

2016

The role of Uncertainty in Categorical Perception Utilizing Statistical Learning in Robots

Nathaniel V. Powell
University of Vermont

Follow this and additional works at: <http://scholarworks.uvm.edu/graddis>

 Part of the [Cognitive Psychology Commons](#), [Robotics Commons](#), and the [Statistics and Probability Commons](#)

Recommended Citation

Powell, Nathaniel V., "The role of Uncertainty in Categorical Perception Utilizing Statistical Learning in Robots" (2016). *Graduate College Dissertations and Theses*. Paper 581.

This Thesis is brought to you for free and open access by the Dissertations and Theses at ScholarWorks @ UVM. It has been accepted for inclusion in Graduate College Dissertations and Theses by an authorized administrator of ScholarWorks @ UVM. For more information, please contact donna.omalley@uvm.edu.

THE ROLE OF UNCERTAINTY IN CATEGORICAL
PERCEPTION UTILIZING STATISTICAL LEARNING IN
ROBOTS

A Thesis Presented

by

Nathaniel V Powell

to

The Faculty of the Graduate College

of

The University of Vermont

In Partial Fulfillment of the Requirements
for the Degree of Master of Science
Specializing in Statistics

May, 2016

Defense Date: March 28, 2016
Thesis Examination Committee:

Jeff Buzas, Ph.D, Advisor
Josh Bongard, Ph.D, Chairperson
Ruth Mickey, Ph.D
Cynthia J. Forehand, Ph.D, Dean of the Graduate college

ABSTRACT

At the heart of statistical learning lies the concept of uncertainty. Similarly, embodied agents such as robots and animals must likewise address uncertainty, as sensation is always only a partial reflection of reality. This thesis addresses the role that uncertainty can play in a central building block of intelligence: categorization. Cognitive agents are able to perform tasks like categorical perception through physical interaction (active categorical perception; ACP), or passively at a distance (distal categorical perception; DCP). It is possible that the former scaffolds the learning of the latter. However, it is unclear whether DCP indeed scaffolds ACP in humans and animals, nor how a robot could be trained to likewise learn DCP from ACP. Here we demonstrate a method for doing so which involves uncertainty: robots perform ACP when uncertain and DCP when certain. Furthermore, we demonstrate that robots trained in such a manner are more competent at categorizing novel objects than robots trained to categorize in other ways. This suggests that such a mechanism would also be useful for humans and animals, suggesting that they may be employing some version of this mechanism.

Keywords: Uncertainty; Active Categorical Perception; Evolutionary Robotics

CITATIONS

Material from this thesis has been submitted for publication to Cognitive Science Society on February 1, 2016 in the following form:

N Powell and J Bongard. (2016). Exploring uncertainty and movement in categorical perception using robots. *The 38th Annual Meeting of the Cognitive Science Society*. Philadelphia. pp. 1-8. In review.

ACKNOWLEDGEMENTS

I would like to say thank you to my parents for allowing me to pursue a master's degree in statistics and supporting me during my time at the University of Vermont. I would like to thank my friends for the emotional support they provided. Finally, I would like to thank Josh Bongard for allowing me to pursue this research in his Morphology, Evolution, and Cognition Laboratory.

TABLE OF CONTENTS

Citations	ii
Acknowledgements	iii
List of Figures	v
1 Introduction	1
1.1 Uncertainty	2
1.2 Active Categorical Perception	3
2 Methods	5
2.1 The Robot	5
2.2 The Controller	6
2.3 The Environment	9
2.4 Evolutionary Algorithm	9
2.5 Analyses	13
3 Results	14
4 Discussion	17
5 Future Work	24
5.1 Effect of the Environment	24
5.2 Shannon's entropy as a Measure of Uncertainty	25
5.3 Increasing Complexity	26
5.3.1 Increased Guess Neuron Complexity	26
5.3.2 Increased Category Complexity	27
5.4 Uncertainty in Humans	27
6 Conclusion	29

LIST OF FIGURES

2.1	Evolved behavior for a robot that interacts with its environment when uncertain to perform categorical perception. Time is tracked through panels A through D. Movement of the arm is tracked in box e. Movement of the object is tracked in box f.	6
2.2	The robot body is constructed from two arms, each composed of an upper and lower limb. The upper limbs are attached to a central body. The limbs are attached to each other and the body with four motorized joint, each of which contains an angle sensor. Thin lines represent line of sight for sensing distance. The robot can be exposed to any one of 14 objects, which are placed either in the line of sight or ‘blind’ to the robot. Objects are either large (large circles) or small (small circles).	7
2.3	The controller of the robot is instantiated as an artificial neural network. The input layer is made up of four proprioceptive neurons (P), four vision neurons (V) and a bias neuron (B). Hidden layer consists of 4 recurrent hidden neurons (H). Output layer: four motor neurons (M) and four guess neurons (G)	8
3.1	The average ability of the best 15 controllers to correctly categorize unseen objects, for the four experimental conditions tested. Error bars represent the standard error for each condition tested.	15
3.2	The distribution of the errors for each condition. The means for each condition are given in the top left corner of each subplot.	16
4.1	Performance of a robot evolved to categorize and maximize movement. Average categorization error is seen in the top graph, average change in joint angles is seen in the middle graph. Performance of robot in categorizing a (A) visible novel object and (B) a blind novel object	18
4.2	Performance of a robot evolved with both C and R to categorize. Average categorization error is seen in the top graph, average change in joint angles is seen in the middle graph, and the variance of the guess neurons is seen in the bottom graph. (A) Performance of robot in categorizing a (A) visible novel object and (B) a blind novel object	19

CHAPTER 1

INTRODUCTION

At the heart of all statistical questions lies the concept of uncertainty. It is the driving force of decision making and how to make inferences about our world through data. Uncertainty, as an internal state, is also a key component in cognition; for example how an agent knows about and learns from its environment and/or makes decisions. This implies that cognition relies on statistical processing in order to exploit the regularities of both the brain of an agent and the environment in which it is situated [11]. Statistical learning of this nature is observed not only in biological agents but artificial ones as well, and allows these agents to cope with the uncertainty of the world in which they exist.

