
SELF-PROGRAMMING ARTIFICIAL INTELLIGENCE USING CODE-GENERATING LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Recent progress in large-scale language models has enabled breakthroughs in previously intractable computer programming tasks. Prior work in meta-learning and neural architecture search has led to substantial successes across various task domains, spawning myriad approaches for algorithmically optimizing the design and learning dynamics of deep learning models. At the intersection of these research areas, we implement a code-generating language model with the ability to modify its own source code. Self-programming AI algorithms have been of interest since the dawn of AI itself. Although various theoretical formulations of generalized self-programming AI have been posed, no such system has been successfully implemented to date under real-world computational constraints. Applying AI-based code generation to AI itself, we develop and experimentally validate the first practical implementation of a self-programming AI system. We empirically show that a self-programming AI implemented using a code generation model can successfully modify its own source code to improve performance and program sub-models to perform auxiliary tasks. Our model can self-modify various properties including model architecture, computational capacity, and learning dynamics.

1 INTRODUCTION

Artificial intelligence technology has sought to emulate human cognitive capabilities within a computational substrate. Machine learning, and more recently deep learning, has enabled computers to derive programs from relevant data in order to address novel tasks for which they were not explicitly programmed.

Advanced artificial intelligence is theorized to require self-modifying programs that continuously imbue themselves with extended capabilities (Schmidhuber, 2006; 1987), surpassing their pre-programmed functionalities to achieve increased computational and predictive effectiveness. Meta-learning approaches have made significant progress along this avenue, from self-referential evolutionary algorithms (Schmidhuber, 1987) to RNN-based learners (Santoro et al., 2016) and modern gradient-based meta-learners (Finn et al., 2017; Khodak et al., 2019; Li et al., 2017). Adjacent work in neural architecture search (Real et al., 2018; Stanley et al., 2009; Zoph and Le, 2016) has led to the development of various algorithms for automatically optimizing the architectural parameters of neural networks.

In recent years, AI-based computer code generation has seen dramatic progress (Chen et al., 2021; Le et al., 2022; Li et al., 2022; Wang et al., 2021) stemming from breakthroughs in large-scale language processing models (Brown et al., 2020; Devlin et al., 2018; Raffel et al., 2020; Vaswani et al., 2017). Modern language models (LMs) trained on computer programming tasks can not only generate syntactically correct computer code, but can also formulate logically correct computer programs from human-language prompts and understand pre-written programs to derive human-interpretable information.

At the intersection of these paradigms, we set out to explore a completely new avenue of research by combining self-modification approaches with code-generation approaches to achieve self-programming AI. Although theoretical formulations of self-programming AI (Schmidhuber, 2006) have been posed in the past, no practical implementation of a fully self-programming AI has been achieved to date under real-world computational constraints.

In this paper, we implement a LM-based code generation model with the ability to rewrite and improve its own source code, thereby achieving the first practical implementation of a self-programming AI system. With free-form source code modification, it is possible for our model to change its own model architecture, computational capacity, and learning dynamics.

Since this system is designed for programming deep learning models, it is also capable of generating code for models other than itself. Such models can be seen as sub-models through which the main model indirectly performs auxiliary tasks. We explore this functionality in depth, showing that our model can easily be adapted to generate the source code of other neural networks to perform various computer vision tasks. We illustrate our system’s ability to fluidly program other deep learning models, which can be extended to support model development in various other fields of machine learning.

2 RELATED WORK

Meta-learning. Meta-learning is the subfield of machine learning that attempts to "learn how to learn". To be precise, let ω denote an assumption on "how to learn" and $p(\mathcal{T})$ be a task distribution, where a "task" $\mathcal{T} = \{\mathcal{D}, \mathcal{L}\}$ for some dataset \mathcal{D} and loss function \mathcal{L} (Hospedales et al., 2022). Then, the goal of meta-learning can be viewed as computing

$$\min_{\omega} \mathbb{E}_{\mathcal{T}}(\mathcal{L}(\mathcal{D}; \omega)).$$

Approaches in this subfield are broadly classified into RNN-based learners (Santoro et al., 2016), gradient-based meta-learners (Finn et al., 2017; Khodak et al., 2019; Li et al., 2017), and self-referential evolutionary algorithms (Schmidhuber, 1987).

