

Towards Stepping Stones in Goal-Directed Search: Improving Image Perceptual Similarity by combining Evolution, Backpropagation, and Fourier Transforms

Jackson Dean, Nick Cheney

University of Vermont
ncheney@uvm.edu

Abstract

Automatically identifying the stepping-stones that will lead to a predetermined final solution presents a significant challenge for optimization algorithms, yet is essential for solving complex problems. This study, inspired by Picbreeder, investigates a variation of NeuroEvolution of Augmenting Topologies (NEAT) which aims to perform an image replication task using Compositional Pattern Producing Networks (CPPNs) without a human in the loop. This challenge is central to many similar problems in evolutionary computation and artificial life, where identifying key intermediate goals known as stepping-stones is crucial but difficult, often requiring precise fine-tuning of solutions. We leverage techniques from deep learning computer vision research: a fitness function based on perceptual-similarity to help avoid deceptive optima, Fourier features to diversify the CPPNs' inputs, and gradient-based backpropagation to balance the exploration of evolutionary search with goal-directed exploitation. Backpropagation has the additional benefit of smoothing the fitness landscape of topological network mutations. Our results indicate that combining these approaches with CPPN-NEAT yields more diverse and higher fitness solutions compared to traditional NEAT. This hybrid method not only preserves diversity, but also leverages the strengths of both evolutionary algorithms and gradient descent to achieve more detailed and accurate image generation. We speculate that this is a promising avenue for algorithm design, where exploiting gradient information can be balanced with maintaining robust diversity in the search process.

Introduction

One of the most intriguing aspects of evolutionary algorithms is their ability to mimic natural evolution. They are capable of discovering diverse *stepping stones* in complex problem domains requiring hierarchical solutions, like neuroevolution and image generation. Even so, identifying solutions that serve as intermediate stepping stones towards a complex final solution is a significant challenge for optimization algorithms because these stepping stones may not look like the final solutions or be on a straight-line path towards that ultimate goal (Lehman and Stanley, 2011a; Woolley and Stanley, 2011; Auerbach and Bongard, 2011). A notable solution to this challenge can be seen in

Picbreeder (Secretan et al., 2008) as users guide the evolution of Compositional Pattern Producing Networks (CPPNs) through selective breeding to generate complex and compelling images. Users often stumble upon stepping stones in an intuitive and serendipitous manner that is difficult to replicate in an automated, goal-directed algorithm. These stepping-stones need to be appropriately diverse in order to represent the full range of potential solutions, while also being high enough quality to enable future improvement and to be viable candidate solutions in themselves.

Evolutionary methods without a human in the loop like NeuroEvolution of Augmenting Topologies (NEAT) are often successful in evolving diverse structures, but they can stumble in navigating the vast solution space to identify architectures that not only perform well but also have potential for significant local improvement. CPPN-NEAT attempts to automatically identify useful stepping stones but fails to do so in the complex problem domain of image matching, because stepping stones have very different pixel patterns than the final target (Gaier et al., 2019).

On the other hand, humans are able to classify images based on *perceptual similarity*. In Picbreeder, this allows users to identify which CPPNs are likely to lead to future improvements. Picbreeder also enables many people to collaborate on the evolution of a given image by branching from one another's evolutionary paths. This type of collective intelligence can be seen as a form of goal-switching, where the identification of stepping stones is made possible by including various perspectives and objectives. While this is an effective approach, Picbreeder scales poorly due to the large amount of time and people involved, suggesting the need for an alternative approach for automated algorithms.

In this work, we use a fitness function from deep learning computer vision research: Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al., 2018). Unlike traditional pixel-matching functions like mean squared error between the raw image and target, LPIPS more closely matches human perception by comparing latent feature representations of the current and target images and has been used successfully in deep learning for image generation (Esser et al.,

2021; Rombach et al., 2022; Yi et al., 2019).

The introduction of open-ended novelty-seeking and Quality Diversity algorithms has provided another way to find stepping stones by exploring a broad spectrum of novel solutions during search. These methods strive to discover promising stepping stones by promoting diversity and may combine this with a pressure to be fit within a wide variety of behavioral/phenotype niches (Pugh et al., 2016; Lehman and Stanley, 2011b; Pugh et al., 2015; Cully and Demiris, 2017; Lehman et al., 2008; Secretan et al., 2008; Stanley et al., 2017; Soros and Stanley, 2014; Stanley et al., 2016).

