

Emergence of Novelty in Evolutionary Algorithms

David Herel¹, Dominika Zogotova¹, Matej Kripner² and Tomas Mikolov¹

¹ CIIRC, CTU in Prague, 160 00 Dejvice

² Charles University, 121 16 Nové Město
david.herel@seznam.cz

Abstract

One of the main problems of evolutionary algorithms is the convergence of the population to local minima. In this paper, we explore techniques that can avoid this problem by encouraging a diverse behavior of the agents through a shared reward system. The rewards are randomly distributed in the environment, and the agents are only rewarded for collecting them first. This leads to an emergence of a novel behavior of the agents. We introduce our approach to the maze problem and compare it to the previously proposed solution, denoted as Novelty Search (Lehman and Stanley, 2011a). We find that our solution leads to an improved performance while being significantly simpler. Building on that, we generalize the problem and apply our approach to a more advanced set of tasks, Atari Games, where we observe a similar performance quality with much less computational power needed.

Introduction

In nature, individuals of one population compete with each other for survival, and only the fittest ones can pass their genes on to the next generation. Similarly, in evolutionary algorithms, it is the agents who are rewarded for being the fittest. In order to realize the rewards, an objective function is needed to evaluate the individual solutions based on their proximity to the optimum (Goldberg and Holland, 1988) and to guide us through the search space to a valid solution. Even though evolutionary algorithms have been proven successful in numerous optimization tasks (Goldberg and Holland, 1988), the deceptiveness of the environment is rather problematic. In reality, increasing fitness does not always reveal the best path and can mislead us from finding the global optimum (Goldberg, 1987).

Aiming to solve this problem, many algorithms were developed (Hutter and Legg, 2006)(Basu and Bhatia, 2006) that significantly reduce the deceptiveness. But the underlying problem remains – the objective function may still misdirect the search toward a dead end. Actually, finding a solution to this problem requires thinking counter-intuitively and abandoning the objective completely, as Novelty Search does (Lehman and Stanley, 2011a). We can think of the Novelty Search as a divergent search technique applied

to an evolutionary computation, which also introduces a new reward system. Instead of rewarding agents based on their proximity to the goal, agents are rewarded for being different from others in the population. In comparison with the traditional objective-based evolutionary processes, Novelty Search algorithms have demonstrated their efficiency in many highly deceptive problems (Lehman and Stanley, 2011a).

In this paper, we propose our own search method called Sugar Search, which aims to reproduce or surpass Novelty Search results without explicitly defining novelty. Thus, the behavioral novelty will emerge as a by-product of the environment and objective function definition. Building on the qualities of the Novelty Search (Lehman and Stanley, 2011a), we reproduce this method on the maze problem, which is identical to the mazes presented in (Lehman and Stanley, 2011a), and use it for comparison with our Sugar Search technique. Following the proof of concept and given the competitive results, we also conduct experiments with different reward densities and a combined fitness-based and Sugar Search approach in the maze environments.

In the second part of this paper, the main objective is to present a more complex problem that our search technique can be applied to. For this purpose, we present Sugar Search in the Atari games environment. We compare the results of our method to the fitness-based approaches, which have shown good results on some Atari games (Bellemare et al., 2012), as well as some reinforcement algorithms like DQN (Mnih et al., 2013) and AC3 (Mnih et al., 2016). Because of the challenges that arise with agent allocation in such a complex environment, we also generalize our approach to this problem.