A select body of prior work within meta-learning establishes the conceptual groundwork for theoretical self-programming AI systems. Notably, Gödel machines (Schmidhuber, 2006) pose a theoretical formulation wherein an intelligent system can continuously rewrite itself in a globally optimal way. Schmidhuber (1987) discusses some of the evolutionary concepts that might underpin such systems.

Automated Code Generation. Vaswani et al. (2017) introduced the transformer, an attention-based architecture that alleviated many of the previous issues with prior sequence-to-sequence models. The success of this architecture and subsequent innovations in Natural Language Processing (Brown et al., 2020; Devlin et al., 2018; Raffel et al., 2020) have led to breakthroughs in AI-driven code generation (Chen et al., 2021; Drori et al., 2022; Le et al., 2022; Li et al., 2022; Wang et al., 2021).

Chen et al. (2021) revolutionized the application of language models to code generation. Le et al. (2022) and Wang et al. (2021) extended the transfer learning concepts introduced in Raffel et al. (2020) to create systems capable of efficiently executing multiple tasks, including code generation and understanding. Li et al. (2022) created a system that was able to achieve top-half performances at various code competitions. Drori et al. (2022) showed that a pretrained language model fine-tuned on code was capable of solving, generating, and explaining university level math problems, as well as problems from challenging high-school math competitions. Although these systems are capable of writing programs for diverse and complex tasks, they are not capable of modifying and improving their own source code.

AutoML. AutoML is a nascent and incredibly complex field that aims to automatically generate machine learning algorithms to solve problems without human tuning. Part of the difficulty lies in the fact that intricate, well-crafted neural networks are often necessary to obtain high performance on even mildly nontrivial problems such as CIFAR-100. Thus, the subfield of Neural Architecture Search(NAS) formed with the goal of creating algorithms capable of efficiently generating suitable architectures for a given problem. There are three approaches typically used for NAS: reinforcement learning-based NAS (Jaafra et al., 2019; Zoph and Le, 2016), gradient-based NAS (Shi et al., 2021), and evolutionary NAS (ENAS) (Liu et al., 2021).

Evolutionary algorithms are based upon the principles of Darwinian evolution and natural selection: a population of algorithms are created and evaluated, the "fittest" (best-performing) survive and are used as "parents" for a future generation of algorithms that are then evaluated, and so on (Liu et al., 2021). Some evolutionary approaches include regularized evolution search (Real et al., 2018; 2020), particle swarm optimization (Haidar et al., 2021), and genetic algorithms (Whitley, 2021).

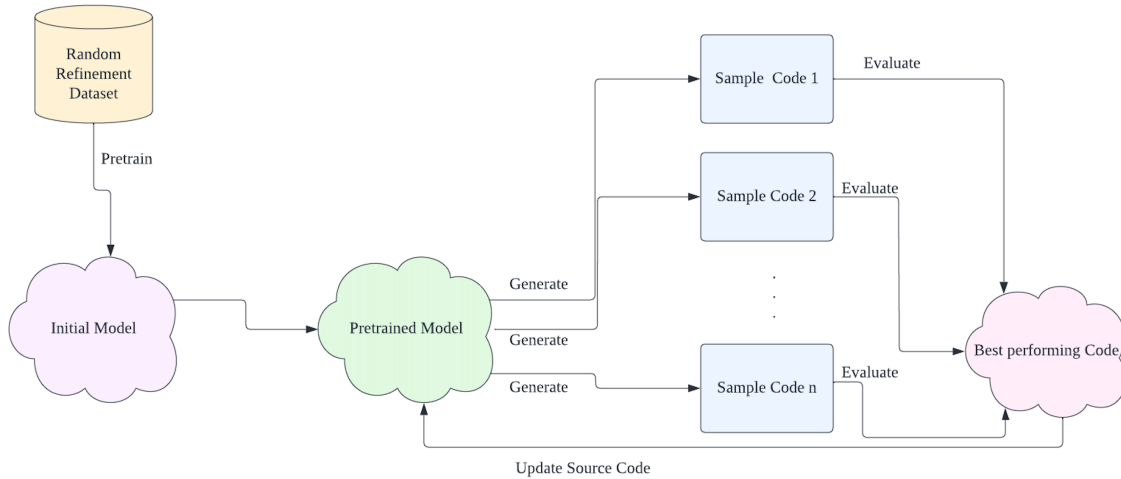


Figure 1: Baseline implementation of our model with a simple genetic algorithm.

