been run to test the effect of constraining degrees of freedom within a developing robot. Lungarella and Berthouze experimented with the freezing and freeing of degrees of freedom in a robot learning to swing in situations with environmental disruptions (2002). They found that the steady unfreezing of joints allowed the robot to exhibit swinging behavior, but it was necessary to refreeze and unfreeze degrees of freedom when environmental disruptions effected the swinging oscillations.

Less obvious but of equal importance to constraints in movement present in infants is the fact that an infant's sensory organs are less acute than that of an adult. Three-month-old infants see in very poor resolution, and are roughly fifty-times less sensitive to contrast (Brown, 2009). Lungarella explains that these sensory inhibitions are advantageous for the infant, because abridging extensive sensory information prevents a premature brain from being overwhelmed by sensory overload. It naturally follows from these observations that a developing robotic agent may benefit from visual constraints as well; the principles applicable to the development of humans may be transferable to artificial learning systems. Our study aimed to address this question by imposing similar constraints on a simulated robot and contrasting its learning success with that of an unconstrained control robot.

Artificial learning was modeled using Category Based Intrinsic Motivation (CBIM) (Lee et. al, 2009). Both constrained and unconstrained learning robots attempted to predict the coloring of a wall, which operated on a rule-based system dependent on the robot's movement and position. Should constraint-facilitated learning be transferable to artificial systems, we would expect the agent developing under constraints to learn to predict its environment more successfully, and consequently, show less prediction error than a robot developing free from constraints.

2. Methods

2.1 Developmental Architecture

Artificial development in our experiments was implemented using CBIM. CBIM is a hybrid method that combines Intelligent Adaptive Curiosity (IAC) (Oudeyer et al, 2007) and Growing Neural Gas (GNG) (Fritzke, 1995) to produce a system in which a developing agent searches situations that will maximize its learning process. IAC acts by storing all situations previously experienced by the robot in a model vector memory. These vectors are split into regions of similar sensorimotor space and contain experts trained to predict outcomes of actions taken based on current sensory data. At each time step, IAC takes in a sensory vector and makes predictions for the sensory vector of the next time step based on generated candidate actions. On the next time step, an action is chosen and error in prediction is calculated and stored in the expert. These stored errors are used to evaluate potential learning progress in choosing future actions. Actions with the highest determined learning progress are chosen and carried out. An additional parameter is used to determine when a random action is chosen instead. In this way, the system can be considered curious but still exhibit random behavior that can lead to the discovery of new situations. CBIM implements IAC using a GNG model to categorize vectors into regions. GNG is a topology learning algorithm that consists of connected units containing vectors. In CBIM, each GNG unit represents a region of IAC. These regions are connected by edges. At each time step, the GNG distribution is adjusted based on Euclidean distance between a sensorimotor vector from the environment and the nearest GNG unit. When error within the GNG units exceeds a given threshold, new units are added to accommodate the sensory data that does not fit the current distribution. Implementation of a GNG error threshold ensures that IAC regions are only added when the robot's environment presents significantly new situations. Otherwise, the GNG regions are dynamic enough to adapt to and learn slight variances in the sensorimotor space. In our experiment, the GNG error threshold was set to 0.60 and the probability of choosing a random action to carry out on the following time step was 15%. At each time step, the potential learning progress is evaluated for 20 candidate actions per possible color focus output.

2.2 Data Analysis

Over the course of our experiments, we collected data on prediction error and color focus. The color focus data were broken into fifty step epochs in which the percentage of time the robot chose to focus on a color over the course of the epoch was recorded. Percentage of time when the focus color matched the wall color was also recorded. The correlation between these two percentages were then compared over the course of the development. Developmental regions in which matched wall color percentage fell close to the color focus potential corresponded to periods when the robot was correctly predicting the color it would see in the next step. Regions where the percentages differed suggested incorrect color predictions.

In CBIM, the prediction error is calculated as the Euclidean distance between a vector prediction of the sensory data of the next step in development and the actual sensory data collected at that next step. At each time step, a list of candidate actions is generated and evaluated to determine learning progress. The majority of the time, the action with the highest learning progress is chosen. From this action, a sensorimotor vector is generated and passed into the GNG region that is closest in Euclidean distance. The exemplar in this GNG unit makes a prediction of the resulting sensory space. Once this prediction is made, the chosen action is taken and Euclidean error is calculated between the prediction vector and the resulting sensory vector. In our experiment, these prediction errors were averaged over 125 step epochs throughout the development. Graphical analysis was run on these average errors in an attempt to observe trends in the data.

