Transfer or no transfer: The key role of learning specificity

Paul Rauwolf and Luc Berthouze
Centre for Computational Neuroscience and Robotics
University of Sussex
Falmer BN1 9QH, UK
{P.Rauwolf,L.Berthouze}@sussex.ac.uk

Abstract

During development, infants garner knowledge about their environment and their own body through explorative behaviours mediated by changing motilities. A key question regards infants ability to transfer knowledge between these various motilities. There is currently little consensus in the literature. In depth perception studies, for example, Adolph showed that knowledge amassed while sitting does not appear to transfer to crawling, and that information collected during crawling does not transfer to walking. In contrast, Witherington and colleagues reported findings suggesting that there is transfer from crawling to walking. Here, we attempt to reconcile the two findings by suggesting, with a simulation, that a key experimental difference could explain the disparity between these studies. The results suggest that an infant lacking transfer between motilities, as suggested by Adolph, can behave in ways consistent with the findings of Witherington and colleagues, and an empirical prediction is derived that could be tested in real experiments. This study highlights the potential importance of learning specificity in development, a concept that could have important implications for developmental robotics.

1. Introduction

Most developmental robotics studies make the more or less explicit premise that knowledge or sensorimotor competences acquired during development should transfer to the next set of competences. This is a central component of the notion of *on-going emergence* (Prince et al., 2005) and of numerous hierarchical schemes (see Vigorito and Barto, 2010, for example). From a developmental robotics viewpoint, whilst presenting an important technical challenge – namely, that of designing a proper framework – this is a useful if not intuitive way of thinking about the problem. Surely, the ability to embed everything that has been learned whilst avoiding the creation of too much redundant knowledge, should help guide and stabilise development.

The developmental psychology literature, however, is not that clear-cut on the topic. Whilst some authors explicitly claim that there is transfer (see Witherington et al., 2005, for example), others claim there is not. Rather than taking this as a potential problem for developmental robotics, we see this as a new avenue for considering this problem. If there is no transfer in the way we typically envisage transfer yet the system is robust, then what exactly keeps the system stable? The helpful contribution of environmental (including other agents, e.g., parental scaffolding) structure notwithstanding, is it possible that it is not knowledge or the sensorimotor couplings themselves that get transferred but the way the system acquire them that get transferred? Or alternatively, could lack of transfer benefit learning?

In an attempt to elucidate this question, we explore two sets of studies from the developmental psychology literature in which authors use seemingly compatible experimental setups to investigate the presence or lack of transfer during development, yet arrive at radically different interpretations.

2. Analysis

2.1 Veritable slope/gap findings

In a series of papers, Adolph exposed a discontinuity in how infants interact with potentially dangerous environments, given experience in a particular motility. She showed that knowledge amassed while sitting does not appear to transfer to crawling (Adolph, 2000), and that information collected during crawling does not transfer to walking (Adolph, 1997). Both experiments will be discussed; however, the analysis will focus on the juxtaposition between crawling and walking, as those are the motilities assessed by Witherington and colleagues (see Section

Table 1: Go-Ratio for crawlers and walkers relative to the boundary angle. Note that these ratios have been extrapolated from a figure from Adolph (1997, p.65).

Angle with respect to boundary condition	1-week crawler	10-week crawler	22-week crawler	1-week walker	10-week walker
+0	1.00	0.92	0.90	0.90	0.85
+5	0.75	0.60	0.20	0.75	0.42
+13	0.55	0.50	0.00	0.50	0.45
+18	0.50	0.45	0.00	0.40	0.20

2.2).

The foundation of Adolph's experimentation lies in the fact that infants learn motilities in a temporal order. Thus, an infant will often prove highly experienced in one motility while inexperienced in another. This allows for an analysis of whether past knowledge is implemented when exploring via the new motility. Hence, Adolph compared experienced sitters to novice crawlers and experienced crawlers to novice walkers, following the typical developmental order. Adjustable cliff-like gaps were used to test when a sitter or crawler would attempt to cross the gap to reach a desired toy and when an infant would refrain. She discovered that there is a higher propensity for experienced sitters to abstain from trying to stretch across large gaps compared to inexperienced crawlers (Adolph, 2000). Additionally, Adolph examined experienced crawlers and novice walkers on declining slopes with alterable grades. By modifying the angle of descent, she analyzed when and if an infant would refuse to descend. Again, she discovered that information acquired during one motility did not appear to transfer to the next. The more experienced crawlers often avoided treacherous declines, while the novice walkers often sprinted into danger. For each infant, a boundary angle was defined as the angle that separated risky from trivial descents, or more formally, the angle at which a child would successfully navigate the ramp 67% of the times he/she was afforded an opportunity (Adolph, 1995, p.737). Once the boundary angle was determined, the infant was tested on declines relative to his/her skill. Table 1 approximates the ratio of occasions a child would descend a slope (the go-ratio), relative to the infants boundary angle. Note that the boundary condition is characterized by how successfully the infant performs:

$$\frac{Successes}{Successes + Failures + Refusals} \tag{1}$$

whilst the go-ratio denotes how often an infant will endeavour to navigate the ramp:

$$\frac{Successes + Failures}{Successes + Failures + Refusals} \tag{2}$$

Table 1 demonstrates that an experienced crawler is much less likely to descend a risky ramp compared to the novice walker. As a 22-week crawler is approximately the same age as a 1-week walker, age should not contribute to this disparity. Furthermore, the ratios between a novice crawler and novice walker are almost identical, though a 10-week walker behaves a bit more apprehensively compared to a 10-week crawler. The data appears to point to a lack of transfer. A novice walker does not seem to utilize the knowledge gleaned by progressing from a novice crawler to an experienced crawler. Rather, the novice walker reacts to dangerous slopes similarly to the way he/she responded as a novice crawler.

2.2 Visual cliff findings

In contrast to Adolph who used veritable cliffs and slopes, Witherington et al. (2005) analyzed the decisions of experienced crawlers and novice walkers on a visual cliff. Each child was placed on a platform and coaxed to cross what visually appeared as a pit. However, the pit was covered by a transparent surface so it was safe for the infant to cross. The lighting of the pit was managed in such a way that the infant was, theoretically, incapable of visually observing the transparent surface. While the experiment implemented both a shallow and a deep visual cliff, only the results from the deep visual cliff will be analyzed as the depth more closely approximates Adolph's descents.

Table 2: Ratios of children who crossed the deep visual cliff.

Motility type	Go-Ratio
Experienced crawler	0.35
Novice walker	0.10

As shown by Table 2, the results of this experiment are in contrast to those of Adolph with novice walkers behaving with higher trepidation. Only 10% of the walkers were willing to cross the visual cliff, while the crawlers bridged the cliff 35% of the afforded opportunities, leading the authors to state:

> In conclusion, we found clear evidence for avoidance of heights on the visual cliff in our sample of newly walking infants, who refused to cross onto the deep side of the cliff even more consistently than similarly aged prewalkers with extensive crawling experience.

2.3 Comparison of the two experiments

Two disparities between Adolph's and Witherington and colleagues' experiments require discussion. The first relates to the grade of the slope. In Adolph's crawler versus walker test, the maximal slope was approximately 40°. In comparison, Witherington and colleagues' slope consisted of an invariable 90° cliff. Secondly, there was a difference in the quality and amount of tactual information afforded to the child by the transparent shield in the visual cliff.

2.3.1 Grade of the slope

In Adolph's experiment, the steepest angle on which infants were tested was 18° larger than their boundary condition. Additionally, tests run on a 36° ramp were referred to as the impossible task (Adolph, 1997). Witherington and colleagues' visual cliff, instead, was consistently angled at 90°. It is posited that while this has created the appearance of discontinuity in infant actions, this needs not be the case.

Table 1 suggests that there is no categorical decision boundary but instead that the slope of the ramp appears to affect an infant's go-ratio relatively smoothly, without discrete jumps. Further, the data point to the fact that the steeper the ledge, the less likely an infant will attempt a descent. This is true for the novice walker as much as for the experienced crawler. Even though Adolph's experienced crawlers appear more apprehensive than the novice walkers, this does not preclude the novice walker from exhibiting some trepidation. What is not known is the go-ratio that Adolph would have discovered with a slope of 90°. Potentially, she would have observed the 0.10 go-ratio that Witherington and colleagues reported. If this were the case, then, for the novice walker, there would be no discontinuity between the two experiments. However, such a hypothesis does not cover the experienced crawler. The experienced crawlers in Witherington and colleagues were willing to cross a 90° ramp, whilst the experienced crawlers in Adolph's tests were unwilling to descend at significantly reduced angles. This does not refute the hypothesis, however, as it can be argued that the environment introduced in Witherington and colleagues biased the actions of experienced crawlers whilst innocuously affecting novice walkers.