Cognitive agents are dynamical systems, which means the the amount of uncertainty in them is inherently high, and being able to use this towards the task of categorization has not yet been employed. The most interesting interaction that leads to uncertainty is that between the body and the brain of a cognitive agent.

The embodied approach to cognitive science holds that the body is a necessary component for the acquisition of adaptive—and, ultimately, cognitive—behavior [3, 7]. Since the establishment of this approach, much work has been dedicated to investigating how the body can do so [8], and quantifying its contribution [4, 6]. A common approach for doing so is to employ robots, in which all aspects of their morphology, control structure, and task

1.1. UNCERTAINTY

environment can be observed and experimentally modified.

1.1 UNCERTAINTY

As mentioned above uncertainty is an important factor that drives inference and decision making. In all it is a measure or state that relates to conviction that something is the case. For example, I am certain that it is the case that I am sitting here writing this paper. It is a very important idea that generally that is generally overlooked in much of machine learning and psychological research. Statisticians make inferences about the world based on their level of uncertainty, which comes from forming hypotheses and performing specific kinds of tests on data and from that discerning valuable information. Hypothesis testing allows for probabilities to be utilized when determining what the correct (or likely correct) choice should be regarding the null hypothesis. For example, if a statistician or scientist were to find a p -value of <0.05 from performing a test, such as linear regression, they could say that the choice to reject the null hypothesis is correct. The null hypothesis in this case would be testing that there is no relationship between the independent and dependent variables. This hypothesis would be rejected (meaning there is a relationship) because of the certainty that the p -value gives towards making an inference.

In very much the same way, determining a level of uncertainty can help cognitive agents make decisions and about their environment at a lower level: the neural level. One example of a cognitive task that shows the importance of uncertainty is an agent categorizing objects in its environment. Agents do this by forming neural correlates and learning about the statistical regularities of their environment in order to distinguish among objects. This is as important task to understand because of the unpredictability of the world. An agent's ability to cope with unpredictability and uncertainty is in fact an aspect of being an intelligent agent. Modeling uncertainty is a difficult task, however, because determining a metric with which to measure it as an internal state of a cognitive agent has not yet been

1.2. ACTIVE CATEGORICAL PERCEPTION

discussed.

There are two methods that intelligent agents use in order to deal with the problem of categorization: Active Categorical Perception and Distal Categorical Perception. Active Categorical perception allows for the agent to manipulate the object as means of categorization. Distal perception allows the agent to visually categorizes the object. In the next section we will explain further these two methods and propose how uncertainty might be able to reconcile how these two methods of categorization might work together.

1.2 ACTIVE CATEGORICAL PERCEPTION

A common cognitive skill investigated from an embodied perspective is that of categorical perception: how does an agent make use of its body to generate the requisite stimuli to learn appropriate categories? Categorical perception is one of many ways in which cognitive agents are able to exploit the statistical regularities of their environment. It allows information about the environment to be encoded in such a way that it is recognized at later times. This skill has been under scrutiny for centuries, most notably from the philosopher John Locke in the form of Malynes's Problem. This problem asks whether someone who was born blind is able to distinguish between objects if he or she learned the categories by touch (ACP) and miraculously gained sight and is tested based on vision (DCP). How would such an individual be able to distinguish between the objects, and could it be done visually having never had such experiences before?

Initially, Beer evolved minimally cognitive agents to achieve this 'active' form of categorical perception (ACP) [1]: the agents used their interaction with their environments to reduce intracategorical differences and magnify intercategory ones. Subsequent studies explored this phenomenon using more complex robot morphologies [2, 12, 13] which physically manipulate the objects to be categorized.

However, sophisticated cognitive agents typically employ distal categorical perception

1.2. ACTIVE CATEGORICAL PERCEPTION

(DCP)—categorizing an object by sight and/or sound from a distance—as it has obvious advantages over ACP, such as rapidity, avoidance of potentially dangerous contact, and success even when physical contact is not possible. This raises questions regarding how animals learn (and how robots should learn) DCP, ACP, and how to switch between them. One hypothesis is that ACP scaffolds the learning of DCP: interactions with objects can structure perception in such a way as to facilitate learning of non-embodied skills [5].

The question remains, however, as to the conditions under which ACP or DCP should be employed. Here it is hypothesized that such switching should be modulated by uncertainty: unfamiliar stimuli should trigger internal uncertainty, which in turn should trigger appropriate action resulting in ACP, which, finally, provides scaffolding for the learning of DCP when next presented with this object. Over a lifetime, this should result in an agent that exhibits increasing instances of DCP and fewer of ACP.

This experiment demonstrate the usefulness of this particular mechanism by training through evolutionary computation simulated robots to perform ACP when uncertain and DCP when certain. We show that, when exposed to novel stimuli, these robots categorize more efficiently than robots trained to categorize in other ways.

CHAPTER 2

METHODS

A series of experiments were conducted in which a simulated, yet embodied, robot attempted to categorize objects in its environment (Fig. 2.1). It is embodied in the sense that, despite its virtual surroundings, actions it can take impacts its environment, and it immediately detects the sensory repercussions of those effects. This leads to an interesting body-environment feedback loop.

An evolutionary algorithm was used to train the robot to perform ACP, DCP, or a combination of the two when exposed to a number of objects. The robot's task was to correctly categorize large objects as large, and small objects as small. When training concludes, the best robot's categorization abilities were tested by exposing it to objects placed in novel locations, and its categorization error in those situations was measured.

2.1 THE ROBOT

Each robot is composed of a three-dimensional mechanical structure—its body—and a neural model that mediates interactions with the environment: its controller.