Qualitative analysis was also run on the GNG models created by both the control group developments and the experimental group developments. The number of GNG units created was compared as well as their general fit to the sensorimotor data over the course of each experiment.

2.3 Developmental World

The simulations were run in an elongated hallway (figure 1). All walls in the world were black except for the wall directly in front of the developing robot. The color of this wall varied depending on the robot's speed, direction, and location. The developing robot was initiated halfway down the hallway centered between the two side walls and facing the colored wall. In order for the robot to continue facing the wall throughout the experimental run, its movement was limited to forward and backward translation. At each time step, the wall color changed based on the robot's movement and position. When the robot was within a distance of two units of the wall, the wall color changed to red. Outside of this distance, the color was only dependent on the robot's speed and direction. If the robot was moving away from the wall, the color was blue. At slow forward speeds in the range [0.0, 0.33], the color changed to cyan. With translation outputs in the range [.34, .66], the wall color was yellow. Forward speeds within the range [0.67, 1.0] caused the wall to become green. In total, there were five possible wall colors accessible to the robot. These colors and the robot behavior triggering them are summarized in Table 1. If at any time the robot stalled against either the colored wall or the wall opposite it, the robot would be reset at a random point in the hallway centered between the side walls.

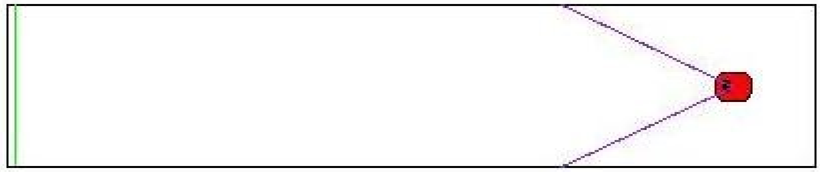

Figure 1. Simulation Environment

	Within 2	Motor output	Direction
1	Yes	-	-
2	No	-	Backwards
3	No	[0 - .33]	Forwards
4	No	[.33-.66]	Forwards
5	No	[.66- 1]	Forwards

Color
Red
Blue
Cyan
Yellow
Green

Table 1. Summary of wall coloring conditions

2.4 Robot Sensorimotor Space

The developing robot was fitted with the Pyrobot camera device. Color match filters were implemented to detect the presence or absence of each possible wall color. The sensorimotor space of the robot in our experiments was designed based on that of the original CBIM experiment. The developing agent was given six inputs and two outputs.

The two outputs corresponded to values of color focus and translational movement. Movement was scaled to a value between 0.0 and 1.0, where values less than 0.50 corresponded to backward movement and values greater than corresponded to forward movement. The color focus output value was implemented in the original CBIM experiment and kept for our experiments. It provided a way for the robot to predict the presence and size of objects of specific colors within its field of vision. In our experiment, it was used to focus on and predict the presence of particular colors during development. Ideally, the robot would choose to focus on colors that were easier to predict earlier in development and move on to focus on harder colors later in development. The color focus output was a value in the range [0.0, 1.0]. This range was split into five equal intervals corresponding to possible colors: [0.0, 0.20] corresponded to red, [0.2, 0.4] to yellow, [0.4, 0.6] to green, [0.6, 0.8] to cyan, and [0.8, 1.0] to blue.

The first five inputs corresponded to the five possible wall colors that the robot was able to see. They were entered in the order [blue, cyan, red, yellow, green]. These inputs were entered as Boolean values; a value of 0.0 meant that the color was not seen and a value of 1.0 meant that the color was seen in the camera view. The sixth input value corresponded the area of the color being seen. This area was given as input only if the color present matched the robot's color focus. Otherwise, the input value was 0.0. As a result, the robot would only be given a wall area if it was focusing on the correct color. This gave the robot the ability to autonomously differentiate between situations where it was predicting the correct wall color and situations where its predictions were incorrect. The overall sensorimotor vector contained eight values:

SM(t) = {blue, cyan, red, yellow, green, wallArea, colorFocus, translation}

2.5 Experimental Process

Both the control and the experimental agent were run for 10,000 time steps. During the last 2,500 steps of development, neither the control nor the experimental agent were constrained. Prediction error was compared only during this period to determine which had better learned to predict the environment. The main developmental period occurred during the first 7,500 steps of each experiment. The control experiments were run with no constraints on sensors or motors. The robot was capable of experiencing all five possible wall colors and able to translate at any speed and direction for the entire development.