These methods explicitly seek out diversity across the different individuals within a population in hopes of maintaining enough breadth to stumble upon useful stepping stones. Borrowing again from deep learning computer vision research, we suggest that an alternative way to foster diversity is to provide a greater breadth of input features to each CPPN, enabling greater feature diversity within a single individual. In particular, using Fourier features as inputs to a CPPN (Tancik et al., 2020) expands the set of images that these models are capable of representing by improving their ability to encode high-frequency details, a key limitation of traditional CPPNs. This enables more diverse subnetworks to grow within the model and allows evolution to select from any one of the many subnetworks that are most useful for a given image, reminiscent of the vast diversity of possible solutions available within a single overparameterized neural network, as highlighted by the lottery ticket hypothesis (Frankle and Carbin, 2018).

While diversity is an essential part of identifying useful stepping stones, exploration alone is likely to yield solutions close to promising stepping stones without fully exploiting them to discover the most viable solutions. Quality diversity algorithms aim to balance both this exploration and exploitation (Lehman et al., 2008; Mouret and Clune, 2015). However the integration of gradient-based optimization techniques presents a compelling avenue for addressing this limitation. By backpropagating error, we can pinpoint and refine useful stepping stones, potentially revealing intermediary solutions that would otherwise remain undiscovered due to their low fitness. This approach is efficient at fine-tuning neural network weights and offers a potent mechanism for rapid, local improvements.

Gradient-based optimization alone typically encourages less diversity and more goal-directedness compared to traditional evolutionary algorithms. This shift can lead to quicker convergence but risks missing good solutions by over-exploiting local optima. Therefore, combining backpropagation with NEAT allows for each method to overcome the limitations of the other. This synergy leverages the strengths of both evolution and gradient-based learning. Ideally, the former excels in discovering innovative, diverse architectures and the latter meticulously optimizes these structures to realize their potential as stepping-stones.

By using evolution to explore diverse CPPN network topologies and backpropagation to finetune each network's weights to its new topology after each structural mutation, we also implicitly generate a form of hierarchical search, or psuedo-meta-learning. This approach enables evolution to better assess the potential solution space around a large structural mutation to the network. Consequently, it becomes easier to discover novel and diverse stepping stones via the topology-mutation-then-weight-backpropagation integrated step than by hoping for subsequent mutations to perform both exploration and exploitation in the areas around this potential stepping stone.

While CPPN-NEAT has been used to great effect in numerous artificial life applications, including the evolution of virtual creatures and evolutionary robotics (Auerbach and Bongard, 2010, 2012; Blackiston et al., 2021; Cheney et al., 2013; Clune et al., 2011; Lai et al., 2021; Pugh et al., 2017; Veenstra and Glette, 2020), we aim to investigate the evolutionary and optimization dynamics of CPPN-NEAT to be able to better navigate the complex fitness landscapes associated with artificial life.

Additionally, performing experimentation in the space of image generation, we are interested in the artistic application of these algorithms. The search for novelty is not only scientifically important, but also necessary for creativity and generative art. Integrating perceptual assessment, Fourier features, and backpropagation with NEAT not only enhances the algorithm's ability to navigate through complex, high-dimensional problem spaces but also significantly enriches its capacity to generate a wide array of unique, visually compelling patterns. By striking a balance between evolutionary diversity and the directedness of gradient-based optimization, we hope to introduce a method for generating art that is not only novel but also detailed and responsive to specific objectives.

Therefore, the primary contributions of this work are to:

- incorporate a deep-neural-network based perceptual fitness function that is not commonly used for evolutionary computation.
- explore the use of additional Fourier feature inputs to overcome the bias towards few low-frequency input signals in CPPN-NEAT.
- integrate backpropagation of error into the NEAT algorithm to set the weights for candidate CPPN architectures for image generation.
- demonstrate that population diversity can be maintained even with the increased exploitation offered by gradient-based optimization.
- generate a diversity of detailed visually compelling patterns with CPPNs, despite the use of a traditionally challenging goal-directed target pattern reconstruction tasks, instead of the divergent open-ended image generation setup employed in the original collective human-in-the-loop Picbreeder framework.

Related work

NeuroEvolution of Augmenting Topologies (NEAT) (Stanley and Miikkulainen, 2002) is a genetic algorithm that evolves neural networks by mutating both their topology and weights. NEAT uses explicit fitness sharing within species to maintain diversity while exploiting discovered optima. In order to overcome the competing conventions problem, NEAT uses historical markings to quantify the compatibility between genomes without expensive topological analysis.