The robot's body was constructed from four equal length cylinders and a small central body constructed from a rectangular solid. The arms adjacent to the main body are con-

2.2. THE CONTROLLER

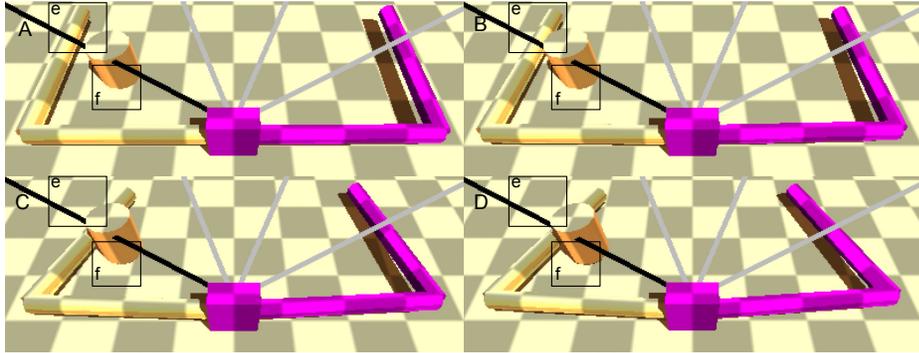


Figure 2.1: Evolved behavior for a robot that interacts with its environment when uncertain to perform categorical perception. Time is tracked through panels A through D. Movement of the arm is tracked in box e. Movement of the object is tracked in box f.

connected by supporting motorized joints to it, as are the forearms connected to the upper arms (Fig. 2.2). There are a total of four motorized joints, yielding four mechanical degrees of freedom. Each of the motorized joints enable the connected object to rotate relative to one another through the robot's coronal plane, which, given its morphology, corresponds to the horizontal plane. This results in the arms flexing in horizontally toward the main body and extending outward from the main body, also horizontally. The arms were placed at a particular height so that they do not end up breaking the visual beams emanating from the robot (black and gray lines in Fig. 2.1 and described below).

2.2 THE CONTROLLER

The controller of the robot is instantiated as a partially-recurrent artificial neural network. There are three layers that make up the neural network: the input layer, a hidden layer, and an output layer (Fig. 2.3).

The input layer consists of two types of sensory neurons: vision neurons (V) and proprioceptive neurons (P). At each time during which a robot is simulated, the angles of the four motorized joints are computed, normalized to real values in the range $[-1, +1]$, and supplied to the four proprioceptive neurons. Likewise, four visual beams are sent out from

2.2. THE CONTROLLER

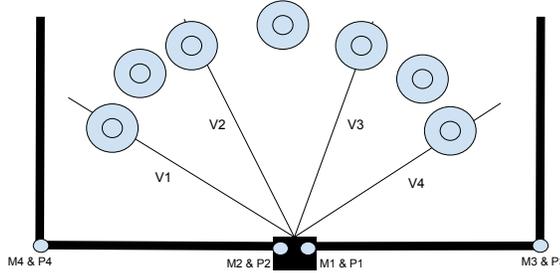


Figure 2.2: The robot body is constructed from two arms, each composed of an upper and lower limb. The upper limbs are attached to a central body. The limbs are attached to each other and the body with four motorized joint, each of which contains an angle sensor. Thin lines represent line of sight for sensing distance. The robot can be exposed to any one of 14 objects, which are placed either in the line of sight or ‘blind’ to the robot. Objects are either large (large circles) or small (small circles).

the main body such that they span the range $[-60^\circ, +60^\circ]$ in front of the robot. The angles between each pair of neighboring beams was set to 40° . While the robot may move its arms, it cannot move the visual beams. However, the beams may be broken by coming into contact with an external object. The length of each beam at each time step is computed and scaled to a real value in $[0, 1]$ such that zero indicates the beam is unbroken, while one indicates that an object can come into contact with the base of the beam.

The hidden layer consists of four fully recurrent hidden neurons (H): each hidden neuron receives input from each of the sensors (in addition a fixed-output bias neuron B) as well as values from the other hidden neurons, including itself. The new value of the i th hidden neuron h_i is computed as

$$h_i = \tanh\left(\sum_{j=1}^4 s_j w_{ji} + w_{bi} + \sum_{k=1}^4 h_k w_{ki}\right) \quad (2.1)$$

where s_j is the value of the j th sensor neuron, w_{ji} is the weight of the synapse connecting the j th sensor neuron to the i th hidden neuron ($w_{ji} \in [-1, +1]$), w_{bi} is the weight of the synapse connecting the bias neuron (value clamped to one) to the i th hidden neuron, h_k

2.2. THE CONTROLLER

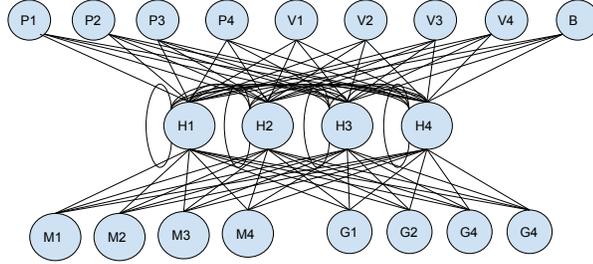


Figure 2.3: The controller of the robot is instantiated as an artificial neural network. The input layer is made up of four proprioceptive neurons (P), four vision neurons (V) and a bias neuron (B). Hidden layer consists of 4 recurrent hidden neurons (H). Output layer: four motor neurons (M) and four guess neurons (G)

is the value of the k th hidden neuron, w_{ki} is the weight of the synapse connecting the k th hidden neuron to the i th hidden neuron, and $\tanh(x)$ brings the hidden neuron values back into the range $[-1, 1]$.

The output layer is comprised of two different types of neurons: motor neurons (M) and guess neurons (G). The value of the i th output neuron is computed at each time step using

$$o_i = \tanh\left(\sum_{j=1}^4 h_j w_{ji}\right) \quad (2.2)$$

regardless of whether it is a motor or guess neuron.

The value of each of the four motor neurons is scaled to the range $[-45^\circ, +45^\circ]$ and then supplied, as a desired angle, to each of the four joints. A proportional-derivative (PD) controller is effected by supplying torque to the joint proportional to the difference between the current angle and the desired angle. The outputs arriving at the guess neurons are employed by the robot to perform categorical perception and are described in more detail below.

2.3 THE ENVIRONMENT

When evaluated, a robot is equipped with a neural network, populated with a particular set of synaptic weights as described above. The robot is then exposed to one of 14 objects, as shown in Fig. 2.3. There are seven possible positions. At each location there are two possible types of objects that may appear: a cylinder with a large or small radius. Objects are placed in such a way that they are either blind to the robot or in its direct line of sight.