The first 7,500 steps of the experimental agent's development was subdivided simultaneously into three motor-restrictive stages and five visual-restrictive stages. In order to eliminate confusion between the robot's chosen movement output and its constrained output, motor constraints were implemented by modifying the candidate actions before they were fed into the regional GNG vectors for determination of learning progress. Since the regions only received candidate actions that were already constrained, the robot was able to learn based on already scaled sensorimotor vectors at each stage of the development. The constraints on motors were divided into three stages of 2,500 steps each. In the first stage, the robot could only move one-third of its maximum speed in the forward or backward direction. In the second stage this constraint was eased to allow the robot to move two-thirds of maximum speed. The constraint was fully released in the third stage.

Colors were distinguished using a combination of the pyrobot camera's match and blobify filters. If the robot was in a stage of development where it was able to see a particular color, its camera would be modified with a match filter that took in RGB values specific to that color and passed any instances of the color within the camera's view into the blue color channel. This channel was picked up by a blue blobify filter in order to determine the area of the wall. To ensure that a blue wall color is not incorrectly identified after previous match filters are implemented, the blue match sensor was the first color sensor passed. The vision constraints were divided into five stages of 1,500 steps each. In the first visual constraint stage, the robot could only distinguish the color blue. If the wall was any other color, the robot was unable to see any color and each color sensor read a value of 0.0. In the second stage, the robot gained the ability to see cyan. In the next three successive stages, it gained the ability to distinguish the colors red, yellow, and green until it reached the final stage in development where vision constraints were effectively released. A summary of the stages of constraints can be seen in Table 2.

Motor Development Stage	Motor output	Visual Development Stage	Additional Color Seen
1	[-0.33, 0.33]	1	Blue
		2	Cyan
2	[-0.66, 0.66]	3	Red
3	[-1, 1]	4	Yellow
		5	Green

Table 2. Summary of visual and motor constraints

3. Results

3.1 Trends in Prediction Error

Prediction error data was averaged over intervals of 125 steps throughout the experiments. Analysis of the prediction error data revealed a few general trends in the robot's ability to predict its environment in both the constrained and unconstrained development. In both the control and experimental developments, prediction error in the final steps of development was lower than prediction error in the first several steps in development. This occurred only as a result of the addition of more GNG units to fit the sensorimotor space of the world. Despite the ultimate decrease in prediction error between beginning and end of development, there were no outstanding trends seen in the change in prediction error over the course of development for either the constrained robot or the unconstrained.

In each unconstrained experimental trial, The prediction error consistently dropped within the first 2,000 steps of development to a value between 0.30 and 0.50 (figure 2). After this period of sharp decrease, the prediction error oscillated between these two values for the remainder of the experiment, showing no further overall decrease in error.

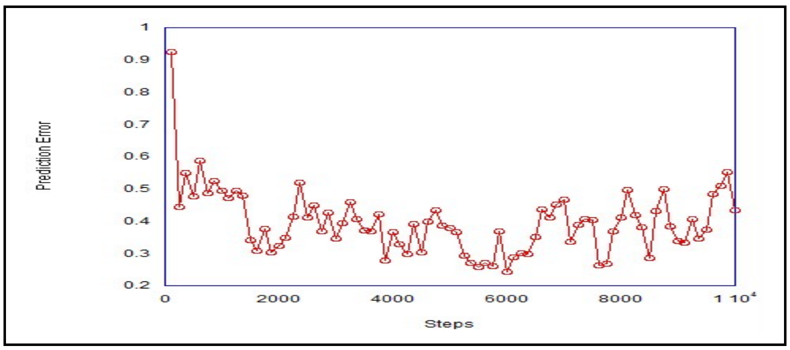

Figure 2. Prediction errors for an unconstrained development

The periods of constraints for the trials run on the experimental agent all showed similar prediction error trends (figure 3). Within the first 1,500 steps of development, The prediction error showed a sharp decline from between 0.60 and 0.90 to between 0.20 and 0.40. Once the development hit step 1,500, a sharp peak in prediction error was observed followed by another general decline. The peak at step 1,500 was seen in every constrained development run and was the most noticeable change in prediction error throughout each development. Similar but less conspicuous peaks were observed in each constrained trial at time steps in which either a motor or a vision constraint was released. These peaks were generally less noticeable and, like the peak at step 1,500, followed by overall decreases in error. As with the control group, the experimental agent development trials all showed a lower prediction error between the end and the beginning of the 10,000 step developments. In the final unconstrained 2,500 steps of development, the experimental trials showed an oscillatory trend similar to that of the control group developments. Prediction error varied between the values of 0.10 and 0.50 with no significant decreasing trend.

