Evolutionary Algorithms using Combined Objective and Novelty Search in Unconstrained Space

Jeff Novak and Kyle Richmond-Crosset

December 2017

Abstract

Novelty search, which incentivizes the evolution of novel rather than fit behaviors, has shown to be effective at solving deceptive problems in constrained behavior spaces. However, in large behavior spaces with a required goal, novelty search can struggle to evolve solutions in a timely manner. We combine novelty search with an objective search to solve deceptive problems in an unconstrained environment. Our experiments add to existing research about the limitations of novelty search and show that a blended objective + novelty search is more efficient than either search alone in unbounded, deceptive behavior spaces.

1 Introduction

Evolutionary robotics aims to mimic natural systems to evolve complex behaviors. Almost all evolutionary algorithms use a goal-based fitness function that differentiates between individuals based on their performance towards a predetermined objective. Individuals that perform the best relative to the fitness function are favored in future generations. In situations in which a goal can be clearly defined and reached by evolving individuals with a steady fitness gradient, objective search has been proven to be efficient and effective.

However, evolutionary algorithms that use objective search struggle to perform well in deceptive situations where the optimal behavior or solution cannot be found by evolving behaviors that have linearly improving fitness [6]. Objective search often converges to local optima in deceptive environments and as a result provides poor solutions to these problems even after many generations [6].

Novelty search, developed by Stanley and Lehman, avoids the deceptiveness problem altogether by ignoring any goal-oriented objective [6]. Instead, novelty is a goal-independent search that rewards novel behavior. Over many generations, this eliminates the problem of local optima, as individuals that produce similar behaviors as past individuals receive low fitness scores using a novelty metric that replaces the traditional fitness function.

In deceptive, constrained environments, novelty search has been found to be more efficient at finding a solution than objective search. In larger behavior spaces, novelty search struggles to evolve solutions in a timely manner because there is no incentive to evolve solutions that reach the goal [2].

The proposed solution, novelty + objective search, is a blended approach that evolves a population based on a fitness function that utilizes both a novelty metric and an objective metric. This combination is perhaps the most logical solution for deceptive environments with large behavior spaces, in which objective search struggles with deceptiveness and novelty search struggles with the size of the behavior space.

Previous research has used a blended search in deceptive, constrained environments [7]. We extend that research to maze environments of various deceptiveness that are unconstrained by outer boundaries. Our experiment provides useful insights about the efficiency of novelty + objective search, and our results extend Stanley and Lehman's previous research about augmenting novelty search [6].

2 Background

2.1 NeuroEvolution of Augmenting Topologies (NEAT)

Our experiments use NEAT, a genetic algorithm that evolves neural networks, as the platform on which objective search, novelty search, and objective + novelty search are tested. NEAT was developed in 2004 by Stanley and Miikkulainen, who argue in their paper that the incremental evolution of neural network topologies can be more successful than fixed topologies at solving complex problems [9]. However, the continuous evolution of the structure of a neural network creates several challenges: meaningfully crossing over topologically different networks, preserving networks with unique topologies until their weights can be optimized, and ensuring topologies are not unnecessarily complex [9].

Stanley and Miikkulainen solve these problems in NEAT using several novel structures. First, NEAT uses historical markers to align genes before crossover, which allows for relatively straightforward recombination and creation of offspring. Second, NEAT creates species of similar topologies and uses explicit fitness sharing to prevent any one species from dominating the population. Finally, because NEAT starts with a minimally complex topology with few or no hidden nodes, networks explore lower levels of behavior space before complexifying. As a result, more complex topologies are only maintained in a population if they improve fitness after several generations of optimization [8].

We use a Python implementation of NEAT developed by Stanley and Miikkulainen [1]. We use NEAT in this experiment for several reasons. First, we recognize NEAT as a tool to evolve complex behaviors, which will be critical to solving complex, deceptive mazes. Second, our experiment focuses on a task in an unconstrained behavior space. The range of behaviors available to any robot is limited by time, the physical and virtual environment, and the ability of the robot to effect change on the environment. While the runtime of each trial must be limited due to practical considerations, we hope to reduce limitations in the other two respects as much as possible. Therefore, in addition to creating an unconstrained virtual environment, we also use NEAT so that the robot is not restricted in its ability to effect change by a predefined and fixed topology. Third, as this research is largely an extension of previous work, we wish to remain consistent with other research that also used NEAT [3, 4, 7].