2.3.2 Tactual input: Hypothesis

In Witherington and colleagues' study, only 45% of the crawlers and 30% of the walkers touched the deep side (Witherington et al., 2005, p.292). However, 5 of the 7 crawlers who crossed the visual cliff, first explored the situation through tactual means, whilst the 2 walkers who ventured over the cliff did not assess the environment tactually. These observations raise the questions of (i) whether tactual information factors in the child's decision to cross the deep side, and (ii) whether the ability to garner information through tactual exploration is transferred through motilities? For example, could the fact that the novice walker is just beginning to learn how to balance in a new motility directly affect his/her aptitude in tactually exploring an environment? Balancing on one foot to explore with the other, or squatting forward to touch a surface would often lead to instability. This would not eliminate a child's explorative capacity, as Witherington and colleagues showed. However, it might attenuate the accuracy of the information accrued. Consequentially, an infant training in the motility of walking might rely upon his/her tactual input less than an experienced crawler with a superior sense of balance. This could create behaviour in line with Adolph's hypothesis while still duplicating Witherington and colleagues findings.

Thus, in this paper, it is posited that the novice walker's go-ratio of 0.10, assessed in Witherington and colleagues' study, would approximate the infants go-ratio on a veritable cliff of 90°. Additionally, it is postulated that if the visual cliff was sloped at angles comparable to those in Adolph's experiments, then similar go-ratios would be observed for novice walkers. Finally, it is hypothesized that the visual cliff biases experienced crawlers and that if an experienced crawler equally depended on visual and tactual input, then the infant's go-ratio would be altered on the visual cliff where the infant visualizes a cliff of 90°, but tactually senses a flat path of 0°.

3. Visual and tactual infant model

Returning to the central question of the presence or lack of transfer between motilities, we seek to reconcile the seemingly contradictory findings of Adolph and Witherington and colleagues by asking whether Witherington and colleagues' observations could actually reflect the behaviour of agents without any transfer capability. To do so, two types of agents were simulated: (1) the experienced crawler and (2) the novice walker. Each agent was trained to approximate the go-ratio of either one of Adolph's experienced crawlers or one of her novice walkers. The agents were implemented as multi-layer perceptrons (MLP) trained using two inputs: visual and tactual perception. The models were tested on simulated veritable slopes (in contrast to simulated visual cliffs). The visual and tactual input was determined by the angle of the slope in the range [0...90] degrees.

3.1 Specificity of the tactual input

As the novice walker begins to learn how to balance in a new motility, it would be erroneous to assume that the capacity to gather tactual information is independent from the child's balance. Lack of balance could affect the infants ability to assess the environment tactually, i.e., a 40° slope may feel like a 50° or 30° slope if the child looses balance while assessing the grade. Thus, in our model, an increased amount of noise was added to the tactual input compared to the visual input. Conversely, it was hypothesized that an experienced crawler would have established the balance necessary for effectively assessing an environment both visually and tactually, and therefore, the model was trained with relatively aligned visual and tactual inputs.

3.2 Methods

For each type of agent, 20 agents were trained to represent the 20 experienced crawlers and 20 novice walkers in Witherington and colleagues' experiment. Each agent consisted of a multi-layer perceptron trained via the backpropagation algorithm. Two inputs were utilized, visual and tactual perception. There were two hidden layers with ten nodes each (these parameters were arbitrarily chosen but should have little bearing on the results). Each network was trained to learn a function related to Adolphs go-ratio for either an experienced crawler or a novice walker, i.e., the agents were developed to mirror a behavioural lack of transfer. The functions which the networks approximated were pieced together from Adolph's experimentations. The average boundary angle for an experience crawler was approximately 20°, while for a novice walker it was 6° (Adolph, 1997, p.58). As shown by Table 1, the approximate go-ratio of an infant at the boundary angle was 0.90, for both motilities. Furthermore, the average go-ratio at +5 the boundary condition was 0.20 for an experienced crawler and 0.75 for a novice walker. A Gaussian function (parameters μ and σ) was used to approximate the ratios from Table 1 and compensate for the scarcity of experimental data points.