An effective genetic encoding to use within NEAT is the **Compositional Pattern Producing Network (CPPN)** (Stanley, 2007). CPPNs are generative (indirect) encodings that compose various activation functions in a feed-forward artificial neural network to transform geometric spatial data. In the case of image generation, the input spatial data are coordinates, and the output data are the R, G, B values at each coordinate. CPPNs compress high-dimensional image data allowing evolution to search in fewer dimensions.

Due to their unconstrained topology and wide expressive capability, CPPNs make great representations for open-ended search processes like those with Interactive Evolutionary Computation, such as Picbreeder (Secretan et al., 2008). Furthermore, there has been notable progress in developing differentiable versions of CPPNs. For instance, **Differentiable Pattern Producing Networks (DPPNs)** introduced by Fernando et al. (2016), have evolvable architectures and differentiable weights. Similarly, Mordvintsev et al. (2018) applied differentiable CPPNs to visualize deep neural network features, adopting a fixed convolutional structure that does not evolve. Adding differentiability opens new avenues for the generative potential of CPPNs, greatly enhancing their utility in complex image generation.

Previous attempts have been made to integrate this efficient gradient-based optimization with the diverse evolutionary search of NEAT. Chen and Alahakoon (2006) introduced **Learning-NEAT**, which backpropagates error to perform efficient local search within NEAT. Ha (2016) also developed a similar methodology called **Backprop NEAT** and solved simple regression problems by backpropagating error through networks evolved with NEAT. His implementation of NEAT used k -medoids to speciate genomes instead of the method from Stanley and Miikkulainen (2002). We empirically found the original speciation methodology outperformed k -medoids and was used for the reported data. Ha also incorporated interactive evolutionary computation, but in this work we attempt to develop a system for balancing exploitation and exploration without a human in the loop.

Although CPPNs are powerful encodings, they tend to be biased towards low-frequency signals. **Fourier features** (Tancik et al., 2020) help to overcome this bias. They resemble positional encodings and contribute to higher performance on computer vision tasks. These features map the input coordinates of an image through a Fourier transformation, effectively converting spatial coordinates into a

rich, high-dimensional space of sine and cosine functions. This transformation allows the network to more easily learn high-frequency functions by representing these complex patterns through simple transformations. These features are a natural fit for coordinate networks like CPPNs. **Fourier-CPPNs** (Tefaldet et al., 2019) apply Fourier features to CPPNs but, unlike in this work, they use the convolutional approach from (Mordvintsev et al., 2018).

Our work is the first to combine all of these approaches into one algorithm. We take the traditional CPPN-NEAT algorithm from Stanley (2007) and incorporate both gradient-based optimization and Fourier features. We evolve our CPPNs to solve a high-dimensional image-replication problem where solution diversity is desirable and use a deep neural network perceptual fitness function that more closely matches human evaluation.

Methods

We modify the classic NEAT algorithm to enhance the capabilities of CPPNs through Fourier features, meta-learning, backpropagation, and an appropriate fitness function. We use Fourier features to expand the representational range of CPPNs to include high-frequency signals essential for encoding diverse images. Then we leverage backpropagation to fine-tune CPPN weights for efficient local optimization. We adopt a meta-learning framework that separates the optimization process into an outer loop for evolving network structures and an inner loop for refining those structures' weights through backpropagation. Additionally, we employ the Learned Perceptual Image Patch Similarity (LPIPS) metric as a fitness function, chosen for its ability to guide search through evolutionary stepping-stones while avoiding deceptive local optima often associated with pixel-matching.

CPPN-NEAT

We used a variation of the NEAT-Python library McIntyre et al. (2024) that we modified to allow error to be backpropagated through networks on the GPU using Pytorch's Autograd Paszke et al. (2017). CPPNs started with no hidden nodes and 1% of initial possible input to output connections. The input coordinates and targets were scaled to 64 pixels squared. There was a 15% chance to add a node and a 15% chance to add a connection to offspring CPPNs. Additionally, a number of connections equal to three plus 1% of each offspring's connections were added to the offspring and 1% of connections were removed, starting with the lowest magnitude connections. Three connections were added to each offspring to accelerate the increase in network complexity. Connections with magnitudes below $1e-3$ were also removed. This approach allows NEAT to grow networks much more quickly than is typical (see Fig. 5).

Fourier Features

CPPNs traditionally only transform a few fixed spatial dimensions and can be biased towards low-frequency patterns (Tesfaldet et al., 2019). By providing additional Fourier transformations of the input coordinates, we can bias CPPNs to represent a wider range of frequencies, including the high frequency patterns found in image data. The features are generated using a different random seed for each run and are added to the CPPNs as additional input nodes.