Related work

There have been several approaches to solve the problem of local minima by encouraging novelty. One of the most well-known ones is the Novelty Search algorithm (Lehman and Stanley, 2011a). It ignores the objective and instead optimizes how unique the individual's behavior is. It

greatly differs from the random search, because it provides higher rewards to a diverse behavior. There is also an archive, which stores the previously explored areas of the search space. The novelty score is then counted as the average distance from the k -nearest neighbors. However, the distance metric is problem-dependent, which could be problematic (Lehman and Stanley, 2011a). Another interesting idea, which promotes the novel behavior, is the Behavior characterization algorithm (Meyerson et al., 2016). This algorithm maps each individual behavior to a vector space. The search is then driven towards diversity in a metric space of these behaviors. Another approach to this problem is Curiosity learning. It works with intrinsic rewards that promote exploration (Pathak et al., 2017). Building on Novelty Search, Surprise Search was later introduced as another divergent search technique working with the notion of surprise. It achieves an efficiency comparable to the Novelty Search and is also able to find solutions more frequently (Gravina et al., 2016).

The ability to generate a diverse set of high-performing solutions in evolutionary algorithms is crucial and has been addressed in many so-called Quality-Diversity algorithms, such as the MAP-Elites algorithm (Mouret and Clune, 2015) or NS with Local Competition (Lehman and Stanley, 2011b). In MAP-Elites (Mouret and Clune, 2015), the search space is discretized into unique regions, for each of which the best-performing solution is identified. The selection of the fittest individuals within a population is restricted to a specific feature, the fittest individuals are kept in a multidimensional archive and only replaced if they are outperformed. This way a diverse group of well-performing individuals with different features can develop in the domain space in a single run. However, the problem of local optima within the elite may still occur over time and the scalability to high-dimensional spaces is also problematic due to the exponential growth of the regions with increasing dimensionality. This problem is overcome by CVT-MAP-Elites algorithm (Vassiliades et al., 2016), which extends the MAP-Elites algorithm and proposes a solution to reduce the number of regions in high-dimensional spaces.

The problem of the deceptive local optima arises not just in evolutionary algorithms, but also in reinforcement learning itself. The reinforcement approaches optimize only for the reward (e.g. DQN (Mnih et al., 2013), A3C (Mnih et al., 2016)) often fail to learn behaviors and strategies to overcome the local optima and solve the task at hand (Jackson and Daley, 2019). To solve this issue, Novelty Search and other approaches such as episodic curiosity (Savinov et al., 2019) were explored. In (Jackson and Daley, 2019) instead of measuring the final position of the agent and comparing it to other individuals of the current generation and the ones in the archive, the Levenshtein distance between the sets of actions performed by the agents

was used to measure the extent of behavioral novelty. The approach of comparing the sequences of actions instead of the final positions removes the need to detect the exact position of an agent for each frame, and therefore makes it easier to generalize Novelty Search to any Atari game.

Our approach: Sugar Search

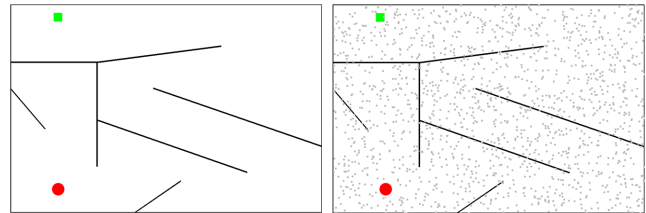


Figure 1: Visualization of the hard map with (on the right) or without the sugars (on the left). The red dot is the starting point, the green dot is the goal, and the black lines are the obstacles. The sugars are represented as gray rectangles, they are placed randomly, and their density is the only hyperparameter. Only the agent who collects the sugar first is rewarded for it.

Even though Novelty Search is conceptually interesting, it seems to employ a certain degree of supervision, which is needed to measure and decide how unique each agent's behavior is. But this is inconsistent with what we observe in nature, where it is rather the environment that formulates, motivates, and rewards individuals for being different.

This difference can be demonstrated on a maze problem specified as follows. A maze contains multiple obstacles and a target. In the environment, there are agents whose goal is to reach the target while avoiding obstacles. If an agent reaches an obstacle, its movement stops.

In the Novelty Search algorithm, the agents that stop at a point far away from the others will receive a higher fitness, which motivates them to explore the maze. On the downside, this requires us to compute the k -nearest neighbors for each agent.