While AutoML techniques are usually restricted to one type of network (such as convolutional neural networks), our system is capable of editing and improving all sorts of neural network source code.

Our Contribution. We extend the methodologies of automated code generation as well as some simple evolutionary ideas from the field of AutoML to create the first functional self-programming AI system. In particular, we use a code-generating transformer model based on Raffel et al. (2020) and a simple genetic algorithm (Whitley, 2021).

3 METHODS

3.1 SELF-PROGRAMMING AI

We present the first practical implementation of a self-programming AI system. We initialize an arbitrary language processing model and train it on a synthetically-generated code corpus dataset, henceforth referred to as the "Random Refinement Dataset". Within this dataset, code samples are included for creating and training transformer-based neural networks. The model learns to parse and generate machine learning code, allowing it to understand its own source code. In an iterative process, the model is trained for a period of time before it is queried to generate a new source code for itself. A new iteration of the model is initialized and trained, and the cycle continues (Figure 1).

Our model is initialized with a standard encoder-decoder transformer model based on T5 (Raffel et al., 2020). The model begins with 2 encoder layers, 2 decoder layers, 8 attention heads, and a feed-forward layer width of 1024. The architecture of the model is easily modified by its self-reprogramming in subsequent episodes.

Each episode of the algorithm is comprised of a training stage and a reprogramming stage. The model can be run for an arbitrary number of episodes to continuously self-improve, although its ability to do so is conditional on having an initial dataset with relevant code samples from which it can effectively deduce how to edit its existing source code.

During the training stage, the model is trained on the random refinement dataset. In our experiments, we train the model on a text-to-text "code refinement" task in order for it to learn code modification. Given an initial source code snippet, the model is trained to generate a modified version of that code snippet. The specific modification applied is arbitrary, but the newly generated or "refined" code should be syntactically and logically correct. The text-to-text transformer model outputs probability distributions over tokens, allowing for effective modeling of non-deterministic mappings for the code refinement process.

The reprogramming stage utilizes an algorithm similar to the *simple genetic algorithm* as described in Whitley (2021). During this stage, the trained model is given its own source code c as input

for code refinement. The model is queried to generate n modified versions of its current source code c_1, c_2, \dots, c_n .¹ Each refined source c_i is executed to check for program validity, and the model instantiated within the refined code is quickly trained for one epoch on a subset \mathcal{S} of the training corpus. Each model’s average loss $\mathcal{L}_{\mathcal{S}}(c_i)$ over this "extended few-shot" training run is used as a proxy to evaluate the potential future performance of each refined source code version. The highest-performing source code c_j is kept as the new source code for the model, and the training stage is repeated with the model defined in this source. To be precise, this c_j satisfies

$$\mathcal{L}_{\mathcal{S}}(c_j) = \min(\mathcal{L}_{\mathcal{S}}(c_1), \dots, \mathcal{L}_{\mathcal{S}}(c_n)).$$

Algorithm 1 Self-Reprogramming Algorithm

Require: Default source code that defines a default model and default training algorithm
Ensure: Implements self-reprogramming artificial intelligence

```

source ← default source code;
while Model is running do
  if source defines a model then
    model ← model as defined in source;
  end if
  if source defines a train() function then
    train() ← train() as defined in source;
  end if
  train(model, dataset);
  for number of candidates do
    candidate ← model.generate(source)
    model, train, dataset ← as defined in candidate
    candidate.metrics = train(candidate) on a small subset of the dataset
  end for
  source ← candidate with the best metrics;
end while

```

3.2 PROGRAMMING OTHER AI

The process to generate code for novel neural networks is similar. We initialize an encoder-decoder transformer based on T5 and train it on a random refinement dataset of code samples for arbitrary convolutional neural networks.