2.4 EVOLUTIONARY ALGORITHM

An evolutionary algorithm was used to train neural networks in the robot as described above. For each evolutionary trial, seven objects from the total set of 14 were chosen as random and fixed as the training set for that run. In a different evolutionary run, seven different objects may be chosen. At the outset of the trial, an initial population of 20 random neural networks were created. Each of these neural network contained random synaptic weights drawn from $[-1, +1]$ with a uniform distribution. Each neural network was then evaluated on the robot seven times, in the presence of one of the seven objects chosen for that run. During each of these seven evaluation periods, the robot was allowed to move for 25 time steps in the simulator.

This population of neural networks was then evolved for 100 generations using a common evolutionary algorithm that balances increasing fitness over time while also maintaining genetic diversity in the population [9].

The robots in this experiment were evaluated against four different fitness functions, leading to four experimental conditions. In the first condition, robots were evolved simply to categorize correctly (C). In the second condition they were evolved to categorize correctly while minimizing movement (CnM). In the third condition they were evolved to categorize

2.4. EVOLUTIONARY ALGORITHM

correctly while maximizing movement (CM). In the fourth and final condition they were evolved to categorize correctly, and to do so by moving when uncertain about the object’s category and remaining still when certain about the object’s category. This condition was referred to as CR , as the robot should establish a correlation (R) between movement and uncertainty. Fifteen independent evolutionary runs, each starting with a different randomly chosen set of seven objects and 20 random ANNs, were run for each of the four conditions.

The fitness functions for these conditions were constructed from combinations of the following terms

$$C = \sum_{t=1}^T \sum_{o=1}^O \sum_{i=1}^G (g_{oi}^{(t)} - s_o)^2 / TOG \quad (2.3)$$

$$K = \sum_{t=1}^T \sum_{o=1}^O \sum_{m=1}^M |(a_{om}^{(t)} - a_{om}^{(t-1)})| / TOM \quad (2.4)$$

$$R = \sum_{o=1}^O \text{Corr}(\vec{G}_o, \vec{m}_o) / O \quad (2.5)$$

$$\vec{G}_o = [\sigma(\vec{g}_o^{(2)}), \sigma(\vec{g}_o^{(3)}), \dots, \sigma(\vec{g}_o^{(25)})] \quad (2.6)$$

$$\vec{m}_o = [\sum_{m=1}^M |a_{om}^{(2)} - a_{om}^{(1)}|, \dots, \sum_{m=1}^M |a_{om}^{(25)} - a_{om}^{(24)}|] \quad (2.7)$$

where

- C denotes how well a given neural network categorizes, averaged over all $T = 25$ time steps, $O=7$ training objects, and $G=4$ guess neurons ($C = 0$ indicates perfect categorization and $C = 1$ indicates the worst possible categorization);
- K denotes the average amount of motion over all T time steps, O training objects, and $M = 4$ motors;
- R denotes the amount of correlation (Corr) between uncertainty (\vec{G}_o) and amount of movement (\vec{m}_o), averaged over all O objects ($R = 1$ indicates the robot moves maximally whenever uncertain and minimally whenever certain);

2.4. EVOLUTIONARY ALGORITHM

- $g_{oi}^{(t)}$ represents the output of the i th guess neuron during the t th time step of exposure to the o th object;
- s_o represents the size of the o th object (small object=-0.5, large object=+0.5);
- $a_{om}^{(t)}$ represents the angle of the m th motorized joint during the t th time step of exposure to the o th object;
- \vec{G}_o represents a vector containing the uncertainties of the ANN when exposed to the o th object, at each time step of the exposure (with the exception of the first time step);
- \vec{m}_o represents a vector containing the amount that the robot moved when exposed to the o th object, at each time step of the exposure (with the exception of the first time step); and
- $\vec{g}_o^{(t)}$ represents a vector containing the values of the four guess neurons generated during the t th time step when exposed to the o th object, and $\sigma(\vec{g}_o^{(2)})$ represents the variance within that vector.

Prediction variance and uncertainty.

Here, we employ the variance among the values of the guess neurons to denote a controller’s uncertainty about the current object’s category. In the machine learning literature, prediction variance is often employed as a proxy for uncertainty [10]. This is because, as long as individual units in a predictive model (here, the guess neurons) are independent, they are likely only to converge on the same prediction when that prediction is correct. This is not unlike a group of people with very different backgrounds generating diverse—and thus mostly wrong—answers to questions that touch on an area of their mutual ignorance, but who only generate similar responses when the question touches on an area of their common

2.4. EVOLUTIONARY ALGORITHM

knowledge. The guess neurons here are independent because each guess neuron has its own synaptic weights connecting it to the input layer.

Condition *C*:

In the first condition, robots were only evolved to categorize, regardless of the amount or type of movement they employed to do so. This was accomplished by evolving robots that maximized the fitness function

$$F_C = 1/(1+C). \quad (2.8)$$

Condition *CnM*:

In the second condition, distal categorical perception was explicitly favored by evolving robots that successfully categorize while also minimizing movement:

$$F_{CnM} = \left(\frac{1}{1+C}\right)\left(\frac{1}{1+M}\right). \quad (2.9)$$

Condition *CM*:

In this third condition, active categorical perception was explicitly favored by evolving robots to successfully categorize while maximizing movement:

$$F_{CM} = M/(1+C). \quad (2.10)$$

Condition *CR*:

Finally, robots were evolved in the fourth condition to employ distal categorical perception when uncertain as to the object's size and to employ active categorical perception

2.5. ANALYSES

when they were sure. This was accomplished using

$$F_{CR} = R/(1+C). \quad (2.11)$$

All of these functions allow the robot to exploit the regularities in it's environment in different ways through different aspects of the neural network and were minimized over evolutionary time.

2.5 ANALYSES

At the end of each evolutionary run the best individual was tested on the 7 remaining objects (see above). The measure that was collected was the error due to categorization (equation C). This was collected for 15 independent individuals from each evolutionary condition. A Mann-Whitney U test was performed to determine the differences between conditions (See Results). A non-parametric test was performed so that no assumptions about the distribution of the categorization error was made.