Our approach is fairly different. We propose a method called Sugar Search, which introduces sugar objects as a form of a reward. Sugars are distributed randomly (and uniformly) over the map, and an agent is rewarded only if he is the first one to collect a sugar object. This way the only form of supervision is the sugar placement and, compared to Novelty Search, we can avoid the costly computation of the k -nearest neighbors and eliminate the archive completely. The concept of the sugar-based search leads to a different pattern in the exploration of the search space than Novelty Search does.

the nearest obstacle on each side, and 4 radar sensors, which indicate the direction to the goal. The value of these sensors is normalized and used as the input to the neural network, which will evaluate whether the agent should move up, down, left or right. The specifics of the NN that was used throughout all the experiments are further described in Neural Network configuration.

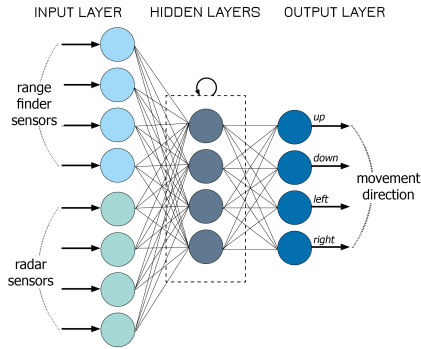


Figure 3: Agent’s neural network in the maze navigation task.

Atari games task

The maze navigation task is a well-suited environment to test our approach and compare it to other works such as Novelty Search. More advanced problems such as walking robots could show the same results.

Further, we decided to implement this idea and generalize it to the environment of Atari games. Atari games have been widely used to compare many machine learning algorithms, especially deep reinforcement learning algorithms (Mnih et al., 2013). But due to the many implementations of Atari games, it is difficult to compare the results of different algorithms. That is why the OpenAI Gym (Brockman et al., 2016) games were chosen at the beginning to avoid any differences in the underlying environment and to see if our ideas could compare with some of the famous algorithms (Mnih et al., 2013)(Young et al., 2019).

Later on, due to the difficulties with the agent localization and tracking in the OpenAI Gym (Brockman et al., 2016), we have decided to test our Sugar Search approach on the MinAtar (Young and Tian, 2019). MinAtar is inspired by the Arcade Learning Environment (Bellemare et al., 2012) but simplifies the games to make experiments with the environments more accessible and efficient. Currently, MinAtar provides analogues to 5 Atari games, which we have extended by Ms. Pacman and Montezuma’s revenge implementation. The environments provide a state representation, where each of the n channels corresponds to a game-specific object, such as the ball, the paddle, and the brick in the game Breakout. Thanks to this environment, we

were able to track the agent and perform Sugar Search.

The goal of the agent is similar to the maze navigation task: to develop such NN that it can achieve the highest score possible. In the MinAtar environment, the agent’s input consists of all the state representations of the objects. In the OpenAI gym, the input is the screen pixels. However, for computational reasons, the screen resolution was down-sampled 8 times to $(210/8) \times (160/8)$ pixels, and the RGB colors were converted into gray-scale. This alteration resulted in 520 input neurons used in the experiments, from the original 100,800.

All the necessary configuration parameters used in the Atari games task are included in Table 1. In the fitness-based search, a score was used to determine the fitness value of the agent. In the Novelty Search technique, 5-nearest neighbors were used. However, to reach even better performance, we had to modify the idea of our approach slightly by re-spawning a new sugar during the game. Because in some Atari games, e.g. Ms. Pacman, the agent has no finish point and its goal is to collect as many points as possible before getting killed. This requires traveling to places that have been explored in the past but may be beneficial to revisit in the present moment. To motivate the agent to behave in such a way, we decided to re-spawn a sugar every l frame.

Neural Network configuration

Table 1: The experiment configuration for both the Maze navigation task and the Atari games of our Sugar Search with RNN.