Next, the model is queried to generate modifications of an initial source code snippet. In our experiments, this is a network with a single hidden layer of 16 neurons. The possible modifications include adding convolutional layers, changing the size of convolutional or hidden layers, and increasing the number of hidden layers.

During the reprogramming stage, the trained model is given the default network as input for code refinement and is then queried to generate n modified versions of this network. Each modification is first checked for program validity, and is then trained for two epochs on subset of data for a classic neural network problem (such as MNIST or CIFAR-10). The objective of each network is to minimize the cross-entropy loss,

$$\mathcal{L}(\theta) = - \sum_{i=1}^n g_i \log(p_i),$$

where θ represents the model parameters, n is the number of classes, g_i is the ground-truth label, and p_i is the softmax probability outputted for the i -th class. The best-performing algorithm is reused as input to the model for further code refinement, and the process continues.

¹For most trials we used $n = 8$.




Figure 2: Loss curves for text-to-text refinement

4 EXPERIMENTS

4.1 SELF-PROGRAMMING AI

We carry out preliminary experiments evaluating a basic implementation of self-reprogramming AI. We procedurally generate a code refinement dataset for transformer models, and use our algorithm to train language processing models on this dataset in order to rewrite their own neural network source code.

In the random refinement dataset, we sample unrefined source codes for T5-based transformer models. Several model design parameters are randomly generated: number of encoder layers (between 1 and 8 inclusive), number of decoder layers (between 1 and 8), feed-forward layer dimensionality (between 64 and 4096), and number of attention heads (between 1 and 16). The key-value dimensionality is set to be the model dimensionality divided by the number of attention heads.

We "refine" this source code by randomly modifying one of the model design parameters. Either one among the encoder depth, decoder depth, and number of heads is incremented or decremented, or the feed-forward width is randomly increased or decreased by a percentage of 1% to 50%, inclusive. The unrefined source code is used as input to the model, and the refined source is used as the corresponding output label.

In the training stage, the model is trained in a conditional generation configuration to generate refined transformer source codes given an unrefined source code as input. The objective of this procedure is to teach the model to modify design parameters in the source code of its own model class. This "refinement" process is more or less random, but the extended few-shot evaluation method used in the reprogramming stage allows us to achieve performance improvements over multiple iterations of our reprogrammed model.

In these experiments, we use a simple greedy search approach to evaluate and select generated source codes that would potentially improve the code refinement model. The model is queried with its own source code as input, and randomly generates 8 refined source code candidates. The candidates are each evaluated using the extended few-shot method, training for one epoch on the first 1024 samples of the code refinement dataset.

Source code candidates that produce errors are discarded entirely, and the source code candidate with the lowest average training loss in extended few-shot evaluation is kept as the new query code. If all candidates underperform the original source code, then the original source is kept. This new query code is given to the code refinement model as input to generate another batch of source code candidates in the following iteration of the main loop. The main loop runs for 10 iterations of training




Figure 3: Loss curves of multiple iterations of self-reprogramming models on convolutional networks for CIFAR-10 and EMNIST. Each graph represented Model 0 refers to the default model. The value of n is the number of candidates tested prior to source code updating. Loss curves for MNIST do not show much improvement as even the input network was able to achieve high performance.

and reprogramming. The resulting loss curves of our self-reprogramming models are show in Figure 2, and an example of a modification made is in Table 1.

Unrefined Source (Input Text)

```
config = transformers.PretrainedConfig.from_pretrained("Salesforce/codet5-small")
config.num_layers = 2
config.num_decoder_layers = 2
config.d_ff = 1024
model = transformers.T5ForConditionalGeneration(config)
```

Refined Source (Output Text)

```
config = transformers.PretrainedConfig.from_pretrained("Salesforce/codet5-small")
config.num_layers = 2
config.num_decoder_layers = 2
config.d_ff = 1331
model = transformers.T5ForConditionalGeneration(config)
```

Table 1: Example of an input-output pair comprised of an unrefined transformer source code and a corresponding refined source code. In the refined source, the feed-forward layer dimensionality parameter is randomly increased by 30%.