Table 3: Go-Ratio of both the experienced crawler (EC) and novice walker (NW) model.

Model type	Angle relative to boundary	Angle	Go-Ratio
EC	0	20	0.8994
EC	+5	25	0.2017
EC	+13	33	0.0005
EC	+18	38	0.0000
NW	0	6	0.9037
NW	+5	11	0.7827
NW	+13	19	0.5461
NW	+18	24	0.3912

For an experienced crawler, $\mu = 18.5$ and $\sigma = 3.8$ were used to to duplicate the go-ratio determined through Adolphs experimentation. Additionally, the function was edited such that any angle less than the mean was allocated a go-ratio of 1. To duplicate the go-ratio of the novice crawler: $\mu = -3$ and $\sigma = 20$. Table 3 lists the go-ratios of an agent that adopts either of these functions. For an experienced crawler, at a 20° decline, the go-ratio is $0.8994 \approx 0.90$. At +5°, relative to the boundary angle, the go-ratio is ≈ 0.20. The other angles also approximate the go-ratio of both motilities shown in Table 1. Thus, if a network learned the novice walker function, then it would act with a go-ratio similar to a typical novice walker noted in Adolph's experimental findings. Figure 1 illustrates the resulting functions.

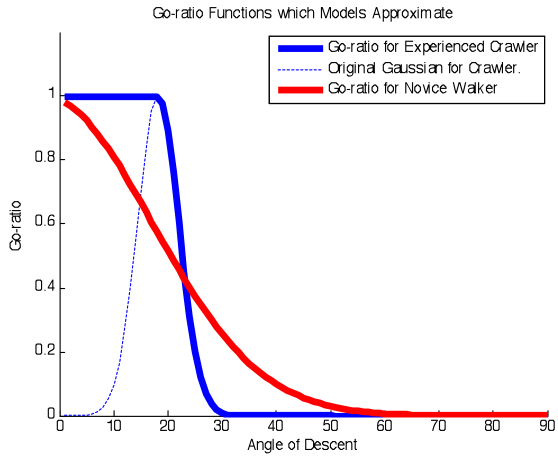

Figure 1: Visualisation of go-ratio functions for experienced crawler and novice walker models. The dashed line represents the Gaussian function of the experienced crawler without thresholding. The blue line denotes the go-ratio function learned by the experienced crawler networks. The red line describes the go-ratio function learned by novice walker networks.

It should be noted that these functions approximate findings which are themselves variable. In her control group, Adolph tested the infants on the impossible task of 36°. The results showed that experienced crawlers would attempt to descend the impossible slope with a go-ratio of approximately 0.05. This is in significant contrast to the go-ratio of 0.00 depicted in Table 1, which is approximated by the learning function. Furthermore, there is no evidence that the go-ratio of an infant follows a Gaussian function. A more thorough quantitative investigation would need to be conducted in order to solidify how the go-ratio changes as a function of slope angles beyond the boundary condition. However, with the data of only four relevant angles being available, a Gaussian function permitted a smooth continuous function and was deemed acceptable for the purposes of this model.

3.3 Training

An infant will develop an expectation of his/her interaction with the environment through exploration. As such, it was decided that a network model would learn its limitations through assessing angled slopes. To account for hypothetical perceptual disparities in how slopes were perceived across the visual and tactual modalities, noisy sensor readings were extrapolated for each given slope angle. A uniformly distributed noise in the range [-1...1] was added independently to both visual and tactual inputs. Furthermore, in the case of the novice walker, as instability might affect tactual input more than visual input (see discussion above regarding balance), a greater range of noise was used for the tactual input (range [-10...10]).

For each network, a thousand input patterns were generated for training. Input patterns were generated using a Gaussian distribution with mean 20 and standard deviation 20. This input pattern bias was meant to capture the fact that a child would, likely, rarely be placed on the edge of 90° cliffs while developing his/her understanding of the relationship between the environment and a particular motility. It could be argued that parentally initiated scaffolding and hypotheses of intrinsic motivation defend this input pattern bias. One such hypothesis of intrinsic motivation refers to a childs tendency to engage in activities slightly beyond his/her threshold (Oudeyer et al., 2000). Thus, a child is consistently attempting to improve through tasks of intermediate difficulty. Given the input patterns, the networks were trained for 20 epochs via online learning with a learning rate of 0.1 (all network parameters were selected following experimentation). The input patterns were uniquely generated for each network. Unique input patterns combined with random initial weights permitted individual differences between the infant models.

