Towards more robust and autonomous AI via introducing creativity and curiosity

Taavi Luik* and Jaan Aru*

Abstract— **In order to create more robust and capable AI systems, we need to narrow the gap between the real world and AI datasets. More diverse and complex datasets would increase the need for more autonomy in learning, leading to more robust systems. We have developed a dataset of ambiguous visual stimuli which existing pre-trained vision models fail to classify correctly. We take inspiration from evolutionary optimization and human psychology, which highlight the benefits of divergent thinking, and are working towards more creative and curious AI systems that can learn more complex tasks.**

I. INTRODUCTION

One of the fundamental issues of optimization is that the search landscapes of many tasks are rugged, containing many local optima. Search cannot always be guided well with an objective function, because objective functions can be 'deceptive', i.e. get stuck in local optima. This happens because of failure to value 'stepping stones', which are intermediate solutions that lead to better solutions in the future [12]. In complex environments and ambitious goals, good solutions might be many stepping stones away. For this reason, signals other than the objective function have been explored, e.g. guiding search via novelty alone [12], surprise (deviation from the expectation [8]), curiosity [4, 14], and more.

Lehman's and Stanley's Novelty Search algorithm has spun off Quality Diversity (QD) algorithms, which, in contrast to traditional optimization algorithms, search for a diverse set of solutions [3]. QD algorithms require manually defined (and/or learned) descriptors for the problem at hand. Importantly, these descriptors are about the phenotype, not genotype - they describe high level properties of solutions, not low level details. These dimensions define a container (an 'archive') where individuals created during search will be stored. Individuals will be assigned descriptors and are entered into the archive if they outperform individual(s) in the same niche or if the niche they fall into is empty. As a result, a second type of selection pressure is created, one that does not increase fitness, but diversity.

The QD framework has been used to teach robots to learn how to walk after incurring damage [1], to produce procedurally generated content for video games [15], as an alternative to train neural networks via gradient descent (neuroevolution, [18]), and also in difficult reinforcement learning settings [6].

Another issue to highlight is that progress in benchmarks does not always translate to progress in real-world environments [17]. Existing benchmarks fail to fill the gap between 'toy' datasets and real-world complexity [9] and model failures can only become apparent when the models are tested on harder tasks [2]. Benchmarks are stateful, which means that there can be a lottery effect - some research directions get overemphasized over others due to luck [5]. This calls for more difficult and diverse benchmarks.

II. THE CONSTELLATION IMAGES DATASET

For these reasons, we have created Constellation images, [11], which are ambiguous visual stimuli resembling star constellations, where everyday objects (typically one) are represented as a dotted outline among randomly placed dots (see Figure 1). Due to the algorithmic nature of these images, new images with varying difficulty can be easily created, requiring only a source image and a pre-trained segmentation model. Difficulty can be varied by reducing image size, by reducing the amount of injected noise or by lowering the distance between the dots that form the outline. A constantly evolving dataset (a 'living dataset' [5]) can help combat model overfitting since it will be harder for models to find and learn shortcuts. Solving these images can be framed in various ways: solutions can be modeled directly by connecting the dots with primitives such as straight lines or curves (or both), or modeled indirectly, e.g. via training a GAN [10]. The signal-to-noise ratio is open to interpretation (meaning that more or less dots can be ignored), the model has to assume the number, size and position of objects. While a ground truth label exists for these images, there is some amount of interpretability or open-endedness in the

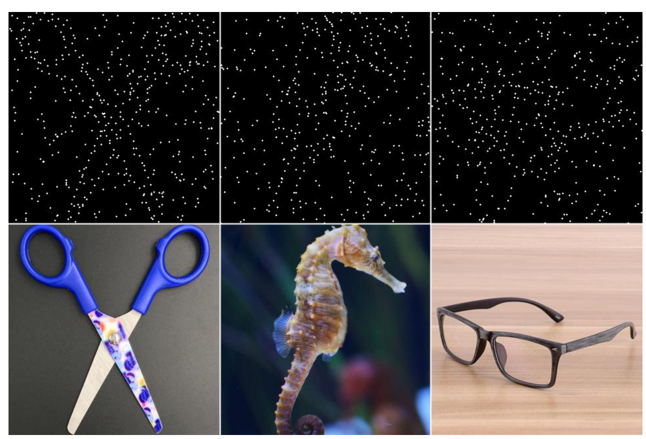

Fig. 1. Examples of Constellation images (top row) and source images (bottom row).

*T. Luik and J. Aru are with Institute of Computer Science, University of Tartu, Estonia taavi.luik at ut.ee, jaan.aru at ut.ee

classification process - the ambiguity present in the images creates the potential for multiple correct answers.

We have the following **research questions** in mind: can the QD framework be used to find the ground truth objects but also plausible alternative solutions? What features are useful to measure diversity and which ones can be learned on-the-fly? How to combine both bottom-up and top-down processing (i.e. guiding search with a specific object class as the target)?

III. Ongoing Work

We took inspiration from work [19] that used an evolutionary algorithm to fit lines on a canvas. An image-text encoder (OpenAI CLIP [16]) was used to guide the placement of the lines so that the output of the two encoders would become more similar. In contrast to that prior work [19], in our case lines are constrained to the dots found on the image.

Since Constellation images can contain hundreds of dots and any pair of dots can form a line, then the number of lines and hence the size of the encoding can rise exponentially. To combat this, we introduce a length threshold - any pair of dots beyond this threshold will not be encoded. As the resulting encoding is discrete, there are no gradients to guide search, necessitating alternative methods to make search more efficient. Ideas from QD framework help but require feature descriptors. We used the MAP-Elites algorithm [13] and with descriptors for line and object counts, but these alone do not suffice. Some prior work in QD has used automated descriptor learning, e.g. with variational autoencoders [7], which we are looking into next.