4.2 PROGRAMMING OTHER AI

In the random refinement dataset, we sample unrefined source codes for convolutional neural networks. Several model design parameters are randomly generated: number of hidden layers n (between 1 and 8 inclusive), number of convolutional layers c (between 0 and 8), number of channels h (between 16 and 512), and size of each hidden layer s (between 16 and 1024). Any of these four hyperparameters is either incremented or decremented for each modification.

The remainder of the process is similar to the self-programming stage. The model is given a initial source code and queried to generate n modified versions. Each modified version is then trained a subset of the target dataset (such as CIFAR-10 or MNIST) before being evaluated on a disjoint subset of the target dataset. The source code with the highest accuracy is then given to the model as a new source, and the process is repeated over the next five main loop iterations.

The performance of the code generated was evaluated on a handful of classic neural network tasks - MNIST, CIFAR10, and EMNIST (ByClass). Each experiment was repeated multiple times with different model setups.

The accuracy of generated networks trained on MNIST increased from roughly 92% to 98% within one training epoch. Accuracy on CIFAR10 improves from 20% to 70% over four epochs. Finally,

accuracy on EMNIST increased from roughly 70% to 87% over four epochs. An example of a modification made by the agent is provided in Table 2. As shown in Figure 3, later models tend to achieve favorable start, end, and intermediate loss values compared to earlier iterations.

| Input Network |
|--|
| Linear(in=784, out=64, bias=True) |
| ReLU() |
| Linear(in=64, out=64, bias=True) |
| ReLU() |
| Linear(in=64, out=10, bias=True) |
| Output Network |
| Conv2d(1, 128, kernel=(3, 3), stride=(1, 1)) |
| Conv2d(128, 256, kernel=(3, 3), stride=(1, 1)) |
| MaxPool2d(kernel=2, stride=2, padding=0, dilation=1) |
| Linear(in=6400, out=64, bias=True) |
| ReLU() |
| Linear(in=64, out=64, bias=True) |
| ReLU() |
| Linear(in=64, out=62, bias=True) |

Table 2: Example of an input neural network and code modifications made by our model. Our system adds two convolutional layers, a pooling layer, and automatically adjusts the number of inputs to the dense layers accordingly.

5 ANALYSIS

Our self-programming results verify that our model is capable of editing its own source code to extend its capabilities.

Furthermore, on all three tasks for programming other models, our system is able to autonomously design networks that achieve effective performance. Our approach is both flexible by way of its free-form programming capabilities and simple in its search and selection procedures. Other methods to search for optimal networks, such as reinforcement learning, may be unstable (Kumar et al., 2020) when applied to different domains, and therefore may require specialized tuning in order to be effective.

Accuracy on MNIST is around 1.5% below SoTA and EMNIST ByClass is about 10% below SoTA, but performance on CIFAR-10 lags significantly behind the current SoTA (Vahidian et al., 2021). This is likely due to the fact that competitive results on CIFAR-10 typically require specialized model designs, which are unlikely to be achieved since our approach uses a simple class of models with greedy search over random code modifications. Performance on these three problems could possibly be improved to by using more sophisticated source code data and search algorithms specialized for computer vision model development.

As shown in Figure 4a, the average performance of generated models on CIFAR-10 tends to increase with the total number of candidates (the product of the number of code search iterations multiplied by the number of candidates per iteration) up to a certain point. Figure 4b shows that the average performance of generated models increases significantly with the number of candidates tried per epoch, and is not as sensitive to the number of code search iterations.

Figure 5 plots the average model size (number of parameters) of generated models over 10 consecutive reprogramming steps. The most significant increases were found to typically occur in the number of convolutional layers. Other significant increases occurred in the size of the hidden layers. The size of the convolutional layers did not increase significantly in any of the three tasks (it was found to fluctuate for CIFAR-10 and MNIST), and the number of hidden layers was rarely found to be increased.