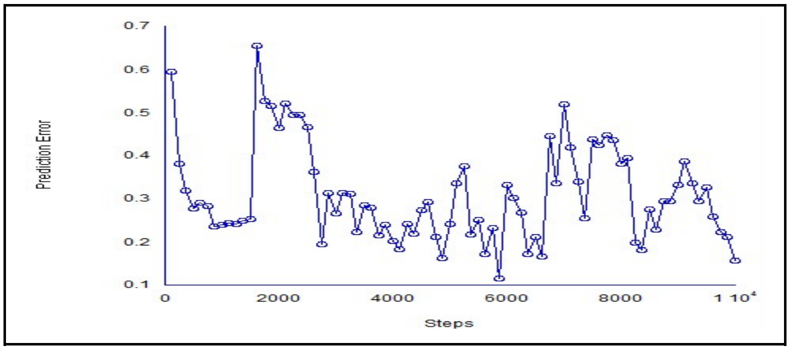

Figure 3. prediction errors for a constrained development

3.2 Statistical Analysis of Prediction Error

A linear regression was fitted to the prediction error of the control agent. Analysis of the regression slope showed no evidence of a non-zero slope. The linear fit residual plot showed a pattern of decreasing residuals to around step 6000, and then consistently increases until the end of the trial. Quadratic fit residuals show a much more even spread (figure 4).

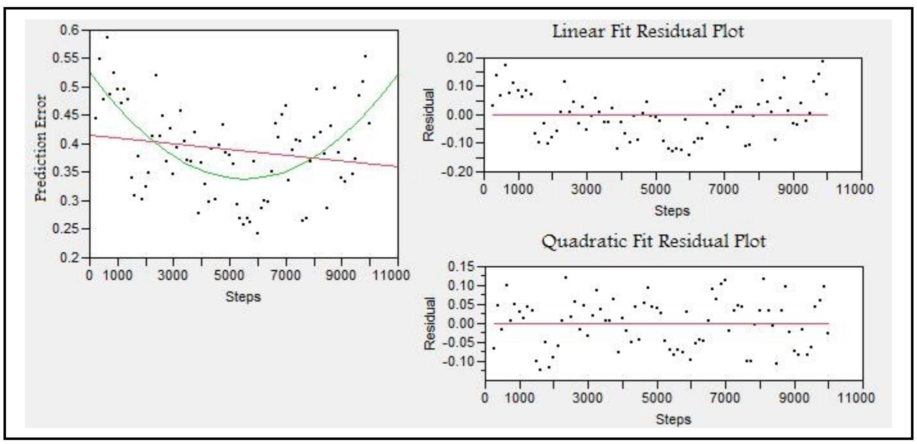

Figure 4. Linear and quadratic fit for prediction error in unconstrained development.

3.3 Comparison of Control and Experimental Groups

In the final 2,500 steps of the experiment, developmental trials in both the control and experimental group showed oscillations in prediction error between the values of 0.10 and 0.60. Trials in the experimental group tended to oscillate closer to the lower boundary of this range while trials in the control group showed peaks approaching the upper bound of this range. There were no consistent trends in prediction error between trials of either the experimental or the control group (figure 5). Some experimental trials performed very well and showed relatively low prediction errors. Others performed poorly and showed increases in prediction error even in the last 500 steps of their development. Similar trends were seen in the control group trials. Through graphical analysis, it was revealed that prediction error in the experimental trials was not consistently lower than prediction error in control group trials; there were regions where the control groups seemed to exhibit lower values of prediction error than the experimental groups.

