Perception as Prediction

Stephen A. Hockema (shockema@indiana.edu)
Psychology and Cognitive Science, Indiana University
1101 E. Tenth St., Bloomington, IN 47405 USA

Abstract

Learning is often about prediction. This paper asks whether perception is also. The main idea is that perception is a stream and that perceivers learn the trajectory through which one moment in that stream turns into the next. A behavioral experiment with children is described that tests two hypotheses developed from this idea.. In the experiment, children briefly watch a transformation (e.g., a triangle increasing in size and/or saturation). If children learn the trajectory of change and if prediction is at the core of perception, then a subsequent statically presented object should trigger the perceptual system to anticipate the "next state". To test this, children were asked to make same/different judgments that should, by hypothesis, be interfered with by the learned trajectory. Children became less able to detect pairs that were the "same" when asked to make judgments about the dimension that they had seen varied. Furthermore, there was evidence that two dimensions could be made more integral by covarying them simultaneously. Both of these results were simulated with a simple connectionist model constructed to embody a predictive mechanism. Taken together, these results lend support to the idea that the perceptual system is designed to make predictions in time and that this architecture gives a "dynamic" aspect to perception.

Introduction

Perception is an interaction *in time* between a dynamic mind and a dynamic world. It is crucial to the understanding of this process that *both* sides of the relation are changing in time. This paper examines two key implications of this fact: 1) an element of *prediction* must be inherent in the perceptual process, and 2) the system should be *adaptable* to a changing environment.

There is ample evidence for the adaptability of the perceptual system in the short and long term. In the short term, the evidence shows up in the form of aftereffects and priming. For example, motion aftereffects (MAEs), illusions of motion (without displacement) that occur after viewing real motion in a certain direction for a short time, are considered to be perceptual adaptations that serve to keep the system "in balance" (Anstis, Verstraten, & Mather, 1998). In the long term, adaptation shows up as the perceptual learning of new psychological dimensions and features and the readjustment of the relative attention paid to existing dimensions (Goldstone, 1998). In brief, our perceptual systems constantly tune themselves to the environment.

Furthermore, our perceptual systems must also be tuned to *anticipate* the future. In the short term, the system must be able to predict, at a very low level, how our environment and the things in it change and appreciate which changes are "normal" and which are unexpected. Normal changes include regular transformations along dimensions, for example changes in position, size, orientation, luminance, pitch, and so on.

There is evidence for the predictive capabilities of the system in many everyday activities, like tracking moving objects behind occluders, navigation through crowds, simple eye-hand coordination, and in several related experimental phenomena where the system exhibits "momentum" when tracking predictable changes. For example, in the phenomenon known as "representational momentum", subjects learn to anticipate the continuation of a transformation, for example an object moving across a computer screen. When the object suddenly disappears and subjects are probed about its final position, there is evidence that the remembered position is shifted forward along the path of the trajectory (Freyd, 1992). This has been taken by some (e.g. Freyd, 1992; Hubbard, 1999) as evidence of a "dynamic representation" that continues to move forward after the real object has disappeared. However, the evidence is also consistent with the possibility that the trajectory of the object is being perceptually *predicted* and thus the perceived position of the object at any given instant in time is actually ahead of its actual position.[1] Such a view would be consistent with an explanation for several other perceptual illusions given recently by Changizi and Widders (2002) and Changizi, Nijhawan, Kanai, and Shimojo (2003).

This paper will present further empirical evidence of a perceptual adaptation that is triggered by predictable variation. A simulation of the experiments using a simple model of a predictive mechanism will then be presented as additional support for these ideas.

Experiment

If an adaptable, predictive mechanism is built into our low-level perceptual systems then it should be possible to prime the system to anticipate and perceive change as in the momentum effects, even when tested with static objects. If the primed change is along a dimension, this might disrupt the perception of the value for this dimension in a subsequently presented static object. If caused by an adaptable, predictive mechanism, this disruption would have two characteristics:

1. a perceptual *shift* in the direction of the primed change; and
2. less certainty about the precise value of the dimension.[2]

---

[1] Of course it is possible that both of these theories are operating.