2.2 Novelty Search

Novelty search was developed by Stanley and Lehman as an alternative to searches that use an objective function [6]. Objective-based searches follow a gradient of increasing fitness towards a single solution. Yet in many deceptive environments, the optimal solution cannot be found by following a linear fitness gradient. Instead, globally optimal solutions may require behaviors that temporarily lower fitness. As a result, objective-based searches often become stuck in locally optimal solutions in environments that have some element of deceptiveness.

Novelty search abandons objectives, instead searching simply for behavioral novelty. Instead of a fitness function to measure individual success, novelty search uses a novelty metric. The novelty metric determines the relative sparseness of any given behavior in the defined behavior space. A novelty archive is maintained with a list of n behaviors with sparseness, ρ, over a certain threshold. The behavior x of a new individual is compared to the k-nearest neighbors in the archive, which is a list of notable past novel behaviors. If the new behavior has a sparseness score over the threshold it is added to the archive. Sparseness is calculated with the following function:

$$\rho(x) = \frac{1}{k} \sum_{i=0}^{k} dist(x, \mu_i)$$

where μ is i-th nearest neighbor of x. In the context of NEAT, the novelty metric is used as a replacement for the fitness function.

2.3 Augmenting Novelty Search

Recent research on novelty search has focused on implementations that retain the advantages of novelty search while also guiding the search towards a goal. The simplest solution, and the solution explored in this paper, is a linear blend of the novelty metric with an objective fitness function [7]. This multiobjective approach uses the objective metric to guide robots towards behaviors that are close to a solution and the novelty metric to avoid local optima.

Another approach is Minimal Criteria Novelty Search (MCNS), a variation on novelty search developed by Lehman and Stanley in which individuals that do not meet a minimum criteria are excluded altogether from reproduction. Tested on a deceptive maze task similar to the environment in this paper, MCNS performed better than both novelty search and an objective, fitness-based search [5].

Gomes, Urbano, and Christensen modify MCNS to create Progressive Minimal Criteria Novelty Search (PMCNS), which uses a ratcheting of the minimum criteria to force novelty search to fully explore smaller behavior spaces close to the solution [4]. In PMCNS,

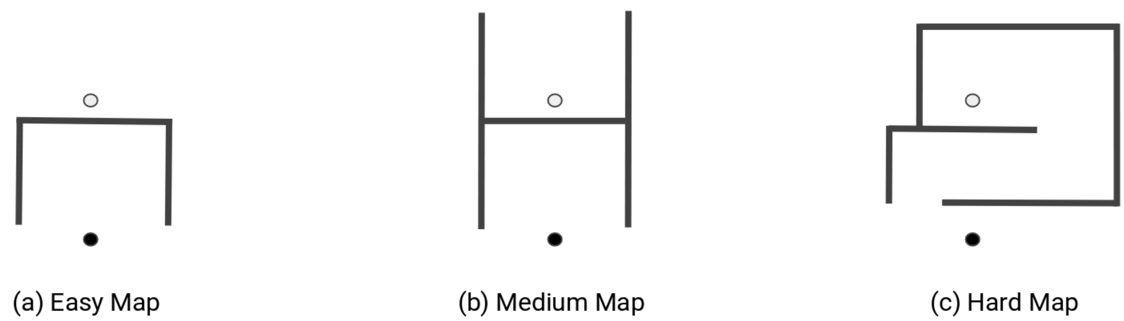

Figure 1: Maze navigation maps. In all maps, the dark circle represents the starting position of the robot and the light circle represents the goal.

the minimum criteria periodically increases over generations and is based on the fitness of the individual in the *nth* percentile. In an environment using swarm robots, PMCNS was found to substantially outperform both novelty and an objective, fitness-based search, and outperformed a linear blend in some respects.

Our experiment focuses on the linear blend augmentation of novelty search. While other research has done related work, our experiment is novel in that it is the first direct comparison between a linear blend search, novelty search, and an an objective, fitness-based search in a deceptive, unconstrained maze environment. Furthermore, our experiment explores how changes in deceptiveness impact the relative success of a linear blend, a topic that has to our best knowledge not been previously researched.

3 Experiments

3.1 Setup

The traditional task used to evaluate NS has been the deceptive maze environment. Following Stanley and Lehman, among others, our experiment follows this pattern. More broadly, we attempt to replicate to the best of our ability the parameters and experimental conditions of previous, related work.