Fourier features work well in combination with backpropagation since the additional inputs require additional parameters, growing the networks beyond the size of what can be effectively evolved via random weight mutations. By using backpropagation for the weight updates, selection in the outer loop can focus exclusively on topology, creating a simpler search landscape for NEAT despite having larger networks. The Fourier feature mapping of input x is given by

$$\gamma(x) = [\sin(2\pi Bx), \cos(2\pi Bx)]^T$$

where B is the matrix of frequencies sampled from $\mathcal{N}(0, \sigma^2)$ and σ is chosen by a hyperparameter sweep. We used $\sigma = 2.0$ and 64 features for all trials with Fourier features.

Meta-learning

We employed an inner loop/outer loop framework similar to meta-learning approaches, where the inner loop focuses on rapid adaptation to the problem domain and the outer loop searches for structures that learn effectively. This framework facilitates a dynamic balance between exploring diverse solutions and exploiting known good strategies to navigate complex fitness landscapes. In the outer loop, we leverage the capability of CPPN-NEAT to evolve structure without expensive topological analysis. This process facilitates fast adaptation as in meta-learning, as NEAT selects for individuals that quickly learn to perform well. The inner loop, on the other hand, focuses on the detailed optimization and refinement of individual solutions using backpropagation. By structuring the optimization process into two loops, the algorithm systematically explores the search space to identify promising regions and then exploits these regions to refine solutions that are not only locally optimal but also diverse across the fitness landscape.

Backpropagation

We use backpropagation to fine-tune each new CPPN's weights. While the evolutionary aspect of NEAT allows the exploration of new solution spaces through topological mutations, backpropagation allows for the precise adjustment of solutions via gradient descent. We use the Adam optimizer (Kingma and Ba, 2017) with a learning rate of $\gamma = 0.01$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 1e-8$. Backpropagation was applied for 2000 steps or until a CPPN's loss did not decrease for 5 steps. Weights were decayed with an L2 penalty of $3e-3$.

The integration of backpropagation into the evolutionary process allows the algorithm to exploit gradient information, enhancing optimization efficiency and leveraging the benefits of modern GPU hardware. This is particularly effective in high-dimensional optimization landscapes, where the gradient can guide search towards promising regions that might be difficult to discover through random mutations alone.

Fitness Function

We used the Learned Perceptual Image Patch Similarity (LPIPS) metric (Zhang et al., 2018) as the fitness function in these experiments. LPIPS uses the difference between features extracted by a trained deep neural network from the target and candidate image to compare perceptual similarity. Previous research has found that CPPN-NEAT struggles to replicate images when using naive similarity metric such as mean squared error (Gaier et al., 2019). Fitness was measured as $1 - \text{LPIPS}(x)$ so that higher values were better, while $\text{LPIPS}(x)$ was used for backpropagation. Our goal when choosing a fitness function was to create a fitness landscape that facilitates the discovery of stepping-stones. A perceptual fitness function, as opposed to a pixel-matching function, has a better chance of rewarding intermediate adaptations that might not directly resemble the target image but are perceptually closer to the target. In this case, slight genetic changes can result in significant but constructive phenotypic variation, rather than catastrophic failure as in the case of pixel-matching. We compared LPIPS to Mean Squared Error (MSE) to test this hypothesis.

We aim to leverage the nuanced measure of image similarity that deep neural networks offer, moving beyond the limitations of traditional metrics that do not account for how humans perceive visual information. While common in deep learning computer vision tasks, LPIPS is rarely used in evolutionary computation or artificial life.

Results

We tested four configurations of NEAT which incrementally incorporate more of our suggested improvements. As a baseline, we test the original CPPN-NEAT algorithm Stanley (2007). Then, we test CPPN-NEAT with Fourier features followed by CPPN-NEAT with backpropagation. Finally, we test a configuration incorporating both Fourier features and gradient-based optimization. All configurations were tested on the five images in Figure 1 and quantitative results were averaged over 40 replication trials. Significance was measured using the Wilcoxon rank-sum test and $\alpha = 0.05$.

Replication Success

To evaluate the efficacy of our proposed modifications to CPPN-NEAT, we systematically compare NEAT configurations on their perceptual similarity to target images.