[2] This is analogous to uncertainty principles in physics, although it also stems from the imperfection that will be inherent in any physical predictive mechanism, especially in circumstances where it has only had a brief exposure to the trajectory that it is trying to learn to predict.

This experiment deals with the second characteristic. It was designed to test the hypothesis that priming the perceptual system with bidirectional change along a dimension could lead to "perceptual spreading" along it. Such an effect should make it more difficult for perceivers to make accurate *comparisons* along this dimension subsequent to experiencing this priming. In particular, the spreading should increase the likelihood that perceivers will see two things that are actually the same as being different on the dimension.

In this experiment, this hypothesis was evaluated using a task where the dimensional change was expansion and contraction in "size". Thus, it was expected that subjects would have more difficulty judging two shapes as being the same size after exposure to change along the size dimension than after being exposed to no change, or to change along an irrelevant dimension. In this experiment, color saturation was used as the control dimension. The participants were preschool-aged children under the assumption that the developing perceptual system is more sensitive to these priming effects.

The use of preschool children was also motivated by the desire to investigate the usefulness of a predictive mechanism to a developing perceptual system. Based on the simple idea that "things that change together go together", the hypothesis was formulated that if the perceptual system adapts to a coherent transformation along more than one dimension, these dimensions might become perceptually "fused", or more *integral* in the sense formalized by Garner and Felfoldy (1970).

Method

The experiment was a between-subjects design. Subjects were randomly assigned to one of four conditions. There were also four possible orderings of the test trials (see below). Subjects were randomly assigned to one of these four orders, which were counter-balanced across the four conditions.

Participants Thirty-two children from the Bloomington, IN area have participated. Children were all between 4 years and 4 years, 8 months of age. (Mean age was 4.3 years.)

Procedure The experiment was divided into three phases, all of which used a computer to present the stimuli.

In the first, "warm-up" phase, subjects were familiarized with the computer and trained to press the space bar whenever they were presented with a pair of shapes on the screen that were the same *size*. This phase was the same for all subjects, regardless of what condition they were in. The items presented included pictures of various common and colorful objects (e.g. balls, flags, hearts, etc). Across trials matching and mismatching objects and matching and mismatching sizes varied orthogonally so as to instruct the child as to the importance of attention to size only.

Subjects were presented with 32 warm-up comparisons, starting at a slow presentation rate and gradually increasing to 1.25 seconds per pair. Whenever a new pair appeared, the computer would emit a short beep. A different pitched beep would sound whenever the space bar was pressed. If an error was made (either the pair were the same size and the child failed to press the space bar, or the pair were different sizes and the child incorrectly pressed the space bar), the computer would emit a lower beep and the warm-up sequence would stop. The experimenter would then explain the mistake to the child, pointing out what they should have done, and then continue the sequence. If, after 32 trials, the child had not yet grasped the task, the training phase was repeated. If they still had not grasped the task after three passes, the subject's data was replaced.

In the second phase, the "priming" phase, all children were shown a simple video animation that lasted 80 seconds and consisted of 20 repetitions of a "transformation" of two shapes. The transformation that children saw depended on their condition assignment and will be described in the next section, however the instructions were identical for all four conditions. Children were told that they were playing a game in which they were supposed to press the space bar whenever the shapes "stop changing" and begin to "change back". The game was simply a ruse to help keep the children paying attention to the crucial animations.

The final phase of the experiment, the "test" phase, was similar to the training phase. Here subjects were again told to press the space bar as quickly as possible whenever they saw a pair of shapes that were the same size. In the test phase, the shapes being compared were either a pair of circles or a pair of squares. Half of the comparisons were the same size and half were different. The shapes were either both blue or both red, with their saturations varying as described in the next section. As in the warm-up phase, the computer emitted a beep whenever a new pair was displayed and a different pitched beep whenever the space bar was pressed. No feedback was provided during the test phase.