Our domain consists of three separate mazes, with varying levels of complexity and deceptiveness (Figure 1). Each maze is situated in an unbounded environment; as the maze structure is the only part of the environment with obstacles, individual behaviors can consist entirely of unimpeded travel away from the maze for the full duration of a run. The mazes are constructed with deceptive cul de sacs that provide local optima in the fitness landscape. In the easy maze, successful robots have to navigate around a cup-shaped obstacle to avoid a local optima. In the medium maze, robots who navigate out of the first obstacle also have to navigate around a second obstacle and turn backwards to approach the goal. In the difficult maze, the cul de sacs are augmented by a maze with a small entrance; individual robots must first enter the maze before navigating the traditional deceptive environment. The difficult maze is a simpler but roughly comparable environment to the maze used by Stanley and Lehman.

For each maze, there is a fixed robot start position and a fixed goal within the maze (Figure 1). The start position of the robot is 6 units below the goal position in the easy and medium mazes and 8 units in the hard maze, due to the increased size of the hard maze. As a reference, at peak velocity the robot is able to travel 100 units over 1,000 time steps, the duration of each individual run. The environment and robots are created within the Jyro simulator.

Each robot is a modified version of the Jyro Pioneer robot. The robot (see Figure 2) has six sonar sensors that indicate the distance to the nearest obstacle, provided the obstacle is within 3 units. In addition, the robot has four pie-slice sensors that fire when the goal is in a particular quadrant relative to the heading and location of the robot, and a single bias node. The robot has two outputs, one for velocity and one for rotation.

For objective search, robots require a fitness function to guide exploration of a behavior space. In this case, fitness is defined as the normalized Euclidean distance between the position of the robot at the end of a run and the goal position. A robot finishes a run

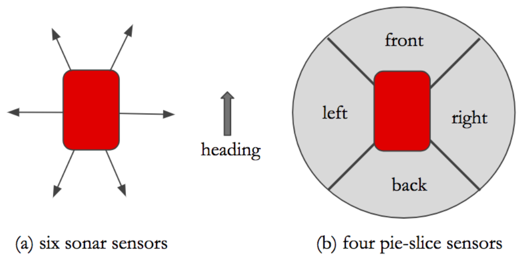

Figure 2: (a) The layout of the sonar sensors. Each arrow is a range finding sensor that indicates the distance to the closest obstacle in that direction, provided the robot is within 3 units of that obstacle. (b) The layout of the pie-slice sensors. Each robot has four such one-hot sensors that act as a compass, activating when the goal is within that pie slice. These sensors are dependent on both the location and the heading of the robot relative to the goal.

after the goal is reached, the robot stalls for more than 10 steps, or 1000 steps, whichever occurs first.

The novelty metric, which we used as the fitness function in NEAT, is based exclusively on the sparseness of the robot's behavior. As this was a maze task with a predetermined goal, the end position of the robot was chosen to represent behavior.

The novelty + objective search implementation uses a linear combination of the objective fitness function and the novelty metric:

$$score(i) = (1 - \rho) * fit(i) + \rho * nov(i)$$

where $\rho \in [0, 1]$. Due to computational limitations, we used $\rho = 0.5$ and did not test other values of ρ. In previous work, Cucco and Gomez found that $0.4 < \rho < 0.9$ produced the best results on a deceptive Tartarus task with a large behavior space, so we assume an even weighting will produce results that will be largely representative of other ρ-values [7].

For each of the three search implementations (novelty, objective, novelty + objective), we ran 40 trials for the easy and medium mazes using NEAT (see Table 1). For the hard maze, we were limited to 11 trials because of computational limitations. Each trial ended upon convergence to the solution or after 250 generations, whichever came first.

Parameter	**Setting**
Generations	max of 250
Population size	200
Input nodes	11
Hidden nodes	at gen 0: 0
Output nodes	2
Prob. to add link	0.2
Prob. to add node	0.1

Table 1: NEAT parameter settings used in the experiments.

3.2 Results

We evaluated our results in a variety of ways with two quantitative methods and one qualitative method. The first quantitative method was the measure of the average generation to convergence and second quantitative measure was maximum average best fitness. Our qualitative measure of evaluation was the end behavior of each individual after the conclusion of a run.