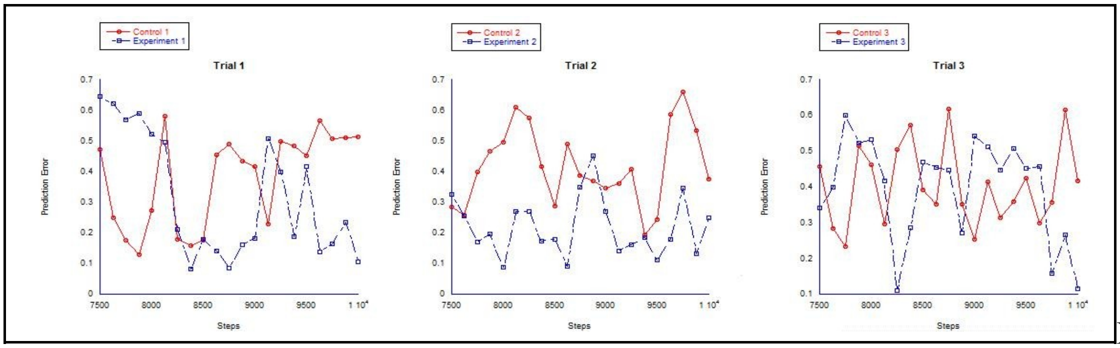

Figure 5. Comparison of prediction error over the last 2500 steps of experiment between constrained and unconstrained development.

3.4 GNG Models

Comparison of the GNG models made by the control and experimental developments showed both similarities and discrepancies. Models in both groups fit the experienced sensorimotor regions well. There were usually either five or six distinguishable regions present. Each region was well represented by a group of connected GNG units, but units between regions were generally not connected. There were exceptions in this trend: in some GNG models, a few units were dispersed randomly outside of experienced sensorimotor regions and some defined sensorimotor regions exhibited small amounts of edges connecting units between one another.

A discrepancy was seen in the amount of GNG units created in the constrained experimental developments versus the unconstrained control developments. Robots in the experimental group consistently produced less GNG regions over the course of their lifetime than robots in the control group. Robots in the experimental group generally ended their development with between 22 and 24 GNG units, while robots in the control group generally ended development with between 26 and 35 GNG units. Despite having less units, however, the experimental GNG models were still sufficient in representing the experienced sensorimotor space in the final 2,500 steps of the development.

3.5 Color Focus

Analysis of the color focus data collected in both groups of experiments showed similar trends in the developing robot's ability to predict and focus on the correct wall color. In the unconstrained agents, the colors cyan, yellow, and green were rarely both focused on and seen. By graphical analysis, it was shown that over the course of development, the percentage of time in which these colors were correctly predicted was often below 5% over a fifty step epoch and rarely exceeded 15% (Figure 6). Further exploration of the focus data reveals that the percentage of time in which the robot chose to focus on these colors usually greatly exceeded the percentage of time in which the robot actually saw these colors.

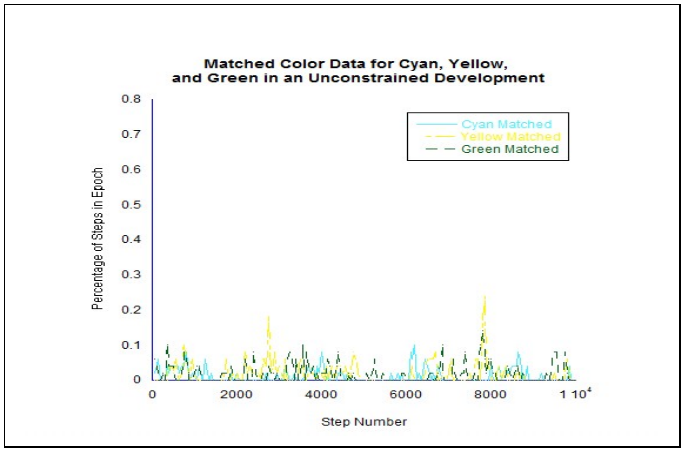

Figure 6. example of color match data for cyan, yellow, and green during an unconstrained development experiment.

The unconstrained robot showed higher prediction successes for the colors red and blue. Analysis of the blue color focus data showed regions in which the color focus matched the wall color between 40% and 65% of the time within a fifty step epoch. Graphical comparison of the percentage of time the robot chose to focus on blue and the percentage of time blue was actually seen showed a close correlation between blue focus and blue presence. Analysis of red color focus data revealed the same trends, but the trends occurred sporadically throughout different developments; different spans of time in each development showed correct matching between focus on the color red and the presence of red, and the length of these spans varied between trials.