3.4 Visual cliff experiment

After the networks were trained, each network approximately mapped input angles to the appropriate go-ratio function. The infant model were then assessed on the simulated visual cliff (i.e., Witherington and colleagues' experimental scenario). This was accomplished by feeding the network an accurate visual angle, but setting the tactual input to zero, that is, the infant model visualized an angled descent, but tactually interpreted no such decline.

4. Results

After the forty networks were trained, the results were analyzed by systematically testing the networks on veritable and visual slopes of altering angles.

Figure 2 depicts sample runs for an experienced crawler network and a novice walker network. The black line illustrates the expected go-ratio of the network shown in Figure 1. The blue line portrays the output of the network on an emulated veritable slope, where the two inputs are identical. The red line represents the go-ratio of the network on a visual slope/cliff, where the tactual input is constantly zero. Note that when the network predicted a go-ratio less than zero, the value was rounded up to zero; if the network predicted a go-ratio greater than one, then the output was rounded to one. In Figure 2(top), since the blue line almost perfectly overlays the expected output, the network has successfully trained to behave on veritable slopes analogous to one of Adolph's experienced crawlers. However, the same network behaves erratically when it must predict a go-ratio on a simulated visual cliff. It should be noted that the visual cliff implemented in Witherington and colleagues' study is only simulated when

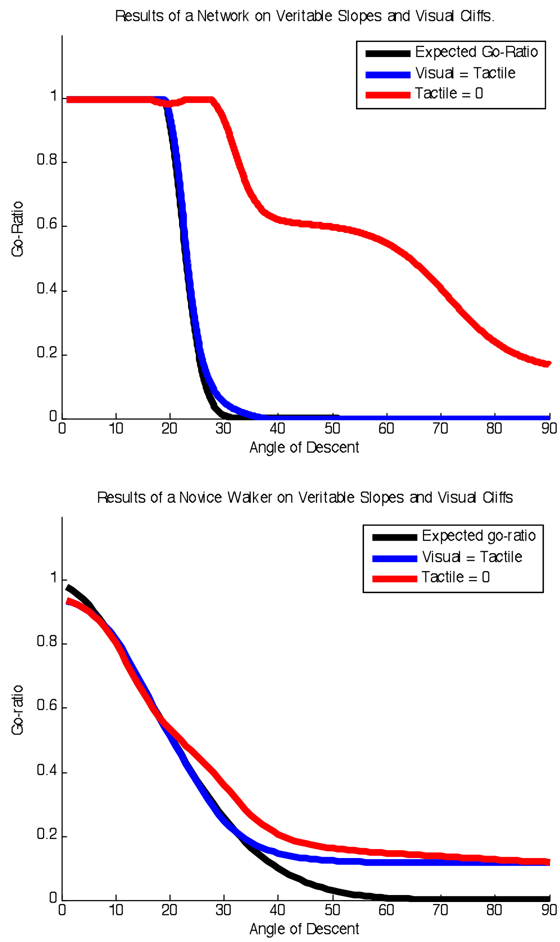

Figure 2: Sample results of an experienced crawler network (top) and a novice walker network (bottom). Black line: the expected go-ratio. Blue line: the network's prediction if visual input = tactual input. Red line: the visual cliff, where tactual input = 0.

the visual input equals 90°.

Comparing the two panels in Figure 2 yields two important observations. First, when visual and tactual inputs are identical (i.e., veritable slopes), the experienced crawler model has more adeptly learned the expected go-ratio. This is due to the fact that the training data was Gaussian distributed, with a bias on the range [0...40] degrees. This range encompasses the pertinent data for the experienced crawler, while it does not extend far enough to cover the more daring novice walker network. Both networks trained well within the training pattern range. The second difference relates to the networks' go-ratio predictions when placed on the slopes. The experienced crawler network is more sporadic compared to the expected go-ratio. The only difference between the two networks, other than the functions, relates to the amount of noise introduced to the tactual input during training. Clearly, with a larger amount of tactual noise, the novice walker network relied more heavily on its visual input when assessing a situation. Thus, when placed upon the visual cliff, it predicted a go-ratio similar to its prediction on the veritable cliff.