CHAPTER 3

RESULTS

The experimental settings are summarized in table 3.1. At the termination of each run, for each condition, the robot with the best fitness is extracted from the population. Each of these controllers is then re-instantiated in the robot and evaluated a further seven times, in the presence of the seven objects that the controller did not see during training. We employed C (Eqn. 2.3) to compute the robot’s average ability to categorize these $O = 7$ novel objects. A controller that obtains lower values of C when exposed to these novel objects is thus exhibiting a better ability to generalize its ability to categorize, compared to another controller with a higher value of C . This is not to be confused with the condition for evolution C .

Fig. 3.1 reports the average generalization abilities of the 15 evolved controllers extracted from the four conditions. Testing for significance was computed using multiple

Table 3.1: Summary of experimental settings.

Experimental conditions:	4 (C ; $C \sim M$; CM ; CR)
Trials per condition:	15
Generations per trial:	100
Size of robot population:	20
Evaluation period of each robot:	25 time-steps
Mutation rate:	25%

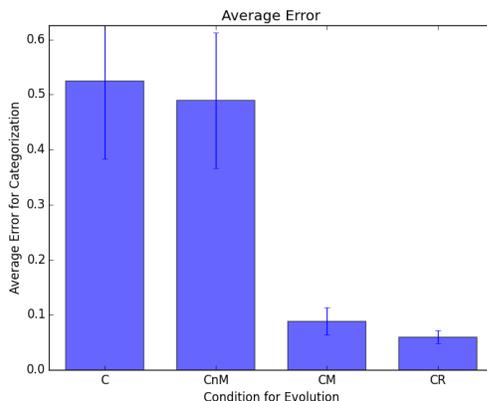


Figure 3.1: The average ability of the best 15 controllers to correctly categorize unseen objects, for the four experimental conditions tested. Error bars represent the standard error for each condition tested.

Mann-Whitney U tests. The distribution and average of the error due to categorization can be seen for each condition in figure 3.2. The tests looked for significant differences between the CR condition and the remaining three conditions (C , CnM , and CM). Significant P-values were found in each comparison. After correcting for multiple comparisons using the Bonferroni method of adjustment, the P-values found between CR and C , CR and CnM , and CR and CM are respectively 1.52×10^{-5} , 1.52×10^{-5} , and 1.18×10^{-3} . Significant P-values indicate that the relative average amount of categorization error performed by individuals in the CR condition is significantly lower than the error of individuals in the remaining three conditions. Therefore, it can be concluded that individuals evolved under the CR condition are better suited for tasks involving categorization of objects than individuals from the other three evolutionary conditions.

Table 3.2: Summary of Median and Spread (Percentile)

Condition	Median	25%	75%
C	0.491	0.415	0.561
CnM	0.452	0.403	0.516
CM	0.087	0.075	0.093
CR	0.065	0.049	0.069

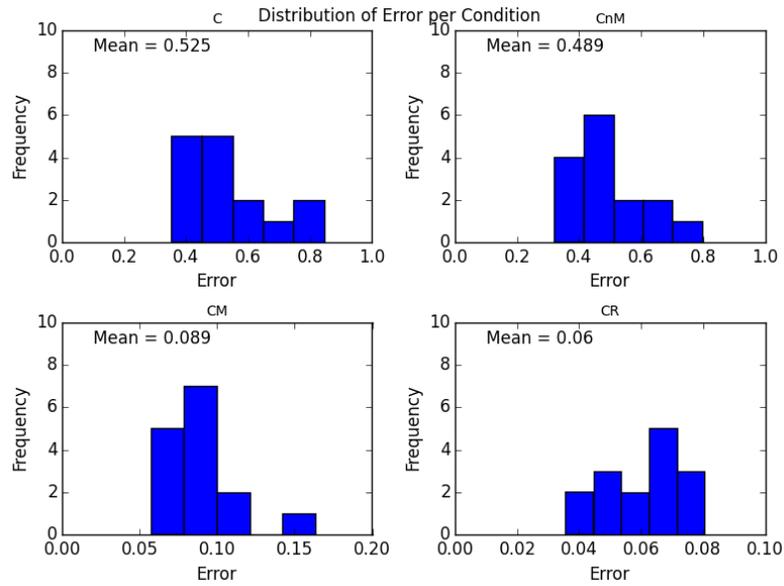


Figure 3.2: The distribution of the errors for each condition. The means for each condition are given in the top left corner of each subplot.

CHAPTER 4

DISCUSSION

Why is C worse than CR ?

The controllers evolved in the C condition performed worse than those evolved for CR for several reasons. Such controllers may have generated little to no motion, enabling rapid and successful categorization of seen objects, while sacrificing the ability to categorize unseen objects. This may have led to overfitting such that novel, seen objects were categorized incorrectly. Conversely, controllers may have evolved that cause much motion. Such a strategy might cause correct, instantaneous categorization which is subsequently lost when the robot comes into contact with the object.

Why is CnM worse than CR ?

Controllers evolved in the CnM are likely incentivized to memorize the categories of seen objects, and ignore the categories of blind objects. This results in overfitting: the robot will not only poorly categorize novel, unseen objects, but novel, seen objects as well. In other words, the robots are deprived of the ability to reduce spurious differences between intracategory objects through motion.

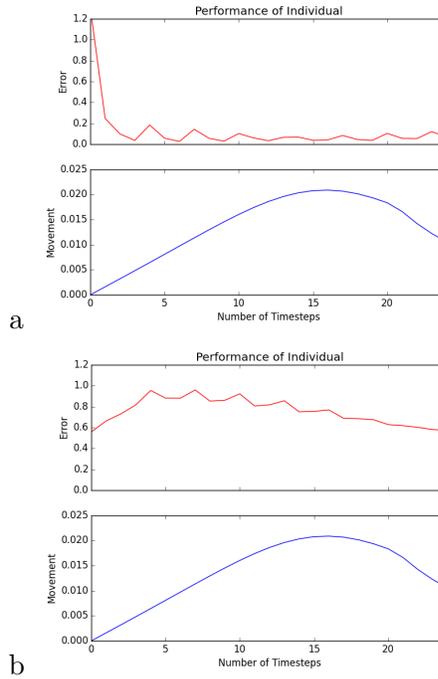


Figure 4.1: Performance of a robot evolved to categorize and maximize movement. Average categorization error is seen in the top graph, average change in joint angles is seen in the middle graph. Performance of robot in categorizing a (A) visible novel object and (B) a blind novel object

Why is *CM* worse than *CR*?