3.2.1 Generation of Convergence

Convergence indicates an individual robot came within 0.2 units of the goal position. A summary of the average generation at which an individual robot reached convergence can be found in Figure 3. In experiments with the easy map, all experiments for each implementation converged within 5 generations and at similar rates that did not have statistically significant differences. However, there were statistically significant differences in mean generations to convergence in the medium maze, and we achieved close to statistically significant results in the hard maze with only 11 trials. Specifically, novelty + objective search converged faster than novelty and objective searches alone. Statistical results have been computed using Student t-tests and can be found in Table 2.

Within the 250 generation limit, 64% of trials using novelty search and 36% of trials using objective search did not converge on the hard maze. Using novelty + objective search on the same maze, only 9% of trials did not converge.

3.2.2 Maximum fitness

Average maximum fitness is a measure of the fitness of the most successful individual in each generation, as measured by distance from the goal at the end

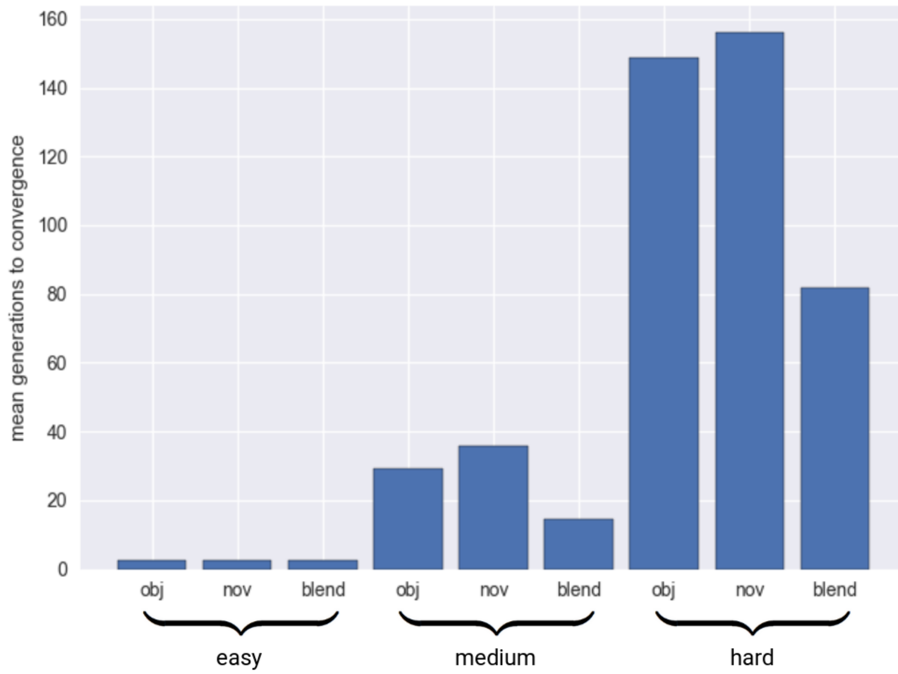

Figure 3: The average generation to convergence for each method and maze. Trials that did not converge within the maximum 250 generations are included in the averages and are assigned the value 250. The generation of convergence indicates an individual robot came within 0.2 units of the goal position within that generation.

	easy			**med**			**hard**		
	obj.	**nov**	**blend**	**obj**	**nov**	**blend**	**obj**	**nov**	**blend**
gen	2.38	2.45	2.30	35.75	29.25	14.6	148.82	165.27	81.91
std dev.	1.79	1.80	1.79	54.95	30.92	11.89	103.25	122.00	75.18
T-Test									
obj.	$p = 1.0$	$p = .85$	$p = .71$	$p = 1.0$	$p = .53$	$\mathbf{p < .01}$	$p = 1.0$	$p = .76$	$p = .11$
nov.	$p = .85$	$p = 1.0$	$p = .85$	$p = .53$	$p = 1.0$	$\mathbf{p < .05}$	$p = .76$	$p = 1.0$	$p = .12$
blend	$p = .71$	$p = .85$	$p = 1.0$	$\mathbf{p < .01}$	$\mathbf{p < .05}$	$p = 1.0$	$p = .11$	$p = .12$	$p = 1.0$

Table 2: Average generation to convergence and results of Students paired t-test. Each generation has a population of 200.