Constellation images can contain objects from thousands of classes, which means that a population-based search across many generations will do many comparisons, so the CLIP score does not necessarily any longer signal the quality of the solution. Some works have used image augmentations and averaged the scores, however this is costly, since CLIP already forms a bottleneck of the search. CLIP is also known to not work well with out-of-distribution tasks [16]. Prompt engineering plays a crucial role and our current manually engineered prompts can be replaced with learned prompts. CLIP prompts are also brittle - we have seen how punctuation and word order strongly influences predictions, further complicating search.

Lastly, evolutionary algorithms require an initial population, and a popular method is to use random initialization. Even with sparse individuals, with most of the values in the chromosome being zeros, the initial solutions are too random and do not populate enough bins in the archive. We have introduced heuristics that create solutions that are constrained to smaller regions in the image to combat this but ideally such heuristics should be found by the algorithm itself, as we want to keep manual engineering to the minimum.

We are also exploring the benefits of and the implementation of autotelic learning, where an agent is forming its own goals and choosing subtasks to engage with. In order to tackle difficult problems, an agent must focus first on those aspects of the task where (most) progress can be made - the ones that are not too difficult nor too easy. These 'progress niches' are not a characteristic of the environment, but a particular agent at a particular time. Focusing on these subtasks creates a causal link between curiosity and learning, leading to a positive feedback loop [14].

IV. Funding

This research was supported by the European Social Fund through the "ICT programme" measure and the Estonian Research Council grant PSG728.

References

[1] Maxime Allard et al. "Hierarchical Quality-Diversity for Online Damage Recovery". Apr. 12, 2022. arXiv: 2204.05726 [cs].

[2] Jeffrey S. Bowers et al. "Deep Problems with Neural Network Models of Human Vision". In: *Behavioral and Brain Sciences* (Dec. 1, 2022), pp. 1–74. URL: https://www.cambridge.org/core/product/identifier/S0140525X22002813/type/journal_article (visited on 12/05/2022).

[3] Konstantinos Chatzilygeroudis et al. "Quality-Diversity Optimization: A Novel Branch of Stochastic Optimization". Dec. 16, 2020. arXiv: 2012.04322 [cs, math, stat].

[4] Antoine Cully and Yiannis Demiris. "Quality and Diversity Optimization: A Unifying Modular Framework". May 12, 2017. arXiv: 1708.09251 [cs].

[5] Mostafa Dehghani et al. *The Benchmark Lottery*. July 14, 2021. arXiv: arXiv:2107.07002. preprint.

[6] Adrien Ecoffet et al. "First Return, Then Explore". In: *Nature* 590.7847 (Feb. 25, 2021), pp. 580–586.

[7] Adam Gaier, Alexander Asteroth, and Jean-Baptiste Mouret. "Discovering Representations for Black-box Optimization". In: *Proceedings of the 2020 Genetic and Evolutionary Computation Conference* (June 25, 2020), pp. 103–111. arXiv: 2003.04389.

[8] Daniele Gravina, Antonios Liapis, and Georgios N. Yannakakis. "Quality Diversity Through Surprise". In: *IEEE Transactions on Evolutionary Computation* 23.4 (Aug. 2018), pp. 603–616. arXiv: 1807.02397.

[9] Klaus Greff, Sjoerd van Steenkiste, and Jürgen Schmidhuber. "On the Binding Problem in Artificial Neural Networks". Dec. 9, 2020. arXiv: 2012.05208 [cs].

[10] Farid Hasanov. "Iterative versus Amortized Inference Solutions to the Constellation Problem". 2022. URL: https://comserv.cs.ut.ee/ati_thesis/datasheet.php?id=75042.

[11] Tarun Khajuria et al. "Constellations: A Novel Dataset for Studying Iterative Inference in Humans and AI". In: *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). June 2022, pp. 5138–5148.

[12] Joel Lehman and Kenneth O. Stanley. "Abandoning Objectives: Evolution Through the Search for Novelty Alone". In: *Evolutionary Computation* 19.2 (June 2011), pp. 189–223.

[13] Jean-Baptiste Mouret and Jeff Clune. "Illuminating Search Spaces by Mapping Elites". Apr. 19, 2015. arXiv: 1504.04909 [cs, q-bio].

[14] P.-Y. Oudeyer, J. Gottlieb, and M. Lopes. "Intrinsic Motivation, Curiosity, and Learning". In: *Progress in Brain Research*. Vol. 229. Elsevier, 2016, pp. 257–284.

[15] Diego Perez-Liebana et al. *Generating Diverse and Competitive Play-Styles for Strategy Games*. June 28, 2021. arXiv: arXiv:2104.08641. preprint.

[16] Alec Radford et al. "Learning Transferable Visual Models From Natural Language Supervision". Feb. 26, 2021. arXiv: 2103.00020 [cs].

[17] Inioluwa Deborah Raji et al. *AI and the Everything in the Whole Wide World Benchmark*. Nov. 26, 2021. arXiv: arXiv:2111.15366. preprint.

[18] Kenneth Stanley et al. "Designing Neural Networks through Neuroevolution". In: *Nature Machine Intelligence* 1 (Jan. 7, 2019).

[19] Yingtao Tian and David Ha. "Modern Evolution Strategies for Creativity: Fitting Concrete Images and Abstract Concepts". Sept. 18, 2021. arXiv: 2109.08857 [cs].