NEAT with both Fourier features and gradient-based optimization ($\text{LPIPS} = 0.71 \pm 0.01$) consistently matched the




Figure 1: The images for representatives of the top four species from the best run of each NEAT configuration. The first column shows the target images. The second column shows a naive classic NEAT approach using MSE as the fitness function, producing simple, blurry images. NEAT-Classic and NEAT-SGD usually find the correct general shape but lack high-frequency variation and detail. NEAT with Fourier features finds more variation in frequencies, but lacks details. Finally, NEAT with both Fourier features and SGD finds both the correct shape and some details, while maintaining interesting diversity between species and producing the closest LPIPS match for all targets ($p < 0.01$). Target images other than the sunrise are from Picbreeder.

target images more closely than other configurations ($p < 0.001$, Fig. 2, 1). As found by Chen and Alahakoon (2006), NEAT with backpropagation (0.62 ± 0.02) performs better than the original NEAT (0.44 ± 0.01). Adding Fourier features (0.43 ± 0.01) to CPPN inputs decreased performance compared with the original NEAT ($p < 0.001$). Classic NEAT with mean squared error fitness fails to escape an early local optima and produces simple, blurry images.

Images in Figure 2 are from runs with 200 generations each. Therefore, the runs with backpropagation do much

more computation than NEAT-Classic and NEAT-FFT, as there is an average of $3202.27 (\pm 6.45)$ forward passes per generation in NEAT with both Fourier features and backpropagation. To ensure this did not bias our results, we also tested the conditions without backpropagation for an equal number of CPPN forward passes by extending those runs to 6,500 generations. We find that in this case, NEAT-FFT (0.51 ± 0.03) and NEAT-Classic (0.50 ± 0.03) still perform much worse than NEAT with backpropagation and NEAT with both backpropagation and Fourier features.




Figure 2: The fitness of the best solution found over generations with 95% bootstrapped confidence intervals. Adding Fourier features to NEAT without backpropagation hurt performance. Backpropagation resulted in much higher performance than NEAT with only random mutations. The most effective strategy was using both Fourier features and backpropagation, as networks were more expressive with the additional inputs and backpropagation was able to exploit that larger input space.

In trials with backpropagation, the inner loop acts as a tool for improving the specificity of fitness evaluations in the outer loop. By optimizing the new offspring network weights, the outer loop receives more accurate information about the quality of those solutions and makes informed selection decisions based on a network’s ability to learn.

Figure 1 shows qualitative results from the best runs of each configuration and target. In the case of the skull image, NEAT-SGD finds visual approximations of the target, although their LPIPS score is still lower than NEAT with both Fourier features and SGD on average. The dolphin image proved consistently difficult to replicate, and the butterfly image was difficult to replicate for all configurations other than NEAT with both Fourier features and backpropagation. Certain local optima were very difficult for NEAT-Classic to overcome, for example the red circle in the case of the apple and the white background in the case of the butterfly. Notably, the images from NEAT with backpropagation have the most visual diversity (see Fig. 3). As predicted, runs using MSE as a fitness function got stuck in early local optima, resulting in simple, blurry images.

Fourier Features We find that NEAT using CPPNs with Fourier features performs worse than the original CPPN-NEAT ($p < 0.001$, Fig. 2). However, with the addition of gradient-based optimization the opposite is true and we see an increase in fitness compared to NEAT with just backprop-

agation ($p < 0.001$). This suggests that simply increasing the size of the input space, and by extension the number of parameters (see Fig. 5), is not enough to increase the representational capacity of CPPNs. The combination of richer inputs with an appropriate optimization technique provides the greatest benefit. Additionally, we find that backpropagation is useful for identifying which input features are the most valuable for a given target image (see Fig. 5).

Diversity

With backpropagation, one might expect the whole population to converge more quickly towards particular solutions, resulting in less diversity. Here, we investigate the genetic and phenotypic diversity within populations to understand how different configurations of NEAT influence the variety of solutions. We plot the overall effect of the NEAT configuration on diversity in Figure 3. Perhaps surprisingly, backpropagating error to update the CPPN weights results in increased genetic diversity. We find that NEAT with backpropagation (2.88 ± 0.02) and NEAT with backpropagation and Fourier features (2.83 ± 0.02) had more diversity than classic NEAT (2.62 ± 0.03) and NEAT with Fourier features (2.59 ± 0.03 , $p \leq 0.01$). There was not a significant difference at the end of evolution associated with adding Fourier features to classic NEAT ($p = 0.19$).