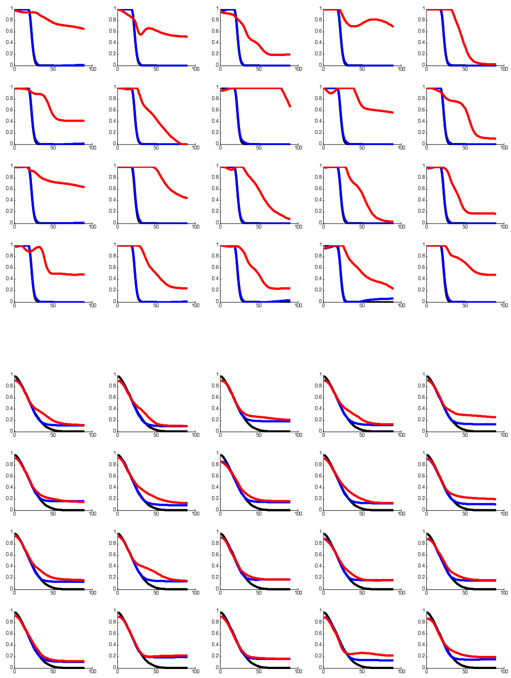

Figure 3: Go-ratios of 20 experienced crawler networks (top) and 20 novice walker networks (bottom).

4.1 Predictions of 20 experienced crawlers and 20 novice walkers

Figure 3 depicts the predictions of all 20 experienced crawlers and 20 novice walkers. As above, the experienced crawler agents go-ratios are significantly affected by whether they are tested on veritable or visual slopes. Comparatively, the novice walkers behave similarly on the veritable and visual slopes.

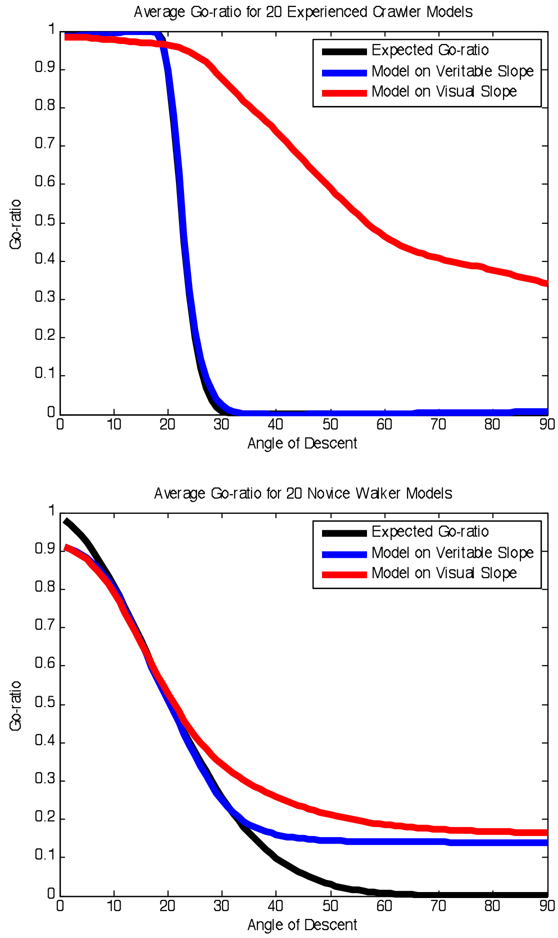

Figure 4: Average go-ratio for experienced crawler agents (top) and novice walker agents (bottom) on veritable and visual slopes.

4.2 Average Go-ratio

The average prediction of the twenty models was calculated for each angled slope in the range [0...90] degrees. The predictions were calculated for both visual and veritable slopes. This created the average go-ratio for each type of slope for both types of models. The average could then be used to assess how many of the networks, on average, would attempt to descend a particular veritable or visual slope. Figure 4 illustrates the findings. Of particular interest are the predictions of the models on the 90° visual cliff. The experienced crawler model implements a go-ratio of 0.3410, compared to Witheringtons discovery of 0.35. The novice walker model implements a go-ratio of 0.1649 compared to Witheringtons 0.10. Additionally, when placed on the veritable slopes which Adolph tested, both models accurately approximate her findings. These results therefore appear to suggest continuity between Adolph's and Witherington and colleagues' work.