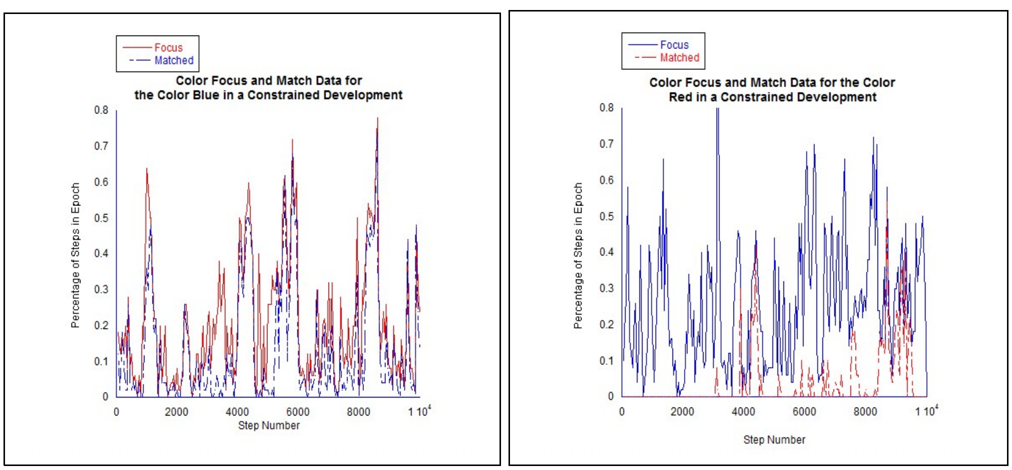

Figure 7. Color Focus and Match Data for Red and Blue in a Constrained Development

In the trials run on the experimental agent, similar red and blue color focus data was seen (figure 7). In addition, the same trends were seen in the color match rates of the colors yellow and green. Cyan showed the same trends for the majority of the development with the exception of a brief period of time. Between the time steps of 1,500 and 3,000 in each constrained development, there is an increase in correct color matching and prediction data for the color cyan. The percentage of time in which cyan focus and presence was correctly matched slightly increased to a consistent value above 5% over a fifty step epoch, and the percentage exceeded 15% at least twice in each trial between these time steps (figure 8).

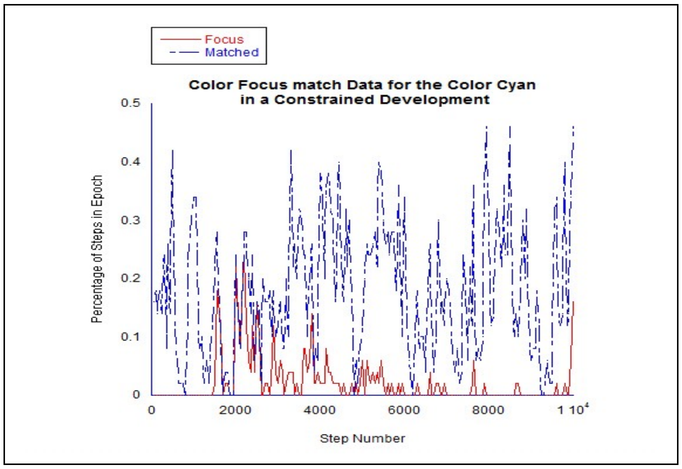

Figure 8. example of color match data for cyan in a constrained development.

4. Discussion

4.1 Prediction Error and Experimental Comparison

In their developmental robotics survey, Lungarella et. al. state that constraints are advantageous to a developing agent because they reduce the complexity of the sensorimotor space for immature systems and prevent sensory overload (2003). In our experiments, we hypothesized that the implementation of visual and motor constraints on a developing robot would allow it to exhibit more successful learning progress than a robot developed without constraints. Our results were inconsistent with our expectations. In one of the three trials, the constrained robot consistently outperformed the unconstrained (figure 5). However, for the other two trials, there is no clear evidence that the constrained robots out-performed the unconstrained robots. Furthermore, prediction error collected from both control and experimental groups showed high variance in all three trials. From this data it was concluded that there was no significant difference in learning between the unconstrained and constrained robots.

To perform statistical analysis on prediction error, we plotted error against step range in the unconstrained robot (figure 4). A linear regression on this data showed no statistically significant evidence of a non-zero slope. That is, there was no evidence of a decrease in prediction error over the course of the experiment. No regressions were calculated on the experimental group's development data because the constrained robot's sensorimotor capabilities were not held constant throughout the trial.

The residual plot for the linear fit on the unconstrained data demonstrated a downward trend approaching a minimum around step 6,000. From here, the residual plot showed an upward trend. This pattern suggested that a linear model is not the best approximation of the data distribution. The quadratic fit residual plot showed an even spread. The data showed that as the development progressed, prediction error would initially decrease, but then after a certain stage in development the pattern would reverse and prediction error would increase. It is possible that our data was confounded by the curiosity mechanism built into the CBIM system. Because robots implemented with the CBIM model are intrinsically curious, situations that are easier to predict become boring to the system. As this occurs, the robot begins to choose candidate actions that are harder to predict. That is, The robot chooses actions that have higher prediction errors. If this is the case, The non-decreasing trend in our prediction error data was an artifact of CBIM's intrinsic motivation.