Figure 1 demonstrates that without backpropagation and Fourier features, the best runs often converge to simple shapes. These runs lack genetic diversity and their phenotypes have converged to more similar images as measured by LPIPS. We measured the pairwise similarity at six points throughout evolution and report the results in Figure 3. NEAT-SGD had the highest phenotypic diversity at the end of evolution (0.59 ± 0.01 , $p < 0.01$). NEAT with both Fourier features and gradient descent (0.55 ± 0.02) had the second most phenotypic diversity, tied with classic NEAT (0.54 ± 0.01 , $p = 0.40$). NEAT-FFT (0.47 ± 0.02) had the least phenotypic diversity ($p < 0.01$). Taken together, these results indicate that adding backpropagation does not cause premature convergence of phenotypes and, in fact, adding just gradient-based optimization can directly increase phenotypic diversity over classic NEAT. We speculate that backpropagation increases robustness, enabling a wider range of mutations to become useful since any given mutation does not need to simultaneously discover new topology and appropriate weights for that topology, as in NEAT.

We observed that evolution, even with the exploitative convergence of gradient-based optimization, discovers distinct phenotypic variations. This mirrors organic life, in which diversity emerges amidst selective pressures. In a different but related setup using the butterfly image in Figure 4, the overall shape of the butterfly is consistent between individuals while there is notable variation in coloration and pattern, akin to the diversity observed in actual butterfly populations. Unlike the runs presented in other figures, CPPNs




Figure 3: Mean diversity by generation, measured using the historical marking genetic distance from Stanley and Miikkulainen (2002) on the left and pairwise LPIPS on the right with 95% bootstrapped confidence intervals. Trials using backpropagation had more genetic diversity than those without. There was no significant effect of Fourier features on genetic diversity. For the phenotypes, backpropagation increased diversity over classic NEAT. NEAT with both Fourier features and backpropagation had the same phenotypic diversity at the end of evolution as classic NEAT, despite having a much higher fitness (see Fig. 2).

in this run started with 50% of connections filled at the start of evolution and 8 initial hidden nodes for an average 268 initial connections, much larger than the CPPNs in Figure 1.




Figure 5: The number of connections per network over generations. We see that despite the number of hidden connections growing at a similar rate between NEAT with both Fourier features and backpropagation as NEAT with just Fourier features, the number of input connections grows less quickly with backpropagation, suggesting that unhelpful input connections are being decayed and pruned. Runs with Fourier features have more input nodes and therefore start with more connections.

Evolved Architectures

We see above that the inclusion of Fourier features and backpropagation improves performance, but it's not clear how the additional features and the ability to easily ignore a given connection with backpropagation may affect the topology of the evolved networks. In Figure 5, we report the number of connections, measured at six different generations. All networks grew quickly due to the three extra connections added to each offspring, but the CPPNs with Fourier features have

more input nodes and therefore a higher number of valid connections, facilitating their faster growth.

We split the data into connections originating from inputs and those originating from hidden nodes. In NEAT with both backpropagation and Fourier features, we see that the number of input connections grows more slowly than it does in NEAT with just Fourier features. NEAT with backpropagation and NEAT with both Fourier features and backpropagation use L2 regularization and magnitude-based pruning to remove unused connections. As a result, input features that are unhelpful for the target are disconnected from the network as their outgoing connections are pruned. This highlights a major benefit of incorporating gradient information in network topology search problems. Not only are the weights fine-tuned, but the networks are also simplified, creating a smoother search landscape for NEAT.

Discussion

Our results affirm that a perceptual similarity function like LPIPS is capable of guiding evolution through this complex search landscape to discover diverse, high-quality images. Unlike pixel-matching, LPIPS selects for features that are more aligned with human visual perception and reward intermediate solutions that may be perceptually similar to the target even if they are not pixel-by-pixel matches.

We also showed that including Fourier features as inputs to CPPNs enhances their capability to represent high-frequency patterns in images. However, the results indicate that simply adding Fourier features without gradient-based optimization degrades performance compared with the original NEAT. This result highlights the necessity of coupling feature augmentation with appropriate optimization techniques to fully leverage the increased representational capacity afforded by those features.




Figure 4: Images from the top genomes of one run using the butterfly target with both Fourier features and backpropagation. This run used a higher resolution and different hyperparameters than those in Figure 1. Images are organized in descending fitness order from the top left to the bottom right. The significant variation in the phenotypes of the population suggest the potential for the wide exploration of the search space around this solution and/or the evolvability of diverse artistic forms for generative art. Some features, like the butterfly’s shape, are canalized in the population while the wing patterns and colors vary between individuals.