In the three tasks involving generating other models, the code-generating model consistently chose to increase convolutional layer counts, which was found to produce significant increases in performance. These deeper convolutional network source codes were naturally reused by the algorithm as input in




Figure 4: The first graph (left) illustrates the average performance of a network on CIFAR-10 based on the total number of candidates tested over all epochs. The second (right) shows the average performance of a model with a varying number of candidates over four epochs and the average performance of a model with 8 candidates over a varying number of epochs.

subsequent code generation queries. In the initial training loop, most networks had no convolutional layers and a small minority had a single convolutional layer. Within four loops, every network had at least 2 convolutional layers and often more. The number of hidden layers did not significantly increase, indicating that deeper feed-forward architectures may not be necessary to achieve increased performance on these problems.

As can be seen in 6, the number of trainable parameters in the best-performing network increased significantly over time. This agrees with our current understanding of neural networks that larger models with more parameters are more expressive and capable of effectively addressing more complex tasks (Hestness et al., 2017; Kaplan et al., 2020).

All in all, our results demonstrate that our system is able to effectively rewrite the code for a variety of types of neural networks. From a self-programming perspective, our self-programming AI system is capable of extending itself with additional functionalities by way of programming sub-models that can perform auxiliary tasks. From a more practical perspective, our neural network code generation approach can be seen as a kind of free-form AutoML solution that can be used to support model development in various task domains.

6 CONCLUSION

We propose and experimentally validate the first functional self-reprogramming AI system. In our experiments, we successfully implement the system and show it can freely modify its own neural network design. Over several iterations of the algorithm, reprogrammed models tend to achieve increasingly favorable performance on a simple supervised text-to-text code modification task.

Furthermore, the code-generating model used in our self-programming system can easily be adapted to program other models, in order to perform various other machine learning tasks.

6.1 FUTURE WORK

Our work presents a number of promising directions for future research. We are interested in extending this work by incorporating training on a large computer code corpus. Not only would this allow our model to adapt to differences in code styles (such as syntax differences or distinctions between machine learning libraries and frameworks), but it would also allow the model to incorporate code and synthesize new approaches to improve itself and its auxiliary sub-models.

Furthermore, an actor-critic implementation in which the code generation model is trained as an actor via gradients from an associated critic model, similarly to CodeRL in (Le et al., 2022), could improve performance and generalization. We plan to implement these ideas promptly in future work.

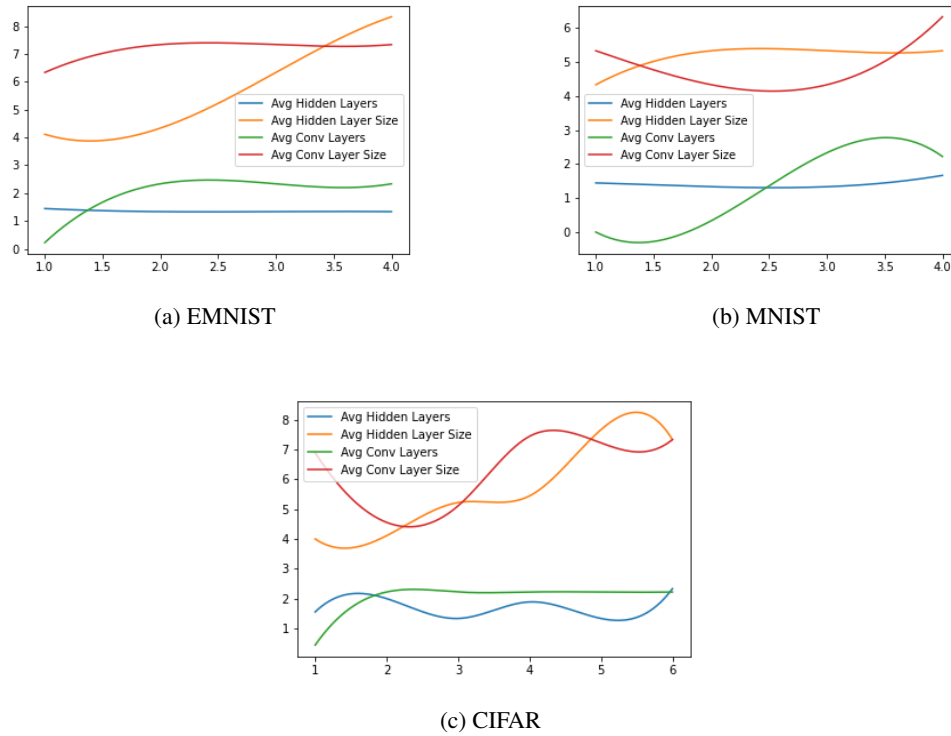


Figure 5: Average values of model design parameters per epoch for convolutional networks on EMNIST, MNIST, and CIFAR. The "avg hidden layer size" and "avg conv layer size" are represented as \log_2 of actual values. The results for MNIST and EMNIST are shown for the first four epochs, because neither network performance nor model design parameters changed significantly in subsequent epochs.