Conversely, controllers evolved in the *CM* condition may suffer from two disadvantages compared to controllers evolved using *CR*. Controllers from *CM* may not be able to afford to hold still when they are certain of a novel object’s class. Moreover, they may have to exhibit so much motion that they end up magnifying spurious intracategory differences, rather than being free to generate less—yet appropriate—movement that reduces those differences. A possible example of one (or both) of these disadvantages is shown in Fig. 4.1. Even though this controller can reduce category error when in the presence of a visible novel object by perhaps moving in a way that does not move this object (Fig. 4.1a), error increases (and then reduces to the original high level) when exposed to a novel blind object (Fig. 4.1b).

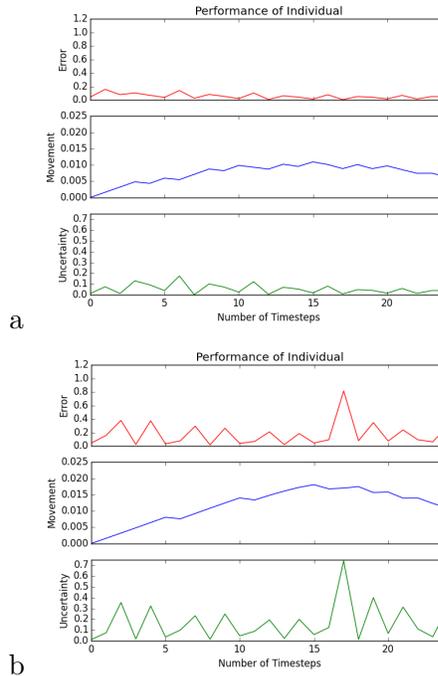


Figure 4.2: Performance of a robot evolved with both C and R to categorize. Average categorization error is seen in the top graph, average change in joint angles is seen in the middle graph, and the variance of the guess neurons is seen in the bottom graph. (A) Performance of robot in categorizing a (A) visible novel object and (B) a blind novel object

How did CR succeed?

Controllers produced in the CR condition presumably outperformed the controllers produced by the other three conditions because they can better employ ACP, or DCP, when that form of categorization is most appropriate. When the robot is certain about an object's category, it should employ DCP without moving: the robot does not need to wait for physical contact with the object to categorize the object. Moreover, if it does contact the object, the resulting sensor data may affect the guess neurons and thus draw the robot away from the already correctly-predicted category. Conversely, when the robot is uncertain as to the object's category, it should initiate movement as rapidly as possible: that is, exactly at the moment that it is uncertain.

Another interesting result from the CR condition is that when a CR -generated con-

troller’s uncertainty is high, it also tends to exhibit higher category error (exemplified by the spike in error and uncertainty in Fig. 4.2). At such times, CR-generated controllers are encouraged to move as much as possible. However, once in contact with the object and the amount of uncertainty starts to decrease, less motion is necessary: the robot is free to perform whatever actions are appropriate to reduce intracategory differences. In contrast, *CM* generated robots are more restricted in the kinds of actions they can employ to reduce these differences: they must generate high-magnitude movements.

What probabilistic properties makes uncertainty important in statistical learning?

Over evolutionary time a number of processes occur that change the internal content of each individuals’ neural network. These changes allow for the controller to become better at categorization, decreasing the overall error from the guess neurons (*CE*). The following expression (4.1) shows that the limit of categorization error converges in probability towards 1 as evolutionary runs (*E*) and initial population (*N*) increase.

$$\lim_{N,E \rightarrow \infty} P(CE = 0) = 1 \tag{4.1}$$

Another such limit that arose from evolution is from the variance of the guess neurons, or the level of uncertainty. The form of this can be seen in equation 4.2. This shows that the limit converges in probability towards 1, or that the probability of the variance of the guess neurons (uncertainty) being 0 converges in probability towards 1. This leads to an increased likelihood of the robot being certain about what category to place the object.

$$\lim_{N,E \rightarrow \infty} P(Var(\vec{g}) = 0) = 1 \tag{4.2}$$

These arise out of the correlation that forms between categorization error and uncertainty. In one example (See figure 4.2b) the error of categorization and uncertainty are correlated ($R = 0.4626$ with $p < 0.001$). This is for one particular object and one particular robot. This, however, is evidence that this relationship does arise over evolutionary time. This needs to be elaborated and confirmed further, but it appears that when the parameters are increased this relationship should become stronger as well (See sources of error). As the number of evolutionary runs increases a correlation between categorization and uncertainty arises. This is evidence that when a robot is certain of an object it is more likely to have lower categorization error. After more evolutionary runs it is assumed that these correlations will become more pronounced and will lead to significant results showing the relationship between categorization error and level of uncertainty.

This result means that as the number of evolutionary runs increases the probability of a robot being able to decrease its uncertainty converges towards a probability 1. This of course is also due to the correlation between movement and uncertainty (i.e. as the robot moves more it is more likely to put the object in a position that it is certain or more familiar with). This correlation gives the robot access to its environment, which allows it to reduce it's own uncertainty (See sections on why *CR* did better above).

The equation referring to C (equation 2.1) is converging to zero, or some minimum that is constrained by the physical body of the robot. All conditions utilized this function (equations 2.8-11) decreased categorization error in a manner similar to equation 4.1, and those which used R as well increased fitness even further. These converge due to the optimization process that occurs over evolutionary time towards some fitness goal.

Uncertainty as a Markov Decision Process

A Markov Decision process is a mathematical framework that allows for a decision about the action state of the robot to be made based on the current internal state of the robot [11].

These processes in and of themselves address the problem of uncertainty in robotic motion and cognition, but not in such a way that uses it directly. Here uncertainty is a potential scaffold between ACP and DCP, which allows for a decision to be made based on the internal state of the neural network in which the robot is experiencing.