4.2 GNG Distributions

The GNG units in both the control and experimental trials were able to separate into distinct regions in sensorimotor space as a result of the sufficiently large Euclidean distances between vectors containing boolean values for the presence or absence of colors. Once the GNG units were spread into these regions, the edges connecting units between different regions were mostly removed through aging. The few that remained at the end of the experiment were the result of insufficient aging; since edge age is dependent on visits to the GNG units connected by the age, ages were not updated unless these units were nearest to the sensorimotor vector at a given time step. Edges between regions were the result of few visits to the units connected by these edges. Despite the presence of lingering edges, however, the GNG distribution fit the distribution of experienced sensory vectors in both the constrained and unconstrained developments. Each cluster of GNG units in the distribution roughly corresponded to a specific color seen in the environment. Figure 9 shows how the colors correspond to the sensory regions.

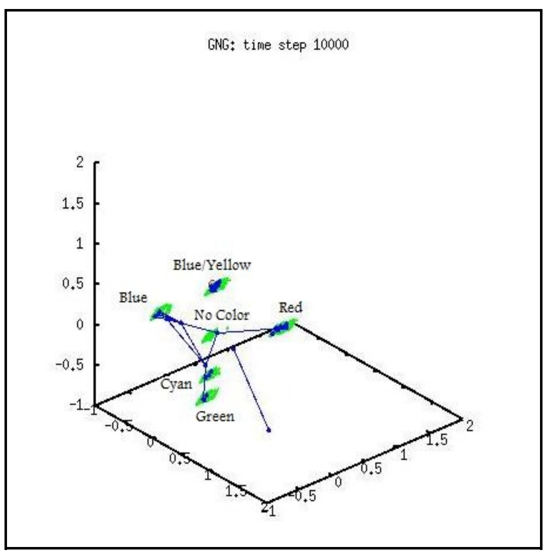

Figure 9. GNG model for an example constrained development. Sensory regions are labeled with colors active in their corresponding vectors.

The discrepancy seen between the number of GNG units created in the control experiments and the number created in the constrained experiments is related to the relative abundance of sensory states available to the developing systems early development. Since GNG units were only made when the error threshold was exceeded, new units were only added when new sensory states were discovered. In the constrained trials, GNG units were added at time steps when constraints were released and new colors were seen. Since these instances were spread over time, the GNG model had time to fit to sensorimotor data before new experiences occurred.

In the unconstrained developments, the robot started with the ability to see every color and move at full speed. Because of this abundance of sensory information, all of the GNG units in these developments were created within the first 1,000 steps. Since there was no time for the distribution of increasing numbers of units to fit to the data before error threshold was again exceeded, more GNG units were made early these trials. The constrained developments produced GNG distributions that fit the sensorimotor space with less units because their sensory space was slowly incremented. The large amount of sensory data available to the control group throughout development led to the production of larger numbers of GNG units to deal with the initial sensory overload.

Aside from the differences in the number of GNG units created to fit the sensorimotor space of the environment, there was no significant difference observed between the models made during the constrained and unconstrained developments. Models for both experiment groups showed similar GNG distributions fitting the sensory states experienced throughout development. Since the final GNG models for both unconstrained and constrained developments were so similar, it can be concluded that there was no difference in the way the two types of developments categorized their sensorimotor space at the end of their lifetime.

4.3 Color Focus

The trends in the color focus data resulted from a combination of the artificial curiosity built into CBIM and the complexity of the world in which the robot was developed. In both the control and experimental trials, the color blue was consistently focused on and correctly predicted to be within camera’s field of vision. The color blue was predicted frequently because it was the easiest color to predict. When the robot was outside of two units from the wall, the wall changed to blue with any backward movement regardless of speed.

The next easiest color to predict was the color red; when the robot was within a two units of the wall, the wall remained red regardless of movement. As a result, prediction of the color red was almost entirely dependent on the wall area input. In both the control group and the experimental group, the color red showed signs of predictability. Because the wall only turned red when the robot was within a small specific region of the world, the robot was only able to learn situations with the color red during short periods in its development. This is why correct prediction of the color red was only sporadically seen in the developments. The time steps in which the robot entered this region during development differed between trials. The resulting color focus data showed varying spans of time in which the robot was able to see and correctly predict the color red corresponding to the brief spans of time in which the robot was within the region close to the wall and able to explore this region.