During the test phase, each pair was displayed for exactly 1.25 seconds. After being given the instructions, and prior to starting, subjects were warned that the speed would be fast and that they should get ready. There were 32 test stimuli in this phase, broken down into four sets of eight. After each set of eight, the computer would pause and the subject would be reminded of the instructions, and told to get ready again. Test trials were presented to subjects in one of four random orders. To ensure that children were still on task, any child that pressed the space bar on at least 20% or more than 80% of the test trials was replaced.

Materials The priming phase was designed to "teach" the perceptual system a predictable trajectory of change. There were four possible animations used corresponding with the four conditions. All animations showed the gradual transformation of a pair of side-by-side triangles. The left and right triangle always transformed identically and in synch with one another. The left triangle was always red and the right triangle was always blue. The four conditions were as follows:

1. <u>Control:</u> increasing and decreasing saturation;
2. <u>Size-Only</u>: increasing and decreasing size;
3. <u>Correlated:</u> both size and saturation increasing and decreasing: the bigger the triangles got, the more saturated, and the smaller they got, the less saturated;
4. <u>Anti-correlated:</u> both size and saturation increasing and decreasing: the bigger the triangles got, the less saturated, and the smaller they got, the more saturated.

Since *size* is at least one of the transformed dimensions in Conditions 2-4, the three conditions will collectively be referred to as the "Size-change" conditions. Conversely, since *saturation* is at least one of the transformed dimensions in Conditions 1, 3 and 4, these three conditions will be referred to together as the “Saturation-change” conditions.

The minimum size (area) of the triangles in the three Size-change animations, as displayed on the monitor, was $1.83cm^2$ (base $1.63cm$, height $2.25cm$); the maximum size was $29.25cm^2$ (base $6.5cm$, height $9.0cm$). The minimum saturation in the Saturation-change animations was 0.1 on a scale of 0 to 1; the maximum was 1.0. For Condition 1, the size of the triangles remained constant at $29.25cm^2$. For Condition 2, the saturation of both triangles remained constant at 0.8. For the Size-change animations, the triangles always started and ended at their smallest point.

In the Testing phase, there were four different types of comparison possible. These were among shapes that were either big, medium or small in size (see Table 1) and high, medium or low in saturation (1.00, 0.45 and 0.20 respectively in the range of 0–1).

Table 1: Actual On-Screen Areas (in cm^2)

Term	Squares	Circles
Big	27.56	9.33
Medium	12.25	5.25
Small	4.86	2.33

The four types of comparison were as follows:

1. Identical: Pair being compared were identical in both size and saturation. There were four variations of this: big/high, big/low, small/high, small/low.

2. Saturation Different: Pair being compared were the same size, but differed in saturation. There were two variations of this: big/high compared to big/low and small/high compared to small/low.

3. Size Different: Pair being compared were the same saturation, but differed in size. There were four variations of this. In the first two, the pair were both high saturation. In the second two, the pair were both low saturation. One of the shapes was always medium size and the other was always either small or big.

4. Both Different: Pair being compared were different both in size and saturation. In these pairs, the bigger shape always had the higher saturation. (This was done for reasons related to other experiments not reported here.) As in the Size Different comparisons, one of the shapes was always medium size and the other was always either small or big.

Results

Children’s errors were classified by the type of comparison trial in which they occurred and the condition to which the subjects were assigned. There were two broad classes of errors: *Misses*, where the shapes were the same size yet the subject failed to press the space bar, and *False Alarms* where the shapes were different size and yet the subject incorrectly pressed the space bar.

Figure 1 shows the average number of Misses broken down by the four priming conditions and two types of relevant test trial (Identical trials on the left and Saturation Different trials on the right in each group). As can be seen, there were significantly more Misses in the Size-Change conditions as compared to the Control Condition ($p < .02$ for the Identical trials, $p < .002$ for the Saturation Different trials).