The ability of CPPNs to recover from significant network topology mutations showcases the robustness of this approach. Adaptability is crucial for maintaining a diverse population capable of exploring various solution pathways as some CPPNs are more able to maintain their performance between topology mutations. We set out to establish a robust framework that balances exploration and exploitation and pave the way for new applications in artificial life, evolutionary art, and beyond, where the complexity and creativity of solutions are paramount. Although our experiments were focused on image generation, many of the characteristics of this problem are shared with research on virtual creatures and robotics. Notably, these problems have deceptive fitness landscapes with difficult-to-find stepping stones and are benefited by maintaining robust population diversity.

Our approach successfully balances exploitation and exploration by integrating backpropagation with CPPN-NEAT. This synergy ensures that the exploration of novel architectures through NEAT is complemented by the precision and efficiency of gradient-based optimization. The inclusion of Fourier features and an appropriate fitness function, coupled with this hybrid optimization strategy, allows for more expressive solutions and enables the generation of detailed im-

ages that were previously out of reach for CPPN-NEAT. We hope that this work will inspire future artificial life and generative art leveraging backpropagation and Fourier features.

Limitations and Future work

The results in this paper are likely dependent on the specific target images, objective function, encoding, and task chosen.

Future work could leverage speciation to explore different parameters for backpropagation. Evolutionary Stochastic Gradient Descent (ESGD) is a population-based optimization algorithm that alternates between SGD and gradient-free evolutionary algorithms (Cui et al., 2018). Much like this work, ESGD leverages the exploration capabilities of evolutionary algorithms and the exploitation efficiency of gradient descent. They allowed evolution to learn different optimizers per individual. That approach could integrate well with an algorithm like NEAT, since each species may benefit from customized learning parameters.

Even though we used a similar evolutionary structure to generative art approaches like Picbreeder, this is not an open-ended algorithm. Given the additional exploitation provided by backpropagation, this approach might be expected to struggle to find stepping-stones or to overfit to the target image, but instead we find that resulting images are both diverse and relatively close to the target. We used a very objective-focused approach in these experiments, but a more open-ended version of this methodology could also produce interesting results. For example, gradient-based optimization could be used to push solutions towards a minimum criterion (Stanley et al., 2016) while continuing to explore novel solutions via evolution. Additionally, the richness provided by Fourier features and backpropagation could be used to improve the methodology of approaches like Innovation Engines (Nguyen et al., 2015) or multi-objective approaches to image replication like Many-objective Optimization via Voting for Elites (Dean and Cheney, 2023).

In this study, we found that our approach is capable of generating high-quality, diverse solutions, but it is left to future research to investigate the evolved structures in more depth. For example, Huizinga et al. (2018) found that Picbreeder produces CPPNs with canalization, hierarchy, and modularity. They found that the dolphin target we used in our study scored low on all three measures and was therefore not robust to mutations, suggesting an explanation for why our approach had a hard time replicating the image. A better understanding of these network characteristics, which are associated with evolvability, would provide further insights into automatic identification of stepping stones.

Acknowledgements

This material is based upon work supported by NSF Grants No. 2008413, 2218063, 2239691. Computations were performed on the Vermont Advanced Computing Core supported in part by NSF award No. 1827314.