The remaining three colors were harder for CBIM to predict because these colors were dependent on both direction and speed. Because these colors were harder to predict, they showed significantly lower color match percentages in the color focus data. The only exception was the match percentage increase in the color cyan between step 1,500 and 3,000 in the constrained developments. Because the robot's movement was constrained to one-third maximum speed in either direction, the only two wall colors accessible to the robot between these time steps were blue and cyan. As a result, cyan became as easy to predict as blue within this brief time period. This is shown by the higher percentages in the matched cyan focus data in this stage.

The absence of clear transitions between successful color focus predictions that were seen in the original CBIM paper is a result of the complexity of our world. Instead of three colors, our robot had to choose from five colors to focus on. The size of the robot's sensorimotor space may also be a factor in the difficulty of color prediction. The number of boolean values within the vector may have made it hard to effectively train GNG exemplars to correctly predict wall color. Because of the difficult sensorimotor vectors, the CBIM brain may have found it harder to categorize situations than it would have given a smaller sensorimotor space.

5. Conclusion

From the data we gathered from our experiments, we were unable to provide conclusive evidence that constraint-facilitated development is more conducive to learning in artificial systems than unconstrained development. If a developing system were learning about its environment, it would be expected to be able to better predict its environment at the end of its developmental phase. From the prediction error data collected in both control and constrained developments, we were unable to see a significant decreasing trend in prediction error over the course of the robot's lifetime. From this data we can conclude that neither the constrained nor the unconstrained robot's showed significant learning progress.

Apart from a difference in the number of units created and the time at which units were added, the GNG models of the unconstrained and constrained developments showed no outstanding differences. Both groups of experiments showed that the sensorimotor space in the developmental world could be well represented with between 20 and 35 GNG unit exemplars. Two conclusions can be drawn from the lack of difference in GNG architecture between the two sets of developments. The first is that a development exhibiting slow incrementation of sensorimotor capabilities resulted in the creation of a smaller number of representative units than a development in which all capabilities are always present. Secondly, because GNG units are able to adapt and move to fit changes in sensorimotor space, the presence of a stage in development of both constrained and unconstrained robots where each have the same capabilities will result in a similar GNG distribution to fit similar sensory environments.

Color focus data from both sets of experiments suggest that neither constrained nor unconstrained robots tended to learn to predict wall color more successfully than the other. It can be concluded that the colors blue and red were both easy to predict within our experimental world. Cyan, yellow, and green were harder to predict for both the control group robots and the experimental robots. The relatively large sensorimotor vector available to our robots may have been at fault. The difficulty may also have been a result of the unforeseen complexity of the learning task.

Each set of data analyzed in this paper shows no conclusive evidence for a difference between the developmental progress of a constrained robot compared to that of an unconstrained robot. In fact, the data suggests that neither constrained nor unconstrained robots showed learning progress over the course of development. Since CBIM has already been shown to work as an intrinsically motivated developmental system, we believe that our results were caused by the complexity of the task. An extended analysis of the effect of variation in parameters such as GNG error threshold may produce more effective results. It may also be effective to reduce the amount of colors possible in the environment or to reduce the sensorimotor vector to include only three color inputs corresponding to combinations of red, green, and blue color values.

References:

Brown and Lindsey, 2009, Contrast Insensitivity: the critical immaturity in the infant visual performance. *Optom Vis Sci. 572-6*

Fritzke, B. (1995). A growing neural gas network learns topologies. In Tesauro, G., Touretzky, D. S., and Leen, T. K., (Eds.), Advances in Neural Information Processing Systems 7, pages 625–632. MIT Press.

Lee, R., Walker, R., Meeden, L., Marshall, J., Category-Based Intrinsic Motivation, *Proceedings of the Ninth International Conference on Epigenetic Robotics,* 2009

Lungarella, M., and Berthouze, L., 2002b, Adaptivity via alternating freeing and freezing of degrees of freedom. *In Proceedings of the 9th International Conference on Neural Information Processing (Singapore),* pp. 492-497

Lungarella, M., Metta, G., Pfeifer, R., Sandini, G., 2003, Developmental Robotics, A Survey. *Connection Science,* 15: 151-190

Oudeyer, P.-Y., Kaplan, F., and Hafner, V. 2007. Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation,* 11(2):265-286.