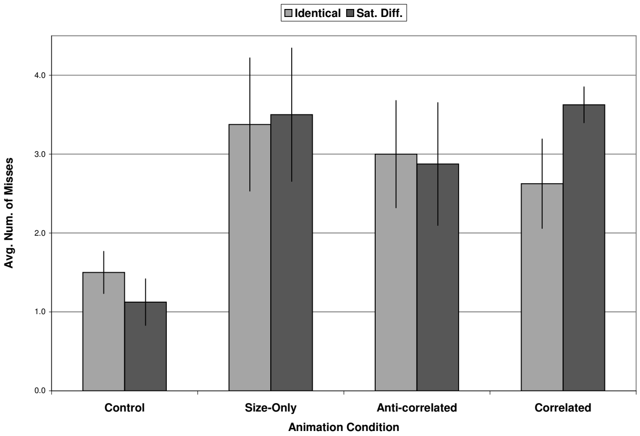

Figure 1: Average Misses for Same-Size trials

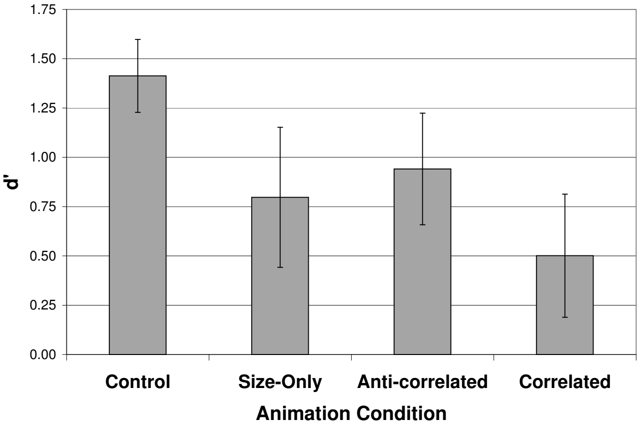

Figure 2: Average disciminibility

An analysis from the perspective of Signal Detection Theory (SDT) was also performed. Figure 2 shows the average d' (discriminability or sensitivity) for each of the four conditions.[3] As can be seen, the three Size-change conditions had a significantly lower average d' value ($p < .03$) than the Control Condition. (In this context, lower d' values mean that it is harder to distinguish Same Size shapes from Different Size shapes.) The three Size-change conditions also had a marginally significant ($p < .06$) higher average β value (criterion for pressing the space bar): they were less likely to respond “Same Size” in general. Also, the Correlated condition by itself also had a significantly lower d' ($p < .01$) and higher β ($p < .03$) than the Control condition.

Figure 3 shows the False Alarms, again broken down by the four conditions and two types of relevant test trial (Size Different trials on the left and Both Different trials on the right). Considering Figure 3, notice that the gap between the number of False Alarms in the Different Size, *Same* Saturation trials and the number of False Alarms in the Different Size, *Different* Saturation trials increased in the conditions where size and saturation were covaried (Correlated and Anti-correlated) over what it was in the conditions where they were varied in-

---

[3] For one subject in Condition 2, d' was infinite because the subject had a Hit Rate of 1.0. For this subject, d' was estimated using an adjusted Hit Rate of $15.5/16 = .9688$, yielding a d' of 3.01.