Parameters	<i>Maze env.</i>	<i>Atari env.</i>
Population size	250	75
Max. generations	1200	10000
Time frame (max. frames)	600	250
Sugars density per pixel	0.3	0.3
l re-spawn time frame	-	5
Hidden layer	1	1
Hidden neurons	32	32
Activation function	ReLU	ReLU
Gaussian noise	$\mathcal{N}(0, 0.1^2)$	$\mathcal{N}(0, 0.1^2)$

Each agent has to develop a neural network, which will allow him to achieve the goal. The challenge is how to develop this neural network. In the original paper (Lehman and Stanley, 2011a) NEAT, an evolutionary algorithm was used (Stanley and Miikkulainen, 2002). However, it has numerous hyperparameters which are difficult to optimize. Instead, we used Elman recurrent neural network (Elman, 1990) (RNN), which is a simple NN with a recurrent hidden layer. We used random Gaussian noise to mutate its weights.

In our algorithm, we randomly initialize weights to the RNN for each agent. When the time limit is reached, we evaluate 10% of the best agents and keep them for the next generation. The rest of the population is modified - random Gaussian noise is added to its weights. After this mutation, the agents are added to the next generation. This approach is much simpler and results in fewer hyperparameters.

The benefits of our concept of the NN are obvious from the experiments on the medium map, where NEAT needed 19.8 generations on average to finish the maze, whereas the simple Elman neural network (Elman, 1990) with our evolutionary algorithm only needed 10.4 generations (these results were averaged from 50 runs). This proves that our approach is much better, and for these reasons we used it in all of our experiments. All the other parameters necessary for the reproduction of the experiments are listed in Table 1.

Results

Maze navigation task

We have evaluated the results from the least to the most complex map: the medium map, then the hard map, and finally the super-hard map in Figure 2. All the results are shown in detail in Table 2. The level of success of the search techniques is measured by the average number of generations needed to reach the goal and the number of successfully finished runs out of 50 (meaning runs where we were able to reach the goal). Note that we use the same way to reference the maps as in the Novelty Search (Lehman and Stanley, 2011a) paper to avoid confusion.

On the medium map, Novelty Search and our Sugar Search performed similarly, while the fitness-based search was approximately two times slower. All of these techniques were able to successfully finish all 50 out of 50 runs. The difference between divergent searches and fitness-based search is statistically significant ($p < 0.01$; Student's t-test). On the hard map, where the deceptiveness of the environment is more compelling, a remarkable difference between all three examined algorithms can be distinguished. Sugar Search performed significantly better than Novelty Search in terms of the success rate and the number of generations needed, as shown in Table 2. The difference between Sugar Search and Novelty Search is statistically significant ($p < 0.05$; Student's t-test). Both divergent approaches largely outperformed the fitness-based search, which was unable to reach the goal in any run.

Similar results were observed on the super-hard map. The fitness-based approach failed to reach the goal in any run again. The divergent approaches were able to solve this maze with a reasonable success rate: 14 for Novelty Search versus 24 for Sugar Search, and with an average number of generations needed: 345.3 for Novelty Search versus 280.4 for Sugar Search. This difference is not statistically significant ($p < 0.05$; Student's t-test). A larger sample of runs

on this map would be needed to confirm the significant difference of our Sugar Search outperforming Novelty Search, which can be observed in the results.

Table 2: Maze navigation task performance in the average number of generations needed to reach the goal and the number of successfully finished runs out of 50 (meaning runs where we were able to reach the goal). The standard deviation is in parentheses. Bold font indicates the best performance for each metric. All numbers are averaged on the successful runs only. Symbols \uparrow (\downarrow) represent that the higher (lower) the better.