References

- Auerbach, J. E. and Bongard, J. C. (2010). Evolving cppns to grow three-dimensional physical structures. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pages 627–634.
- Auerbach, J. E. and Bongard, J. C. (2011). Evolving complete robots with cppn-neat: the utility of recurrent connections. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 1475–1482.
- Auerbach, J. E. and Bongard, J. C. (2012). On the relationship between environmental and morphological complexity in evolved robots. In *Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation, GECCO '12*, page 521–528, New York, NY, USA. Association for Computing Machinery.
- Blackiston, D., Lederer, E., Kriegman, S., Garnier, S., Bongard, J., and Levin, M. (2021). A cellular platform for the development of synthetic living machines. *Science Robotics*, 6(52):eabf1571.
- Chen, L. and Alahakoon, D. (2006). Neuroevolution of augmenting topologies with learning for data classification. In *2006 International Conference on Information and Automation*, pages 367–371.
- Cheney, N., MacCurdy, R., Clune, J., and Lipson, H. (2013). Unshackling evolution: Evolving soft robots with multiple materials and a powerful generative encoding.
- Clune, J., Lipson, H., et al. (2011). Evolving three-dimensional objects with a generative encoding inspired by developmental biology. In *ECAL*, pages 141–148.
- Cui, X., Zhang, W., Tüske, Z., and Picheny, M. (2018). Evolutionary stochastic gradient descent for optimization of deep neural networks.
- Cully, A. and Demiris, Y. (2017). Quality and diversity optimization: A unifying modular framework. *IEEE Transactions on Evolutionary Computation*, 22(2):245–259.
- Dean, J. and Cheney, N. (2023). Many-objective optimization via voting for elites. In *Proceedings of the Companion Conference on Genetic and Evolutionary Computation*, pages 131–134.
- Esser, P., Rombach, R., and Ommer, B. (2021). Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883.
- Fernando, C., Banarse, D., Reynolds, M., Besse, F., Pfau, D., Jaderberg, M., Lanctot, M., and Wierstra, D. (2016). Convolution by evolution: Differentiable pattern producing networks. *CoRR*, abs/1606.02580.
- Frankle, J. and Carbin, M. (2018). The lottery ticket hypothesis: Training pruned neural networks. *CoRR*, abs/1803.03635.
- Gaier, A., Asteroth, A., and Mouret, J.-B. (2019). Are quality diversity algorithms better at generating stepping stones than objective-based search? In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 115–116.
- Ha, D. (2016). Neural network evolution playground with back-prop neat. *blog.otoro.net*.
- Huizinga, J., Stanley, K. O., and Clune, J. (2018). The emergence of canalization and evolvability in an open-ended, interactive evolutionary system. *Artificial life*, 24(3):157–181.
- Kingma, D. P. and Ba, J. (2017). Adam: A method for stochastic optimization.
- Lai, G., Leymarie, F. F., Latham, W., Arita, T., and Suzuki, R. (2021). Virtual creature morphology-a review. In *Computer Graphics Forum*, volume 40, pages 659–681. Wiley Online Library.
- Lehman, J. and Stanley, K. (2011a). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223.
- Lehman, J. and Stanley, K. O. (2011b). Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 211–218.
- Lehman, J., Stanley, K. O., et al. (2008). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE*, pages 329–336.
- McIntyre, A., Kallada, M., Miguel, C. G., Feher de Silva, C., and Netto, M. L. (2024). neat-python.
- Mordvintsev, A., Pezzotti, N., Schubert, L., and Olah, C. (2018). Differentiable image parameterizations. *Distill*, 3(7):e12.
- Mouret, J.-B. and Clune, J. (2015). Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*.
- Nguyen, A., Yosinski, J., and Clune, J. (2015). Innovation engines: Automated creativity and improved stochastic optimization via deep learning. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, pages 959–966.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. (2017). Automatic differentiation in pytorch.
- Pugh, J. K., Soros, L., and Stanley, K. O. (2016). Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI*, page 40.
- Pugh, J. K., Soros, L. B., Frota, R., Negy, K., and Stanley, K. O. (2017). Major evolutionary transitions in the voxelbuild virtual sandbox game. In *Artificial Life Conference Proceedings*, pages 553–560. MIT Press.
- Pugh, J. K., Soros, L. B., Szerlip, P. A., and Stanley, K. O. (2015). Confronting the challenge of quality diversity. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, pages 967–974.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695.

- Secretan, J., Beato, N., D'Ambrosio, D. B., Rodriguez, A., Campbell, A., and Stanley, K. O. (2008). Picbreeder: evolving pictures collaboratively online. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1759–1768.
- Soros, L. and Stanley, K. (2014). Identifying necessary conditions for open-ended evolution through the artificial life world of chromaria. In *ALIFE 14: The Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 793–800. MIT Press.
- Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8(2):131–162.
- Stanley, K. O., Cheney, N., and Soros, L. B. (2016). How the strictness of the minimal criterion impacts open-ended evolution. In *Artificial Life Conference Proceedings*, pages 208–215. MIT Press.
- Stanley, K. O., Lehman, J., and Soros, L. (2017). Open-endedness: The last grand challenge you've never heard of. *While open-endedness could be a force for discovering intelligence, it could also be a component of AI itself*.
- Stanley, K. O. and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127.
- Tancik, M., Srinivasan, P. P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J. T., and Ng, R. (2020). Fourier features let networks learn high frequency functions in low dimensional domains.
- Tesfaldet, M., Snelgrove, X., and Vazquez, D. (2019). Fourier-cppns for image synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*.
- Veenstra, F. and Glette, K. (2020). How different encodings affect performance and diversification when evolving the morphology and control of 2d virtual creatures. In *Artificial Life Conference Proceedings 32*, pages 592–601. MIT Press.
- Woolley, B. G. and Stanley, K. O. (2011). On the deleterious effects of a priori objectives on evolution and representation. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 957–964.
- Yi, X., Walia, E., and Babyn, P. (2019). Generative adversarial network in medical imaging: A review. *Medical image analysis*, 58:101552.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. (2018). The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595.