One argument for why the *CR* condition potentially employs this kind of algorithm is that it is based in a form of reinforcement learning(Q-learning). Robot controllers are reinforced by surviving to the next generation, thus correlating the neural network with the state that gives the optimum fitness. Reinforcing the relationship between uncertainty and movement allows for a change in action state by the level of uncertainty within the network, which in turn allows decisions to be made based off uncertainty. The decision to either actively categorize or passively categorize being subject to change due to the level of uncertainty. This is offered as a possible explanation for how uncertainty can be viewed in this paradigm, but not as a definitive fact.

Possible Sources of Error.

One potential source of error is that only 100 generations of evolution were performed with a relatively small initial population of 20 individuals. This may not have allowed for significant optimization to occur in all conditions. The very short evaluation period may also be a source of error, because the *CM* and *C* conditions may allow for the discovery of exaggerated movements that yet, given enough time, eventually reduce intracategory differences. More generally, longer evaluation periods will allow for a greater range of movements that may then better help clarify the relationship between movement, categorization, and uncertainty.

One of the crucial aspects to this experiment is that the guesses that the *G* neurons are outputting are independent of one another. If this is not the case, then this would be a major source of error. In order to make sure of this in the future steps can be taken that

measure the association between the guesses.

Another potential source for error is due to constraint on the amount of optimization that can occur. This is due to the physics of the world in which the robot lives, as well as its body. For example no matter how fast it can move to change the position of an object it can't do it instantaneously(the maximum amount of force to be exerted by the joints is static), so therefore uncertainty cannot be 0 and categorization error cannot be 0. But these values can still get close. This means that they can converge as N and E get sufficiently large, but they will never actually reach the ideal fitness.

CHAPTER 5

FUTURE WORK

In this chapter four avenues of future work are described.

5.1 EFFECT OF THE ENVIRONMENT

The environment plays a crucial role in the agent's ability to categorize objects. Because the objects and locations are chosen at random for testing there is a chance that a trial will consist of only visible object or only blind objects (all blind objects is impossible in this experiment due to the fact that there are less blind objects than visible objects).

In order to test whether the location of the objects in the environment has an effect on the agents' ability to categorize a study could be formed that would train these same 4 conditions but only on either blind or visible objects. This would allow for analysis on controlling for environmental condition, which would also have something to say about the affect of uncertain and certainty. This being because the blind runs would be cause the agent's to be more uncertain than the visible objects. This would also be able to be done on small vs large objects.

Another aspect that would be interesting to control for in this case would be the effect of average movement on their ability to categorize in different environments. There should

5.2. SHANNON'S ENTROPY AS A MEASURE OF UNCERTAINTY

be an interaction between movement and being tested on blind runs if the hypothesis is correct. A robot should move more on blind runs due to its increased level of uncertainty, and this should be seen more so when controlling for objects placement in the environment.

5.2 SHANNON'S ENTROPY AS A MEASURE OF UNCERTAINTY

Shannon's entropy is another measure of uncertainty in a way similar to how variance of the guess neurons has let us calculate the level of uncertainty in the agent. It is more formally a way that calculates how much information is being encoded by the system. The definition of Shannon's Entropy is as follows:

$$H(P) = E[-\log_2 p(x)] = -\sum p(x) \log_2 p(x) \quad (5.1)$$

This measure will be used in a manner similar to what was done in the current experiment, and would be substituted for $Var(\vec{g})$ as a measure of uncertainty. $P(x)$ in this equation representing the probability that the guess neuron takes a given value.

If this allows better fitness to arise in individuals over evolutionary time then we would be able to say that entropy is a better measure of uncertainty than the variance of the guess neurons. Entropy is of common use in the field of robotics, but it has not been used in such a way that it measure the internal state of uncertainty within an embodied agent like what is proposed for the future of this project.

5.3 INCREASING COMPLEXITY

5.3.1 INCREASED GUESS NEURON COMPLEXITY

One avenue that would be worth pursuing is increasing the complexity of the guess neuron vector. Increasing the number of guess neurons would allow for many more guesses at one time than the current study, and would affect how the level of uncertainty changes. This would allow for study of this, and how uncertainty plays a role in the relationship between ACP and DCP as the complexity of the guess neurons increases.

In order to test the relationship that exists between the number of guess neurons and uncertainty's role in the categorical task from the present experiment several different experimental conditions would need to exist. One such condition for training would be manipulating the number of guess neurons in the neural network. In the present study there are 4 G neurons, in order to determine if there is a relationship between this number of categorization other numbers of guess neurons would have to be studied. An example of the number of neurons that could be tests are $2G$, $3G$, $5G$, $10G$, $20G$. This would allow for a trend to emerge if one exists.

The results from these different runs of the current experiment with different length guess neuron vectors would then be compared with one another in order to determine the relationship. As the task gets more complicated it is hypothesized that individuals trained with some optimum number of guess neurons would have higher fitness. This is because if uncertainty is low there would be more of a chance that the guesses are correct. This would also be of benefit in situations or tasks that require ACP and DCP. This could also lead to the agent to be able to categorize based on more than one feature of the object in question.

5.4. UNCERTAINTY IN HUMANS

5.3.2 INCREASED CATEGORY COMPLEXITY

Increasing the complexity of the task for categorization would involve creating either multi-dimensional objects (i.e. objects with more than one aspect involved in categorization) or more than 2 sizes of objects.

If the current study were to generalize to multi-dimensional objects, it is hypothesized that creating a matrix of guess neurons would be sufficient for categorization (See Matrix below). This would allow for each column to represent a feature of the object in which the controller can be either certain or uncertain. An example of an object fitting this description could be the same objects from this experiment but their weight also is a feature that can be used to categorize it (i.e. heavy-large, light-large, heavy-small, light-small). This could, in theory, be expanded to include many other features. One important aspect that would have to be kept in mind is how the different vectors of guess neurons are correlated with one another, and what that might tell about the features of the object and vice versa.

$$G = \begin{bmatrix} g_{11} & g_{12} & g_{13} & g_{14} \\ g_{21} & g_{22} & g_{23} & g_{24} \\ g_{31} & g_{32} & g_{33} & g_{34} \\ g_{41} & g_{42} & g_{43} & g_{44} \end{bmatrix}$$

Increasing the complexity of the task and neural network would help to generalize this idea to concepts more closely representing higher level cognition and would give an embodied explanation for the ways categorization is instantiated in cognitive agents.