Maze Results	Generations \downarrow			Finished runs (out of 50) \uparrow		
	<i>med. map</i>	<i>hard map</i>	<i>super-hard map</i>	<i>med. map</i>	<i>hard map</i>	<i>super-hard map</i>
Fitness-based approach	17.6 (9.18)	-	-	50	0	0
Novelty approach	10.4 (5.08)	237.7 (172.12)	345.3 (183.42)	50	23	14
Sugar Search (ours)	11.7 (5.96)	153.3 (97.53)	280.4 (101.78)	50	35	24

Reward density

Naturally, while working on the maze problem, an interesting question arose. How does the density of the reward distribution on the map affect the results? Our hypothesis was that having as many sugars as there are pixels on the map would produce the best results since theoretically a higher density of sugars should lead agents to the goal more efficiently.

From a series of experiments on the medium map with a varying sugar density parameter as shown in Figure 4, we have proven this hypothesis valid. Placing sugar on every pixel produces the best results, and reducing the sugar density generally slows down the process of reaching the goal. Interestingly, performing a small-scale reduction of the reward density still produces comparable results.

These findings can be beneficial when applied to a similar problem defined in a continuous space, where the sugar placement cannot be realized per pixel for computational reasons.

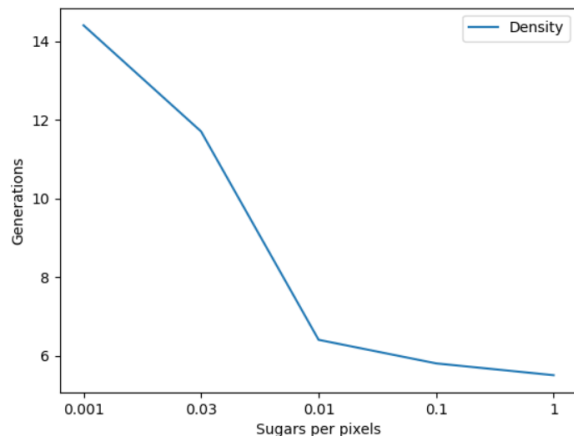


Figure 4: Average number of generations needed to reach the goal (*the lower the better*) and the sugars per pixel density trade-off curve.

Combining Sugar Search with the fitness-based approach

We could observe that in a deceptive maze problem, the divergent search, especially our Sugar Search, completely outperforms the objective-driven search. But what if we combined these two approaches?

We wanted to explore this idea further and decided to design a combined approach, where the agent is rewarded not only for being novel (divergent search) but also for moving towards the goal (fitness-based search). This combined approach requires us to define the weights of each component that would produce the best results, meaning we need to find an optimal alpha.

$$p(x) = \alpha * distance + (1 - \alpha) * sugars$$

In order to find an optimal alpha and observe how this weighted approach stands in comparison with the original Sugar Search and other techniques, we have executed a series of experiments on the medium and the hard maps. To reduce the computation time, as a proof of concept, we have excluded the super-hard map from the following experiments. Based on the results of the basic Sugar Search algorithm, we can expect to observe similar trends on the super-hard map as the hard map.

From the results shown in Table 3, we can conclude that the combined approach leads to improved performance, specifically in the average number of generations needed to reach the goal. On the other hand, an improvement is not observable in the agent’s success rate. It was 29.3 on average on the hard map, which is slightly worse than Sugar Search alone.

Table 3: Maze navigation task with the weighted approach: performance in the average number of generations needed to reach the goal and the number of successfully finished runs out of 50 (meaning runs where we were able to reach the goal). The standard deviation is in parentheses. Bold font indicates the best performance for each metric. All numbers are averaged on the successful runs only. Symbols \uparrow (\downarrow) represent that the higher (lower) the better.