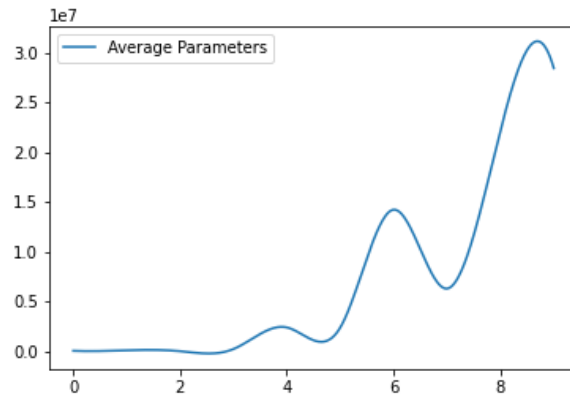


Figure 6: Average number of parameters in best-performing model over 10 epochs. Analysis was conducted on CIFAR-10. Networks in the other two tasks converged in four or fewer epochs.

REFERENCES

- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Sandhini Agarwal, Amanda Askell, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *arXiv*, 2020. URL <https://arxiv.org/abs/2005.14165>.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harrison Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code. *CoRR*, abs/2107.03374, 2021. URL <https://arxiv.org/abs/2107.03374>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018. URL <http://arxiv.org/abs/1810.04805>.
- Iddo Drori, , Sarah Zhang, Reece Shuttleworth, Leonard Tang, Albert Lu, Elizabeth Ke, Kevin Liu, Linda Chen, Sunny Tran, Newman Cheng, Roman Wang, Nikhil Singh, Taylor L. Patti, Jayson Lynch, Avi Shporer, Nakul Verma, Eugene Wu, and Gilbert Strang. A neural network solves, explains, and generates university math problems by program synthesis and few-shot learning at human level. *PNAS*, 119, 2022. URL <https://www.pnas.org/doi/10.1073/pnas.2123433119>.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *CoRR*, abs/1703.03400, 2017. URL <http://arxiv.org/abs/1703.03400>.
- Ali Haidar, Matthew Field, Jonathan Sykes, Martin Carolan, and Lois Hollaway. Pspso: A package for parameters selection using particle swarm optimization. *SoftwareX*, 15, 2021. URL <https://www.sciencedirect.com/science/article/pii/S2352711021000510>.
- Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md. Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. Deep learning scaling is predictable, empirically. *arXiv*, 2017. URL <https://arxiv.org/abs/1712.00409>.
- Timothy Hospedales, Antreas Antoniou, Paul Micaelli, , and Amos Storkey. Meta-learning in neural networks: A survey. *IEEE Transactions On Pattern Analysis And Machine Intelligence*, 44, 2022. URL <https://www.computer.org/csdl/journal/tp/2022/09/09428530/1twaJR3AcJW>.
- Yesmina Jaafra, Jean Luc Laurent, Aline Deruyver, and Mohamed Saber Naceur. Reinforcement learning for neural architecture search: A review. *Image and Vision Computing*, 89, 2019. URL https://www.sciencedirect.com/science/article/pii/S0262885619300885?casa_token=JH1xifIR77UAAAAA:1TCmzZmNzJyLAtMWHaCcLehNgxmyD9YjhANEM9SZcIdwGYcCmxVPgFXthXOr0RwycTDvPwj8BA.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Rewon Child Benjamin Chess, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv*, 2020. URL <https://arxiv.org/abs/2001.08361>.
- Mikhail Khodak, Maria-Florina Balcan, and Ameet Talwalkar. Adaptive gradient-based meta-learning methods. *NeurIPS*, 2019. URL <https://arxiv.org/abs/1906.02717>.
- Aviral Kumar, Abhishek Gupta, and Sergey Levine. Discor: Corrective feedback in reinforcement learning via distribution correction. *Advances in Neurological Processing Systems*, 2020. URL <https://arxiv.org/pdf/2003.07305.pdf>.
- Hung Le, Yue Wang, Akhilesh Deepak Gotmare, Silvio Savarese, and Steven C.H. Hoi. Coderl: Mastering code generation through pretrained models and deep reinforcement learning. *arXiv*, 2022. URL <https://arxiv.org/pdf/2207.01780>.