of a run. In contrast to convergence, average maximum fitness measures relative success over generations rather than absolute success. Results are displayed in Figure 4, with the exception that results for the easy maze are excluded because all runs regardless of search implementation method converged within 5 generations at similar rates. For the medium maze, all trials for novelty + objective search converged within 55 generations, while the worst case objective search took more than twice that long (143 generations). One trial with novelty search did not converge, and we see consistently lower maximum best fitness scores across generations. Experiments with the hard maze had similar results. 72% of novelty + objective trials converged in less than 70 generations while only 27% of trials in novelty and objective searches converged within 70 generations. In both the hard and medium mazes, novelty + objec-

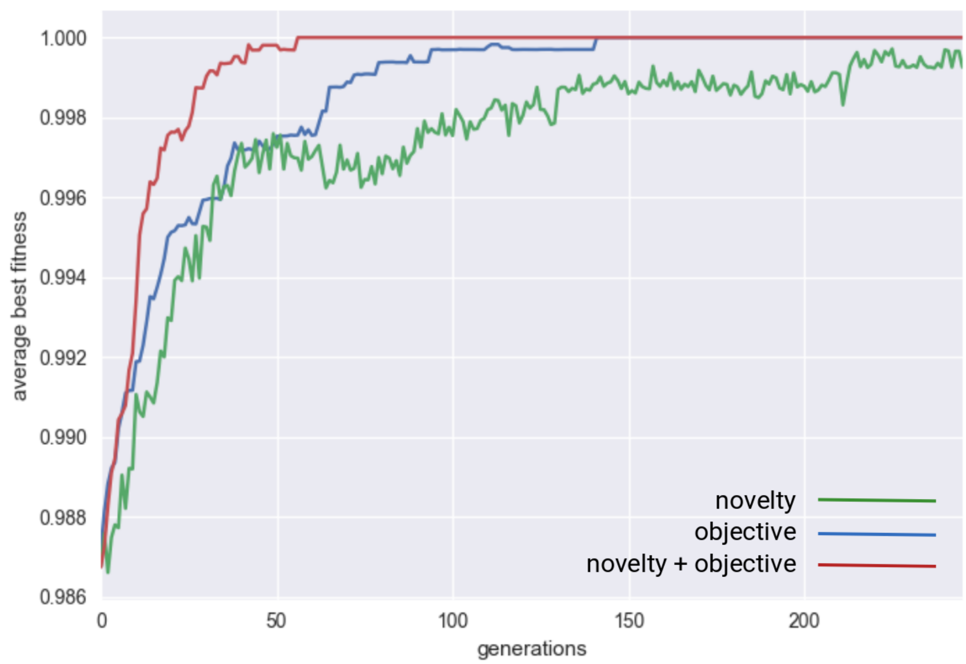

(a) medium maze

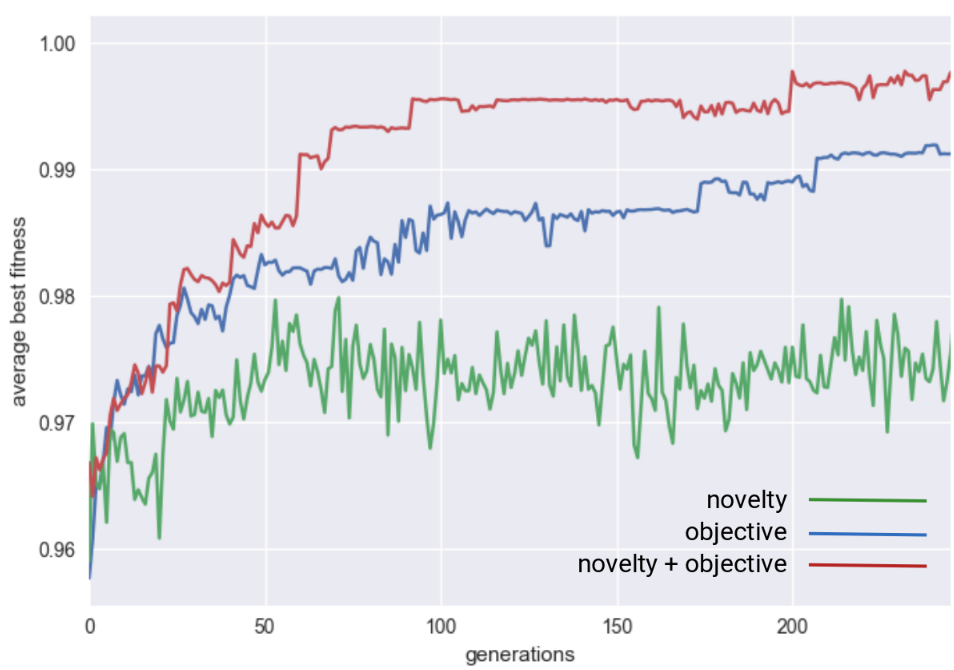

(b) hard maze

Figure 4: Average maximum fitness in the population for each generation (40 runs for the medium maze and 11 for the hard maze) for each search implementation method.