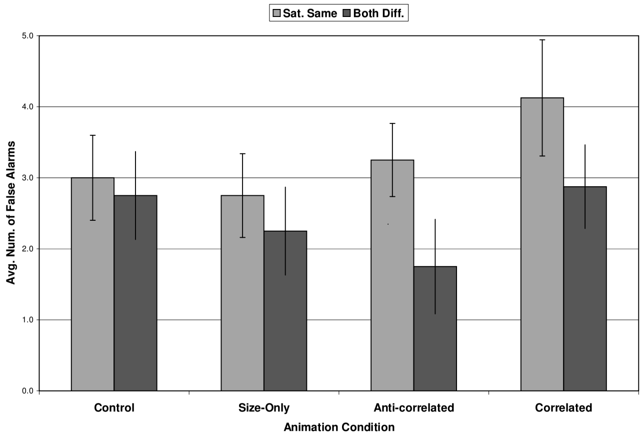

Figure 3: Average False Alarms for Different-Size trials

dependently. The graph shows the group averages and the error bars are the standard error of these means. For each subject, the difference between their number of False Alarms on the Different Size, Same Saturation trials and the Different Size, Different Saturation trials was computed. This difference was found to be significantly above zero in both the Correlated ($p < .006$) and Anti-correlated ($p < .028$) conditions. Comparing to the Control and Size-only Conditions, it can be seen that no such difference existed where the dimensions were varied independently during the priming event. When the within-subject difference scores of the two "covaried" conditions are compared to the difference scores of the two "independent" conditions, the "covaried" differences are significantly bigger ($p < .04$). Thus, differences in the irrelevant dimension of saturation had a larger effect on subjects' comparisons of size during test when size and saturation were covaried together during the priming event.

To see if the effect of the priming event looses its potency over time, the probability of error on any given test trial was correlated with how long after the priming phase it occurred. Table 2 shows the correlation coefficients between error frequency and trial. In no case was the correlation significant. Thus, performance did not change throughout the test phase.

Table 2: Correlation between trial number and error prob.

Condition	r
Control	$-.12$
Size-only	$.12$
Correlated	$-.17$
Anti-correlated	$-.29$

Discussion

The results appear to support the first hypothesis of perceptual spreading. Priming with size transformations in the three Size-change conditions led to a decrease in the ability to detect shapes that were the same size in the Test phase. This is consistent with "spreading" (bidirectional propagation) along the size dimension due to the anticipation of change by the perceptual system. Such spreading apparently leads to difficulties in comparison.

The speeded comparison task focusing on just size was chosen specifically to induce errors — it is known that children have difficulty ignoring variation on an irrelevant dimension to focus on just one dimension (Smith & Evans, 1989). The point of interest here was how the type of errors would be affected by the prior experience with a systematic trajectory of change. It is noteworthy that in the Control and Size-only conditions, performance was *not* significantly different for test comparisons where the saturation was different than it was for comparisons where the saturation was the same. Thus, size and saturation were fairly separable in these two conditions. The fact that the gap in the number of False Alarms widened in the two "covaried" conditions over the two "independent" conditions is taken as support for the second hypothesis: It appears that coherent, predictable change along both the size and saturation dimensions caused them to become more integral, such that saturation differences had significantly more of an effect on size comparisons. This finding has developmental implications, providing a possible account for how low-level *features* and *properties* that start out as perceptually distinct can congeal into perceptual *dimensions* when they are experienced as covarying in a predictable and coherent way. One thing that is very interesting about the present results is that it took only a relatively short exposure (80 seconds) to such regularity to produce this effect! Further, the effect appeared to decay slowly (not measurably) over the course of the experiment. Taken together, the quick adaptation and slow decay of the effect suggest that an interesting avenue of future research will be to explore the relationship between the amount of experience with certain types of transformation and the duration of the adaptation. This type of predictive learning of correlations may well be a potent part of the developmental process.

Signal Detection Theory: An Alternative Explanation

Signal Detection Theory (SDT) is a type of analysis that assumes that the system responsible for making decisions is inherently noisy, such that, for example in the case of this experiment, even when the shapes are obviously very different in size there is still a chance that the subject will respond "Same Size". What makes SDT really useful here is that it provides a way of separating subjects' propensity to respond "Same Size" (their *bias* or *criterion*) from their ability to tell the difference between the Same Size test trials and the Different Size trials. The predicted perceptual spreading effect should result in a decrease in this second ability, not merely a change in bias. And, importantly, there was a significant difference in discriminability (d') triggered by experience with the Size-change transformations. Subjects had a harder time distinguishing between Same Size and Different Size test trials in these conditions, as the spreading hypothesis would predict.