-
- Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, Thomas Hubert, Peter Choy, Cyprien de Masson d'Autume, Igor Babuschkin, Xinyun Chen, Po-Sen Huang, Johannes Welbl, Sven Gowal, Alexey Cherepanov, James Molloy, Daniel J. Mankowitz, Esme Sutherland Robson, Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu, and Oriol Vinyals. Competition-level code generation with alphacode, 2022. URL <https://arxiv.org/abs/2203.07814>.
- Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. Meta-sgd: Learning to learn quickly for few shot learning. *CoRR*, abs/1707.09835, 2017. URL <http://arxiv.org/abs/1707.09835>.
- Yuqiao Liu, Yanan Sun, Bing Xue, Mengjie Zhang, Gary G. Yen, and Kay Chey Tan. A survey on evolutionary neural architecture search. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*, 2020. URL <https://arxiv.org/abs/1910.10683>.
- Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V. Le. Regularized evolution for image classifier architecture search. *CoRR*, abs/1802.01548, 2018. URL <http://arxiv.org/abs/1802.01548>.
- Esteban Real, Chen Liang, David R. So, and Quoc V. Le. Automl-zero: Evolving machine learning algorithms from scratch. *arXiv*, 2020. URL <https://arxiv.org/abs/2003.03384>.
- Adam Santoro, Sergey Bartunov, Matthew M. Botvinick, Daan Wierstra, and Timothy P. Lillicrap. One-shot learning with memory-augmented neural networks. *CoRR*, abs/1605.06065, 2016. URL <http://arxiv.org/abs/1605.06065>.
- Juergen Schmidhuber. Goedel machines: Self-referential universal problem solvers making provably optimal self-improvements. *Lecture Notes in Computer Science - Adaptive Agents and Multi-Agent Systems II*, 3394, 2006. URL <https://arxiv.org/abs/cs/0309048>.
- Juergen Schmidhuber. Evolutionary principles in self-referential learning. on learning now to learn: The meta-meta-meta...-hook. Diploma thesis, Technische Universitat Munchen, Germany, 14 May 1987. URL <http://www.idsia.ch/~juergen/diploma.html>.
- Xian Shi, Pan Zhou, Wei Chen, and Lei Xie. Efficient gradient-based neural architecture search for end-to-end asr. *ICMI-MLMI*, 2021. URL <https://dl.acm.org/doi/10.1145/3461615.3491109>.
- Kenneth Stanley, David D'Ambrosio, and Jason Gauci. A hypercube-based encoding for evolving large-scale neural networks. *Artificial life*, 15:185–212, 02 2009. doi: 10.1162/artl.2009.15.2.15202.
- Saeed Vahidian, Mahdi Morafah, and Bill Lin. Personalized federated learning by structured and unstructured pruning under data heterogeneity. *IEEE International Conference on Distributed Computing Systems Workshop*, 2021. URL <https://ieeexplore.ieee.org/document/9545941>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *CoRR*, abs/1706.03762, 2017. URL <http://arxiv.org/abs/1706.03762>.
- Yue Wang, Weishi Wang, Shafiq R. Joty, and Steven C. H. Hoi. Codet5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation. *CoRR*, abs/2109.00859, 2021. URL <https://arxiv.org/abs/2109.00859>.
- Darrell Whitley. A genetic algorithm tutorial. *Statistics and Computing*, 15, 2021. URL <https://link.springer.com/article/10.1007/bf00175354>.
- Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. *CoRR*, abs/1611.01578, 2016. URL <http://arxiv.org/abs/1611.01578>.