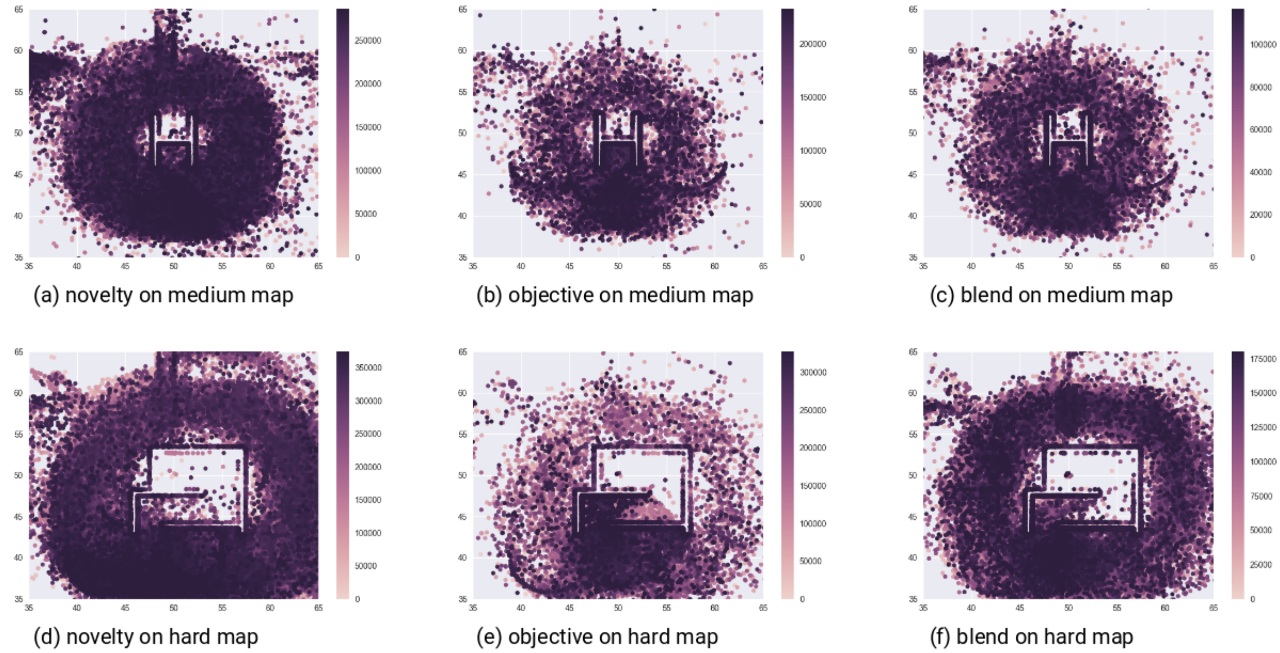

(a) novelty on medium map (b) objective on medium map (c) blend on medium map

(d) novelty on hard map (e) objective on hard map (f) blend on hard map

Figure 5: Each scatter plot displays the endpoints of most individuals over all generations of all trials of a given search implementation method. The few individuals with end behaviors far from the obstacle are not included in the plots. Lighter colored dots indicate individuals in earlier generations, while individuals in later generations are darker. Plots for the easy map are essentially equivalent and not shown.

tive search consistently outperformed novelty search and objective search across most generations.

3.2.3 Exploration Behavior

Figure 5 depicts the end position of every individual of all trials for each implementation method for the medium and hard mazes. The graphs reveal larger patterns of behavior that are difficult to describe quantitatively. It is important to note that the behavior space is unconstrained, and individual robots can and did travel up to 100 units away from the start position (50,44) within 100 steps. However, end positions of the vast majority of runs were within the boundaries of the axes of the scatter plot graphs (Figure 5). This crowding behavior around the start position may seem counterintuitive for novelty search, as individual behaviors with endpoints in the vast behavior space far from the start position are valued more than other behaviors. We hypothesize that the behavior shown for novelty search is the result either of a gradual exploration of the behavior space that will occur over many thousands of generations, sensors that only give meaningful data near the obstacles, or some combination of both.