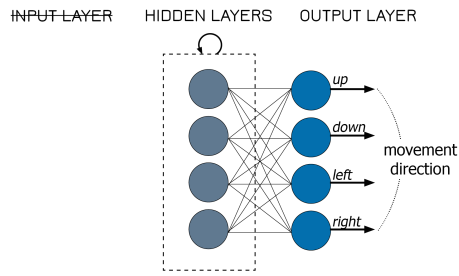
Maze Results	Generations \downarrow		Finished runs (out of 50) \uparrow	
	medium map	hard map	medium map	hard map
Fitness-based approach	17.6 (9.18)	-	50	0
Novelty approach	10.4 (5.08)	237.7 (172.12)	50	23
Sugars approach	11.7 (5.96)	153.3 (97.53)	50	35
Weighted with alpha 1/4	8.0 (7.97)	148.7 (104.25)	50	32
Weighted with alpha 2/4	5.4 (4.09)	122.4 (86.55)	50	30
Weighted with alpha 3/4	6.7 (5.59)	139.1 (90.87)	50	26

Agents without sensors

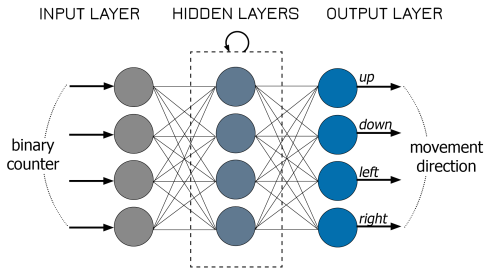
The idea behind this section is to prove that the agents are really processing the information from their sensors to avoid hitting an obstacle and that they are not simply memorizing a sequence of optimal actions instead. This problem could arise when we take into consideration the amount of memorization and the overall generalization in terms of the agents’ performance in the mazes.

To prove this, we have decided to experiment with a blind agent. The original agents were blinded - meaning they had all their inputs removed and had to rely solely on their memory in the hidden layer, as shown in Figure 5a. Later, we have modified this concept further in order to elevate the memory and created a modified blind agent that had a simple binary counter as its input as shown in Figure 5b. This binary counter is a timer, which helps the agent memorize sequences of actions.

From Table 4 we can see that the blind agents can navigate throughout the map successfully without having any information about the environment, simply by memorizing a sequence of actions. However, having information from the sensors about the nearest obstacle and the location of the goal significantly accelerates the process of reaching the



(a) A blind agent relies solely on the memory in the hidden layer.



(b) A modified blind agent with a simple binary counter on the input.

Figure 5: Modified agents without the sensors and their NN.

goal. It is also important to note that the blind agents are not adaptable and will fail completely when placed in a slightly different starting point or if the task is modified.

Table 4: Maze navigation task for agents without sensors: performance in the average number of generations needed to reach the goal. The standard deviation is in parentheses. Bold font indicates the best performance for each metric. Symbols \uparrow (\downarrow) represent that the higher (lower) the better.

Maze Results	<i>Sensors</i> \downarrow	<i>Binary counter</i> \downarrow	<i>No input</i> \downarrow
medium map	11.7 (5.96)	37.8 (50.88)	96.1 (36.91)
hard map	153.3 (97.53)	597.4 (354.65)	1154.7 (592.78)

MinAtar results

In this environment, we ran experiments on 7 different Atari games to test and compare the fitness-based approach, Sugar Search, and its generalization - Pixel Novelty to other techniques like Novelty Search (Lehman and Stanley, 2011a) or DQN (Mnih et al., 2013). The results in Table 5 show that Sugar Search does not always lead to a higher score compared to the fitness-based approach, but it performs comparably well with Novelty Search consistently across different games, as shown in Table 5. The scores achieved

with Pixel Novelty are also competitive or comparable to other approaches, meaning that Pixel Novelty is a valid generalization of Sugar Search.

From the results in Table 5, we can conclude that the achieved scores differ based on the characteristics of each game, where some approaches seem to be more fitted than others. Generally speaking, we can observe a pattern in Atari games with no distinctive local optima, such as Breakout or Space Invaders, where a simple fitness-based search outperforms both our Sugar Search and Novelty Search. Interestingly, Pixel Novelty also performs rather well in these games (it even achieved the highest score in Breakout).

Regarding Sugar Search, it outperformed both the fitness-based search and Novelty Search by far in games like Freeway or Ms. Pacman and was even able to compete with algorithms like DQN (Mnih et al., 2013) despite its simplicity. This demonstrates that our search technique is very powerful and able to produce competitive results.