4.3 *Continuity between Adolph and Witherington et al.*

Figure 5 compares the average network go-ratios with the experimental data obtained by Adolph and Witherington and colleagues. The average experienced crawler and the average novice walker were tested in the appropriate environments. For the angles tested by Adolph, the average models were assessed on veritable slopes, while to mimic the experimental process implemented by Witherington and colleagues, a 90° visual cliff was adopted. The results for the experienced crawler and the novice walker matched Adolph's and Witherington and colleagues' experimental findings.

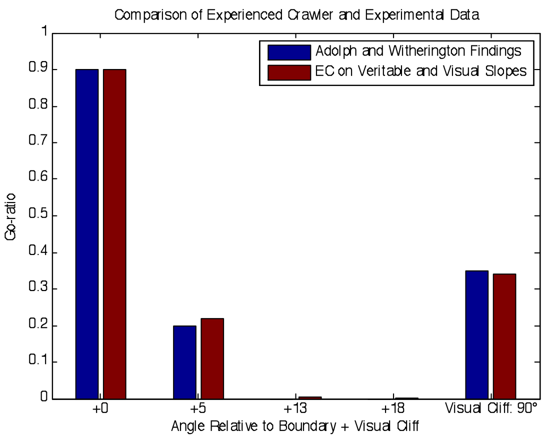

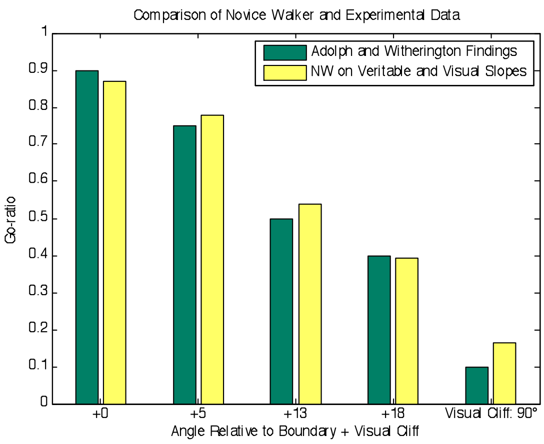

Figure 5: Comparison of average experienced crawler model (top) and novice walker model (bottom) with experimental data. The first four experimental findings relate to Adolph's veritable slopes. The last represents Witherington's visual cliff.

Thus, the models of experienced crawlers and novice walkers, based on Adolph's no transfer hypothesis, successfully unified the data obtained in the two experiments. The largest discrepancy existed on the visual cliff between the prediction of the novice walker model and the experimental findings of Witherington and colleagues who noted a go-ratio of 0.10 when the model predicted a go-ratio of 0.1649. However, the experimental findings are somewhat variable and the pattern of unification is clearly recognizable. The novice walker still behaves with significantly more trepidation than the experienced crawler.

5. Discussion

5.1 *Summary of results*

Adolph hypothesized a behavioural lack of transfer from crawling to walking. This was supported by novice walkers' willingness to rush down dangerous veritable slopes while experienced crawlers avoided such risky behaviours. Witherington and colleagues' work on the visual cliff attempted to contradict Adolph's hypothesis by demonstrating that experienced crawlers traversed visual cliffs more frequently than novice walkers. However, if an infant's ability to accurately assess the environment through tactual means is not transferred from crawling to walking, then it has been shown that it is possible for Witherington and colleagues' experimental findings to be reproduced, given Adolph's no transfer hypothesis.

5.2 *Proposed experiment*

Adolph's experimental data illustrate the non-discrete nature of infant go-ratios. Even novice walkers appear capable of partially distinguishing between slopes. Less than half of the novice walkers that would descend at the boundary angle would do so at +18° (see Table 1). The experimental data garnered from Witherington and colleagues show that novice walkers cross the visual cliff on 10% of occasions. A typical boundary angle for a novice walker is +6°. Therefore, the visual cliff is typically +84° beyond this boundary. Is it possible that a novice walker would average a go-ratio of 0.10 on a 90° veritable slope? It has been hypothesized that part of Adolph's lack of transfer hypothesis encompasses the ability of an infant to appropriately assess visual and tactual input. While the experienced crawler may be capable of accurately appraising visual and tactual input, the novice walker's assessment of tactual input may prove more volatile. This is speculatively due to balance concerns which might cause the novice walker to learn to rely more heavily on visual rather than tactual input during the early stages of exploration in the new motility. In order to test whether the usefulness of tactual input transfers from crawlers to walkers, it is suggested that a visual cliff apparatus is adopted to test infant's go-ratios on angles of descent not limited to 90°. The results could then be compared to Adolph's findings on veritable slopes. The model presented in this paper predicts that an experienced crawler will more frequently cross a visual decline than a veritable decline. Furthermore, a novice walker will cross a visual and veritable decline at similar rates. This prediction could be tested by assessing experienced crawlers and novice walkers on contrastingly angled visual slopes. For example, it is posited that an experienced crawler will cross a 30° visual slope more frequently than a 30° veritable slope, while the novice walker will cross each with a similar rate. If this occurred, then it is shown that Witherington and colleagues' experiment does not contradict the lack of transfer theory proposed by Adolph. Additionally, this could elucidate one of the reasons why infants behave as if they lack transfer.