As with mean generations to convergence, the scatter plot graphs for the easy map are essentially equivalent. In the medium map, there are clear differences between the different search implementations. Novelty search produces a generally even and dense scattering of points around the obstacle. Objective search produces many individual behaviors that get stuck in local optima (in the cul de sac below the goal), and other individuals collectively create a mysterious mustache formation. Novelty + objective search avoids local optima, as can be seen by the less dense area of endpoints in the cul de sac below the goal. At the same time, the distribution of points is less evenly scattered around the obstacle than novelty search. Thus, novelty + objective search creates robot behaviors that appear to be a combination of goal searching behavior and novelty.

4 Discussion

The experimental results show that combining novelty search and objective search in an unconstrained, deceptive maze environment is more effective than novelty or objective alone at reaching a desired goal. This is unsurprising and confirms previous research that shows the effectiveness of augmenting novelty search to incorporate some goal-finding behavior. It is clear from our results that novelty search struggles in unconstrained environments and objective search struggles in deceptive environments. In contrast, novelty + objective search evolves individuals with a multiobjective task: reach the goal and reach end positions not previously explored. Our results show that this implementation method overcomes both the vast behavior space problem associated with novelty search and the local optima problem associated with objective search.

The results also show that the deceptiveness of the task is directly related to the relative performance of novelty + objective search. The easy map was deceptive in a sense that it had a local optimum, but practically it was not deceptive enough to create clear differences between novelty + objective search and the other two methods. The only significance differences between novelty + objective and the other methods appeared in the more deceptive medium and hard mazes. Therefore, we conclude that there is a certain threshold of deceptiveness that must be reached before novelty + objective can outperform the other two methods in an unconstrained environment. We hypothesize that this threshold varies based on the hyper-parameters of the algorithm (in our case, NEAT) and the type of deceptive task.

This paper adds to the body of existing research about novelty search and its various augmentations. Previous research explored the ability of a linear blend on large but constrained behavior spaces and found that it performed better than objective and novelty alone. Our research shows that a linear blend continues to outperform in an entirely unconstrained physical maze environment. Along with other research, we now have a clear understanding of which search methods perform best in various environments:

- deceptive environments:
  - small behavior spaces: novelty or augmented novelty search
  - large behavior spaces: augmented novelty search (e.g. linear blend, MCNS, PMCNS)
- non-deceptive environments:
  - small behavior spaces: objective search
  - large behavior spaces: objective search)

Outside of simulations, much of the world is unconstrained and deceptive. Therefore, our results can meaningfully contribute to how robots navigate through the physical world. For example, many small search and rescue robots tasked with moving to a specific location in unknown and deceptive terrain might perform more effectively with a linear combination of novelty and objective search than with a pure objective search method.

An obvious next step is to run more tests on our hard maze in order to get more significant results. Furthermore, the construction and testing of more mazes would be interesting, especially those mazes that are impossibly deceptive or complicated for objective and novelty search but may be solvable by novelty + objective. Finally, we used a very simplified method of combining novelty and objective searches with a static linear ratio between the novelty and objective metrics. The use of a more complicated function to model the novelty/objective ratio, including some form of annealing, might produce even better results.

References

[1] The NEAT users page. www.cs.ucf.edu/~kstanley/neat.html. Accessed: 2014-05-01.

[2] Giuseppe Cuccu and Faustino Gomez. When novelty is not enough. *EvoApplications.*

[3] John Dinh and David Wilikofsky. Noveleat: Combining novelty search and objective-based search. *Swarthmore Computer Science Department, CS81*, 2012.

[4] Paulo Urbano Jorge Gomes and Anders Lyhne Christensen. Progressive minimal criteria novelty search. *Advances in Artificial Intelligence IBERAMIA*, 13:281–290, 2012.

[5] Joel Lehman and Kenneth Stanley. Revising the evolutionary computation abstraction: Minimal criteria novelty search. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO).*

[6] Joel Lehman and Kenneth Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2), 2011.

[7] Jean-Baptiste Mouret. *Novelty-Based Multiobjectivization*, pages 139–154. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

[8] Kenneth Stanley. Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research*, 21, 2004.

[9] Kenneth O. Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2), 2002.