Table 5: MinAtar games performance in the average number of the achieved game score. Bold font indicates the best performance for each metric. All numbers are averaged from 100 generations/episodes. Symbols \uparrow (\downarrow) represent that the higher (lower) the better.

Game	<i>Fitness-based Search</i> \uparrow	<i>Sugar Search (ours)</i> \uparrow	<i>Novelty Search</i> \uparrow	<i>Pixel Novelty (ours)</i> \uparrow	<i>DQN</i> \uparrow
Seaquest	8	7	6	5	20
Breakout	13	5	5	21	9
Asterix	12	10	14	10	17
Freeway	16	53	49	28	50
S. Invaders	43	10	14	21	45
MsPacman	1897	3580	1962	2480	1099
Montezuma's revenge	0	41	22	22	46

OpenAI Gym results

In the OpenAI Gym environment, experiments on 12 Atari games in total were performed with both Pixel Novelty and the fitness-based approach. The specific games to be included in the experiments were selected to represent equal subsets of games, half of which contains multiple obvious local minima (e.g. Montezuma's revenge) and is expected to show the advantages of our approach. The results of the experiments are presented in Table 6 in comparison with the fitness-based approach and reinforcement learning

algorithms like DQN (Mnih et al., 2013) and A3C algorithm (Mnih et al., 2016).

Table 6: OpenAI Gym games performance in the average number of the achieved game score. Bold font indicates the best performance for each metric. All numbers are averaged from 100 generations/episodes. Symbols \uparrow (\downarrow) represent that the higher (lower) the better.

Game	<i>Fitness-based</i> \uparrow	<i>Pixel Novelty (ours)</i> \uparrow	<i>DQN</i> \uparrow	<i>A3C FF</i> \uparrow
BeamRider	1188.6	900.4	8672.4	13235.9
Breakout	15.4	24.1	303.9	551.6
Enduro	86.9	52.2	475.7	82.2
Pong	21.0	24.3	16.2	11.4
Seaquest	780.4	480.3	2793.9	2300.2
S. Invaders	1075.7	1205.3	1449.7	2214.7
Riverraid	2390.5	1810.7	4065.3	10001.2
Freeway	30.0	0.4	25.8	0.1
Gravitar	2350.4	700.6	216.5	269.5
Zaxxon	6240.1	5400.4	831.0	2659.0
Venture	440.5	200.3	54.0	19.0
Montezuma's revenge	0.0	10.0	50.0	53.0

In the light of the results, rewarding the agents solely for generating new pixels in Pixel Novelty leads to the same or better results as the score-driven reward system in the fitness-based approach. These overall positive results are the outcome of the agent's behavior in Pixel Novelty: as the agent tries to generate a new unseen screen content, he conquers obstacles on the way and by doing so also achieves a reasonable score.

Additionally, when compared to DQN and A3C algorithms, we are able to achieve a higher score in 5 out of 12 games. Furthermore, it is important to highlight that our results were achieved with a significantly greater efficiency, since we only needed 20 million frames and 1 CPU for 24 hours, whereas the results from DQN and A3C were trained on 320 million frames, which is 16 times more.

We have proven that our simple evolutionary technique is a scalable competitive alternative to the reinforcement-learning algorithms (Salimans et al., 2017) and that it can achieve similar results even with a divergent-driven approach and with a largely improved efficiency.

Conclusion

The idea of placing rewards within the environment to precede the problem of local minima performs surprisingly well in the maze experiments, especially taking into consideration that this concept completely disregards the objective (score). As the experiments have shown, our technique outperformed the popular Novelty Search (Lehman and Stanley, 2011a) in both the success rate and the solution speed, furthermore, it is also conceptually simpler and easier to implement. Following this proof of concept of our technique, we have also explored the effect of multiple local minima and different reward densities on the search efficiency. The results indicate that the sparse rewards model could be applied to problems in a continuous or a higher-dimensional space.