The assumption of inherent noise that underlies SDT also provides another way to explain the data. It is possible that priming the system with changing size ("dynamic priming"), increased the internal noise in the perceptual system related to size judgements. (In SDT terms, this amounts to increasing the variance of both the "signal absent" and "signal present" distributions.) This would show up in the analysis as a decrease in d', as seen. This provides a potentially useful description at a different level of abstraction. Indeed, the "uncertainty" stemming from the proposed predictive mechanism might provide a mechanistic explanation (at a lower level) for the increased noisiness.

Yet another possibility, and perhaps a simpler account of the present data than the adaptive prediction hypothesis, is that the dynamic priming of size increased the children's sensitivity to size differences, enabling them to make finer-grained distinctions of size. In this case, the inherent noisiness of the system would have more of an effect. Subjects would be more sensitive to very slight discrepancies in size caused by noise and would be less likely to respond "Same Size" in general (i.e., they would both Miss more and False Alarm less). This would show up in the analysis as a difference in their response bias, β.

The fact that there was a marginally significant increase in β in the three Size-change conditions means that this hypothesis cannot be ruled out here as a possibility.[4] It should be noted that this alternative only offers an explanation for the overall decrease in performance in the three Size-change conditions relative to the Control condition, but it does not address the effects related to increased Integrality. Nor does this bias-shift account explain the (more significant) shift in d', which indicates that subjects in the size-change conditions really did have more difficulty distinguishing the Same Size test trials from the Different Size trials. For this, an adaptive, prediction mechanism still seems to be reasonable.

A motion aftereffect? The analysis of the probability of error as a function of trial is interesting because it shows that the induced effect does not decay rapidly. If it were a typical motion aftereffect, 80 seconds of exposure to the animation motion might be expected to trigger on the order of 10 seconds of aftereffect, as motion aftereffects typically decay with the square root of the time exposed to the inducing motion (Anstis et al., 1998). Yet the error rate showed only a very slow decay, lasting over a minute.

Furthermore, the present effects are also set apart from typical motion aftereffects in that they occur after seeing *bidirectional* motion. Motion aftereffects typically occur in the opposite direction from the direction of inducing motion (Anstis et al., 1998). The theory behind this is that the visual system adapts to correct what it (mistakenly) takes to be drift in the neurons detecting motion in the inducing direction and lowers their weight relative to neurons sensitive to motion in the opposite direction. If this accurately reflects what really happens (in a MAE), then bidirectional motion should not produce the aftereffect because there will be no incentive for the visual system to "suspect drift" in the first place, the opposing motions will cancel each other out.

Thus, it would appear that the effect observed in this experiment is a new and different type of perceptual adaptation that is related to traditional motion aftereffects, but underwritten by a potentially different mechanism.

Model Simulation

The basic principle of perceptual prediction was embodied in a connectionist model consisting of a simple recurrent network (Elman, 1990). Its task was to actively sample its "sensory" input and try to predict how it changes in time. By its nature, such a network allows for "supervised" learning in the sense that it can validate its predictions by what eventually happens. Thus, whenever the model encounters consistent, gradual, continuous variation (for which it cannot already account), it might actively train itself to predict this variation. The model was constructed such that its predictions could be fed back into its input units, enabling extrapolation in time and giving perception "temporal extent".

There is not room in this paper to go into the details of the input representation and training procedure. Basically, a shape was represented by its values on the size and saturation saturation dimensions. There were three different training sequences corresponding to the four animation conditions in the experiment. (The Correlated and Anti-correlated conditions were equivalent in the model representations, so they were combined into a single simulation condition called "Covaried".) For each pattern in a sequence, the network was trained to predict the next pattern. A small amount of noise was injected into this process.