5.3 Lack of transfer and future questions

This paper only briefly discussed how or why a lack of transfer occurs. The volatility of acquiring input in a new motility may partially attribute to the cause of the behaviour noted by Adolph. However, this likely only brushes the surface of potential reasons. If a child could accurately garner information from an experienced motility and a novel motility, would the child still display a lack of transfer? If a child is not willing to risk failure, how would s/he learn the different potentials of crawling and walking? It would be disadvantageous for a crawler to transfer all knowledge between motilities as some knowledge is specifically useful only in the medium of crawling. How would the child/development decide what knowledge is transferable except through trial and error? This trial and error would then behaviourally appear as a lack of transfer. Finally, as shown by Table 1, whilst a 1-week crawler and a 1-week walker have almost identical go-ratios, a 10-week walker is more hesitant than a 10-week crawler. Is this due to the fact that the child is implementing some transfer after a period of testing via trial and error, or could this phenomenon relate to the age of the child? The neural developmental process of a 10-week walker may permit quicker assimilation of knowledge compared to the 10-week crawler. Either way, significantly more experimental exploration is required.

6. Conclusion

It was hypothesized that Witherington and colleagues' experimental findings were not in conflict with Adolph's lack of transfer hypothesis. An existence proof was developed in which an infant was modelled through the training and testing of a multi-layer perceptron. It was posited that a novice walker's instability would hinder the assessment of tactual information. This could then cause the agent to rely more heavily upon other, more dependable, inputs. Conversely, an experienced crawler could regulate his/her stability sufficiently to incorporate tactual information. A model of this hypothesis was developed by training agents to meet Adolph's transfer-less findings. It was shown that a model of an experienced crawler and a model of a novice walker, both of which lacked transfer, could behave in such a way as to replicate Witherington and colleagues' experimental findings. Thus, it is not evident that Witherington and colleagues have contradicted Adolph's no transfer hypothesis. The behavioural disparity between infants implementing different motilities still requires significant exploration. While, hopefully, this paper aids in assimilating behavioural data, it only speculatively touches on the reason behind such behaviour. Hypotheses such as Adolph's sway model are still a long way from general acceptance (Adolph, 2000). By implementing experiments, like the one proposed in this paper, hopefully, behavioural consensus will be reached.

References

Adolph, K. (1995). Psychophysical assessment of toddlers' ability to cope with slopes. *Journal of Experimental Psychology*, pages 734–749.

Adolph, K. (1997). Learning in the development of infant locomotion. *Monographs of the Society of Research in Child Development*, pages 1–162.

Adolph, K. (2000). Specificity of learning: Why infants fall over a veritable cliff. *Psychological Science*, pages 290–295.

Oudeyer, P.-Y., Kaplan, F., and Hafner, V. (2000). Intrinsic motivation systems for autonomous what. *IEEE Transactions on Evolutionary Computation*, pages 1–23.

Prince, C., Helder, N., and Hollich, G. (2005). Ongoing emergence: A core concept in epigenetic robotics. *Proceedings of The Fifth International Workshop on Epigenetic Robotics*, pages 63–70.

Vigorito, C. and Barto, A. (2010). Intrinsically motivated hierarchical skill learning in structured environments. *IEEE Transactions on Autonomous Mental Development*, 2(1):1–15.

Witherington, D., Campos, J., Anderson, D., Lejeune, L., and Seah, E. (2005). Avoidance of heights on the visual cliff in newly walking infants. *Infancy*, pages 285–298.