We have subsequently generalized our search technique and studied how it performs in a more general and popular set of tasks - Atari games. Our technique performed comparably well with a basic fitness-based search and some of the most common reinforcement learning algorithms (DQN (Mnih et al., 2013), A3C (Mnih et al., 2016)). We consider it to be a great achievement considering that our results were computed on much fewer frames than DQN or A3C.

In summary, we believe that it is rather the environment that formulates, motivates, and rewards individuals for being different and novel as we can observe in nature. Also supported by the fact that the concept has been proven valid on multiple experiments in this paper. This suggests that new ways of defining novelty, which avoid any form of supervision, can generally be successful and could be worth exploring in the future.

Acknowledgements

We would like to thank GoodAI for the research funding and valuable advice. We would also like to express appreciation to Daniela Hradilova for the useful discussions and suggestions. Our work is part of the RICAIP project that has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 857306.

References

- Basu, S. K. and Bhatia, A. K. (2006). Mitigating deception in genetic search through suitable coding. In *Proceedings of the 13th International Conference on Neural Information Processing - Volume Part III, ICONIP'06*, page 946–953, Berlin, Heidelberg. Springer-Verlag.
- Bellemare, M., Naddaf, Y., Veness, J., and Bowling, M. (2012). The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. (2016). Openai gym. *arXiv preprint arXiv:1606.01540*.

- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14(2):179–211.
- Goldberg, D. E. (1987). Simple genetic algorithms and the minimal, deceptive problem.
- Goldberg, D. E. and Holland, J. H. (1988). Genetic algorithms and machine learning. *Machine Learning*, 3(2):95–99.
- Gravina, D., Liapis, A., and Yannakakis, G. (2016). Surprise search: Beyond objectives and novelty. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016, GECCO '16*, page 677–684, New York, NY, USA. Association for Computing Machinery.
- Hutter, M. and Legg, S. (2006). Fitness uniform optimization. *IEEE Transactions on Evolutionary Computation*, 10(5):568–589.
- Jackson, E. C. and Daley, M. (2019). Novelty search for deep reinforcement learning policy network weights by action sequence edit metric distance.
- Lehman, J. and Stanley, K. (2011a). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19:189–223.
- Lehman, J. and Stanley, K. O. (2011b). Evolving a diversity of creatures through novelty search and local competition.
- Meyerson, E., Lehman, J., and Miikkulainen, R. (2016). Learning behavior characterizations for novelty search. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016, GECCO '16*, page 149–156, New York, NY, USA. Association for Computing Machinery.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. In Balcan, M. F. and Weinberger, K. Q., editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1928–1937, New York, New York, USA. PMLR.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. (2013). Playing atari with deep reinforcement learning.
- Mouret, J. and Clune, J. (2015). Illuminating search spaces by mapping elites. *CoRR*, abs/1504.04909.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. (2017). Curiosity-driven exploration by self-supervised prediction. *ICML'17*, page 2778–2787. JMLR.org.
- Salimans, T., Ho, J., Chen, X., Sidor, S., and Sutskever, I. (2017). Evolution strategies as a scalable alternative to reinforcement learning.
- Savinov, N., Raichuk, A., Marinier, R., Vincent, D., Pollefeys, M., Lillicrap, T., and Gelly, S. (2019). Episodic curiosity through reachability.
- Stanley, K. O. and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127.
- Vassiliades, V., Chatzilygeroudis, K. I., and Mouret, J. (2016). Scaling up map-elites using centroidal voronoi tessellations. *CoRR*, abs/1610.05729.
- Young, K. and Tian, T. (2019). Minatar: An atari-inspired testbed for thorough and reproducible reinforcement learning experiments. *arXiv preprint arXiv:1903.03176*.
- Young, K., Wang, B., and Taylor, M. E. (2019). Metatrace actor-critic: Online step-size tuning by meta-gradient descent for reinforcement learning control.