Sixteen networks were trained and tested in each condition on the same sequences of test pairs that the children saw in the experiment. The model made *comparisons* as follows: Given two patterns to compare, one was chosen to go first, passed through the network and the outputs were buffered. Then the other pattern was passed through the network and its outputs on the nine size units were compared to the buffered outputs of the first pattern using a cosine distance metric. If they fell within 30 degrees of one another, the process was stopped. If not, the buffered outputs from the two patterns were then fed back in as inputs (keeping the same context layer activations).[5] This process repeated until either the output vectors eventually came within $\lambda = 30$ degrees of one another or 10 iterations had gone by. Whenever this process terminated, the following value was calculated:

$$\text{How Different} = \text{Iterations Required} + \frac{1 - cos(\theta)}{1 - cos(\lambda)}. \qquad (1)$$

This equation gives an estimate on how different the two (*dynamic*) representations were from one another. The more iterations that were required, the more different the patterns were. θ is the final angle between the two output vectors.

Finally, the model decided whether or not to say "Same Size" for the patterns based upon how different they were from one another. The likelihood of saying "same" was inversely proportional to the score computed by Equation 1. This was operationalized with:

$$\text{Say Same} = (rand < \exp(\frac{-\text{How Different}}{\beta})) \qquad (2)$$

where "*rand*" was a random number uniformly distributed between 0 and 1, and β was a "bias" parameter that was set to 2 for this simulation.

Results and Discussion Figures 4 and 5 show the average Misses and False Alarms (respectively) over the simulation runs. These are comparable to Figures 1 and 3.

As can be seen, the basic trends that were present in the child data were also present in the model simulations. In particular, there were significantly more misses in the Size-change conditions as compared to the Control condition. Furthemore, the degree of integrality significantly increased in

---

[4] This has been taken up in other experiments not reported here.

[5] The prediction of dynamic spreading stems from this.

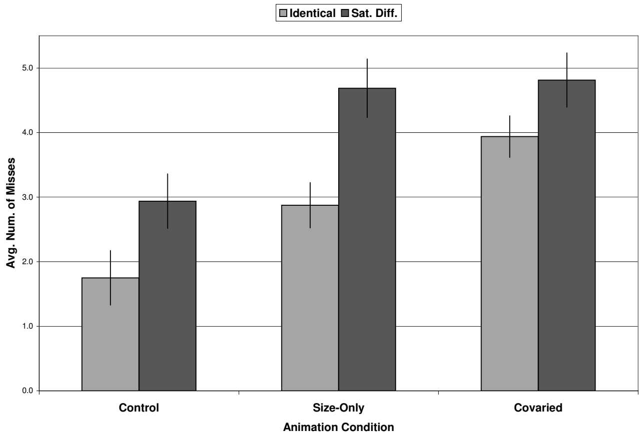

Figure 4: Misses from model simulation

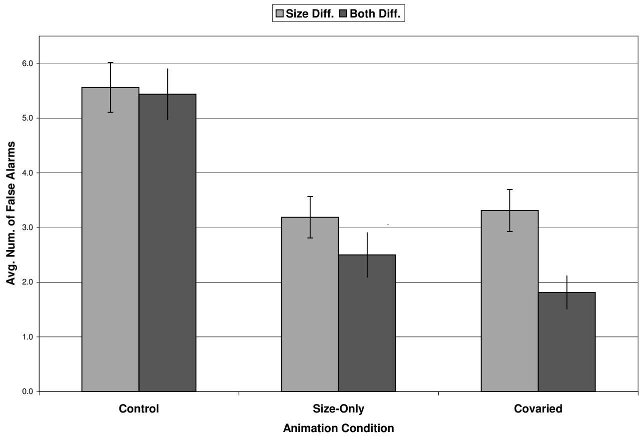

Figure 5: False Alarms from model simulation

the Covaried condition as seen in the increased False Alarm gap of Figure 5. This was due to the fact both the size and saturation dimensions shared the hidden layer and the network learned that each of these dimensions was a good cue to predict the next state of the other.