5.4 UNCERTAINTY IN HUMANS

In order to test this hypothesis in a more naturalistic setting, the affect of uncertainty on the categorical perception in humans can be studied. This can be done by observing

5.4. UNCERTAINTY IN HUMANS

their behavior during categorization tasks at times of certainty and uncertainty.

The potential experiment that could be done in order to test this consists of having human participants categorize how much liquid is in a container. They would be filled with either more than half or less than half of the amount of water it takes to fill the entire container. In order to illicit uncertainty in participants the containers will be either opaque or translucent. Opaque containers will illicit the highest level of uncertainty because participants would not be able to visually inspect the contents.

This draws many parallels from the current experiment. The level of uncertainty should bridge the gap between ACP and DCP in a similar manner if that is indeed how humans also categorize objects. Though there wouldn't be direct access to the neural activity in the participants, the observed behaviors would be enough to safely determine whether or not this hypothesis was correctly being studied. For instance, if ACP and DCP were bridged in a similar manner, then the number of instances of ACP (manipulating the object) would be greater in high uncertainty trials. Trials with high uncertainty being those with opaque containers. This would also make stronger the claim that uncertainty is fundamental in categorical perception with statistical learning.

These follow-up questions and methods would help to solidify the idea that uncertainty plays a crucial role in categorization. This is especially true if generalization to not only more complex tasks and objects, but also naturalistic examples like categorization in humans, is possible.

CHAPTER 6

CONCLUSION

Here it is shown that it is possible to explicitly train robots to exhibit ACP when uncertain and DCP when certain, and that such robots outperform other robots trained to perform ACP or DCP at will but without the ability to do so based on uncertainty; trained to always perform DCP (categorize without moving); or trained to always perform ACP (categorize via movement). Further, this approach does not require us to dictate how the robot should interact with its environment; it is free to discover its own strategies for reducing intracategory differences through physical interaction.

This work thus helps to clarify the relationship between three competencies necessary for any embodied agent that wishes to categorize rapidly and successfully: the ability to categorize, the ability to interact with the world, and the ability to determine when such interactions are and are not necessary for categorization.

In future work hopes to employ more sophisticated optimization methods, such as multiobjective optimization, which enable better tradeoffs between multiple fitness terms. We would also like to investigate the kinds of physical interactions generated by the controllers, and how such actions differ based on different circumstances. Further, we wish to investigate how the successful controllers described here gradually acquired DCP, ACP, and/or whether the acquisition of one scaffolded the subsequent acquisition of the other.

Another issue that could be addressed is whether variance of the guess neurons is the only way to measure the internal state of uncertainty. There are other measures that could work, and determining if they work as well as variance as a measure of uncertainty would be interesting. One such measure of uncertainty is Shannon's Entropy.

On top of these issues, it also is recommended that the idea presented here is scaled up. As was discussed previously, increasing the complexity of the task or the way in which guesses are given about an object can amplify the results found here. This would allow for a more solid foundation to the extent to which uncertainty plays a role in categorical perception and statistical learning.

Finally, we wish to explore whether such dynamics relate to how humans acquire these competencies. This would be done by studying situations that causes uncertainty in humans and see how the task would be resolved. For example, asking individuals to categorize the amount of liquid (mostly full or mostly empty) within containers that are opaque and translucent. Not being able to see into the container would presumably cause a state of uncertainty in the participant. If they bridge the gap between ACP and DCP in a similar manner without explicitly being told to we would be able to state that uncertainty does hold importance in categorical perception.

BIBLIOGRAPHY

- [1] Randall D Beer. The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4):209–243, 2003.
- [2] Josh Bongard. The utility of evolving simulated robot morphology increases with task complexity for object manipulation. *Artificial life*, 16(3):201–223, 2010.
- [3] Rodney A Brooks. Elephants don’t play chess. *Robotics and Autonomous Systems*, 6(1):3–15, 1990.
- [4] Helmut Hauser, Auke J Ijspeert, Rudolf M Fuchslin, Rolf Pfeifer, and Wolfgang Maass. Towards a theoretical foundation for morphological computation with compliant bodies. *Biological Cybernetics*, 105(5-6):355–370, 2011.
- [5] Max Lungarella and Olaf Sporns. Information self-structuring: Key principle for learning and development. In *Procs of the Intl Conf on Development and Learning*, pages 25–30. IEEE, 2005.
- [6] Chandana Paul. Morphological computation: A basis for the analysis of morphology and control requirements. *Robotics and Autonomous Systems*, 54(8):619–630, 2006.
- [7] Rolf Pfeifer and Josh Bongard. *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT press, 2006.
- [8] Rolf Pfeifer, Fumiya Iida, and Max Lungarella. Cognition from the bottom up: on biological inspiration, body morphology, and soft materials. *Trends in Cognitive Sciences*, 18(8):404–413, 2014.
- [9] Michael Schmidt and Hod Lipson. Age-fitness pareto optimization. In *Genetic Programming Theory and Practice VIII*, pages 129–146. Springer, 2011.
- [10] H Sebastian Seung, Manfred Opper, and Haim Sompolinsky. Query by committee. In *Procs of the fifth annual workshop on computational learning theory*, pages 287–294. ACM, 1992.
- [11] Burgard Wolfram Thrun, Sebastian and Dieter Fox. *Probabilistic Robotics*. MIT press, 2005.
- [12] Elio Tuci, Gianluca Massera, and Stefano Nolfi. Active categorical perception of object shapes in a simulated anthropomorphic robotic arm. *IEEE Transactions on Evolutionary Computation*, 14(6):885–899, 2010.

BIBLIOGRAPHY

- [13] Karol Zieba and Josh Bongard. An embodied approach for evolving robust visual classifiers. In *Procs of the Genetic and Evolutionary Computation Conf*, pages 201–208. ACM, 2015.