General Discussion

At the core of this paper is the idea that our perceptual systems are oriented around *transformations*. Transformations contain rich information about the structure of the world. Indeed, it is through this temporal structure that we perceive atemporal structure. (For example, movement is necessary for the detection of occluding edges.) Furthermore, perception itself is a process: it has temporal extent. There is never a single "instant" when we achieve a percept.[6] Moreover, given that the objects of our perception can be changing at the same time we are perceiving them, it behooves us to learn to anticipate their transformations.

The experiment presented herein was designed with this in mind. It provided evidence for a perceptual adaptation in response to brief experience with a predictable trajectory. It showed that there is a dynamic component involved in perception, even with *static* shapes that are *perceptually present*. It also showed that experience with coherent transformations might have developmental consequences, given that the adaptation caused two dimensions that were initially fairly separable to become more integral. The model simulation then tried to flesh out one way a simple predictive mechanism could simultaneously explain both of these effects.

Taken together, these data support the idea that perceptual adaptations go beyond being temporary adjustments to unusual environments and can have important developmental consequences. Indeed, this is the real stuff of development. Long term effects are achieved as the accumulated result of many small "tweaks" to the system occurring on a situation-by-situation basis. Thus, understanding how the perceptual system can and does adjust in short time windows to specific situations should be useful towards increasing our understanding of the type of lasting changes that the architecture can achieve. These results are admittedly only a first step in that understanding, but they do indicate several directions for future work in this regard.

Acknowledgements

Thank you to Linda Smith, Michael Gasser, Robert Goldstone and Robert Port for their feedback on this research.

References

Anstis, S., Verstraten, F. A. J., & Mather, G. (1998). The motion aftereffect. *Trends in Cognitive Sciences*, *2*(3), 111–117.

Changizi, M. A., Nijhawan, R., Kanai, R., & Shimojo, S. (2003). Perceiving-the-present and a general theory of illusions of projected size, projected speed, luminance contrast and distance. (Under Review. Found online on 12/1/2003 at URL http://www.geocities.com/changizi/pp3.pdf.)

Changizi, M. A., & Widders, D. (2002). Latency correction explains the classical geometrical illusions. *Perception*, *31*, 1241–1262.

Elman, J. (1990). Finding structure in time. *Cognitive Science*, *14*, 179–211.

Freyd, J. J. (1992). Dynamic representation guiding adaptive behavior. In *Time, action, and cognition: Toward bridging the gap.* Dordrecht, The Netherlands: Kluwer Academic Press.

Garner, W. R., & Felfoldy, G. L. (1970). Integrality of stimulus dimensions in various types of information processing. *Cognitive Psychology*, *1*, 225–241.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Hillsdale, NJ, USA: Lawrence Erlbaum Associates.

Goldstone, R. L. (1998). Perceptual learning. *Annual Review of Psychology*, *49*, 585–612.

Helmholtz, H. L. F. von. (1866). Concerning the perceptions in general. In *Treatise on physiological optics* (Vol. III, pp. 1–37). (Translated by J. P. C. Southall, 1925, *Optical Society of America*. Reprinted in 1962, New York: Dover Publications, Inc.)

Hubbard, T. L. (1999). How consequences of physical principles influence mental representation: The environmental invariants hypothesis. In *Fechner Day '99: The end of 20th century psychophysics — Proceedings of the 15th annual meeting of the International Society for Psychophysics.* Tempe, AZ, USA: International Society for Psychophysics.

James, W. (1890). The Perception of Time. In *The Principles of Psychology* (Vol. 1, pp. 605–642). Cambridge, MA, USA: Harvard University Press.

Smith, L. B., & Evans, P. (1989). Similarity, identity, and dimensions: Perceptual classification in children and adults. In *Object perception: Structure and process.* Hillsdale, NJ, USA: Lawrence Erlbaum Associates, Inc.

---

[6] These ideas have a long history in psychology, going back through Gibson (1979) at least as far as James (1890) and Helmholtz (1866).