Feedback-guided Dataset Shaping for Automated Downstream Task Optimization

Lukas Nolte[1 [0009−0008−6846−5302]], Marten J. Finck[2 [0009−0008−9808−2222]], Sören Pirk[2 [0000−0003−1937−9797]], and Sven Tomforde[1 [0000−0002−5825−8915]]

[1] Intelligent Systems, Kiel University, Hermann-Rodewald-Straße 3, 24118 Kiel, Germany

[2] Visual Computing and Artificial Intelligence, Kiel University, Neufeldtstraße 6, 24118 Kiel, Germany

Abstract. Self-learning agents in open-world environments often face data scarcity due to the diversity and unpredictability of real-world conditions. Traditional datasets often over-represent common cases while under-representing rare but critical events, leading to biased learning and poor generalization. To address this, we propose dataset shaping. Dataset shaping aims to use generative models, such as multi-task diffusion models (MTDMs), to generate and refine synthetic data-label pairs through a feedback loop. By dynamically adjusting the data composition according to the performance of a self-learning agent-based downstream task model (DSTM), generative models expose agents to diverse and challenging scenarios, which leads to increased robustness and adaptability. Consequently, dataset shaping enhances generalization, particularly in applications, such as autonomous driving and robotics, where reliable performance in novel conditions is essential.

Keywords: dataset shaping · generative AI · agentic AI · computer vision · multi-task diffusion models · active learning.

1 Introduction

In open-world scenarios, self-learning agents (e.g. as core of Organic Computing systems [7]) often struggle from data scarcity due to the vast diversity and unpredictability of real-world environments. Unlike controlled settings where curated datasets provide the required coverage of relevant situations, open-world agents must generalize across a vast range of conditions, many of which are rare or difficult to capture. For example, in autonomous driving, training data can include common urban scenes, but they lack sufficient examples of extreme weather conditions, unusual lighting, or rare road events, leading to performance gaps in real-world deployments. Data collection in open-world settings is often constrained by ethical, logistical, and financial challenges, making it infeasible to manually gather and annotate the necessary volume of training data. Moreover, traditional datasets often do not include data points of the long-tail of a distribution – usually common examples dominate, while critical yet infrequent

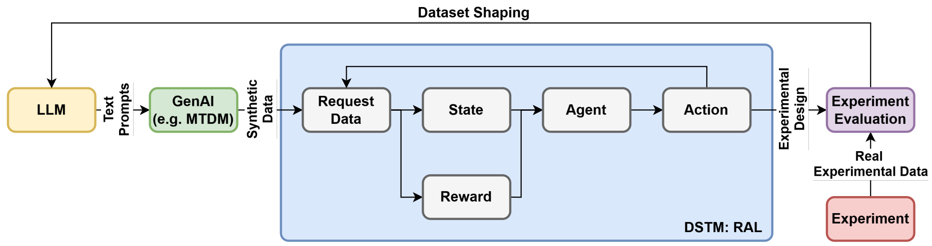

Fig. 1. Overview of the proposed agentic AI-based feedback-guided dataset shaping framework. Text-prompts are generated to produce labeled synthetic datasets with a generative model (GenAI). A reinforcement active learning-based downstream task model (DSTM: RAL) is trained on this generated synthetic data and evaluated on a real dataset. Based on the agent's decision, new prompts are generated, dynamically reshaping the synthetic data distribution to refine the DSTM training process. This iterative dataset shaping enables a targeted exploration, allowing the agent to systematically adjust and optimize data characteristics for improved downstream performance.

edge cases remain underrepresented. Without sufficient exposure to these rare but crucial data points, self-learning agents are prone to biased learning, poor generalization, and failure in novel or unexpected situations. Overcoming this data shortage is therefore essential to ensure that agents can operate reliably and safely in open-world environments.

We aim to address the data bottleneck by introducing dataset shaping. Dataset shaping aims to leverage generative models that generate synthetic data-label pairs, which are continuously refined through a feedback loop (Figure 1). Thereby, the synthetic data is used to train a self-learning agent-based downstream task model (DSTM). After every epoch, the agent's action decisions are evaluated on a target dataset. Data points for which the agent does not perform well are used to request new data by generating text prompts to shape the generation of synthetic data. Thus, a targeted exploration of the agent is established. This process ensures that self-learning agents are exposed to a progressively challenging and diverse set of training samples, addressing common issues like class imbalance and the long-tail distribution problem. Hence, through dataset shaping, the training curriculum of self-learning agents evolves dynamically, allowing them to generalize better to novel situations and ultimately achieve higher performance with minimal reliance on manually curated datasets.

2 State of the Art

Self-learning agents face various challenges related to the scarcity of data in open-world scenarios. Analogous, computer vision tasks, such as classification, segmentation, and depth prediction, often require large-scale, carefully annotated datasets. While synthetic data generated through computer graphics techniques has been used to mitigate this challenge, achieving photorealism and accurately modeling physical properties remains computationally demanding [1]. Recent advances in generative neural networks offer an alternative approach, enabling the synthesis of new images [11] through models such as BigGAN [2]. However, these methods suffer from limitations such structural inconsistencies [3, 5] or lack of pixel-precise alignment between images and their corresponding labels [13]. Only a few methods directly generate multiple data modalities, with DatasetGAN [12] and BigDatasetGAN [6] being notable examples that facilitate multi-task generative modeling for dataset creation. Finally, diffusion models [10] have recently emerged as the leading approach, excelling in image generation tasks [8]. Another approach to addressing data scarcity is active learning, which dynamically adapts models by implementing a feedback loop and selecting the most informative samples while minimizing the amount of labeled data required [9].

3 Research Agenda

Our working hypothesis is that generative AI models, such as Stable Diffusion, produce smoother and more uniform data distributions for solving the problem of inherently imbalanced long-tail distributions present in real-world training datasets, compared to existing methods targeting this challenge.

To address this problem, our aim is to employ novel AI techniques with targeted balancing strategies, such as LoRA adaptors or active learning-inspired approaches such as ActGen [4]. State-of-the-art generative models can synthesize data conditioned on textual prompts and their content. The goal of this proposal is to leverage the text-based synthetic data generation for dataset shaping and apply it as a tool for targeted exploration of a self-learning agent in a DSTM. Once trained, the generative model is integrated into the DSTM training loop. During training of the DSTM, the generative model generates batches of data based on text prompts that are changed during exploration and training of the DSTM agent. With the objective to automatically shape datasets for training self-learning agents the key goals of this project are: (1) to develop generative models, e.g. multi-task diffusion models (MTDMs), that allow synthesizing real-world-like data, along with their labels; (2) to generate a diverse set of text prompts by a pretrained large language model (LLM) for shaping a dataset with generative models; (3) to employ generative models as data generator for training DSTMs with a focus on creating diverse and balanced data distributions; (4) to validate the DSTM agent's decision based on a real datasets to assess its performance; (5) to define a feedback loop that generates new text prompts based on the training outcome of the DTSM agent – the training data distribution is automatically refined for the next stage of training.

With dataset shaping we aim to define a training curriculum for self-learning agent-based DSTMs that leads to optimal performance and robustness for real world applications. For example, when training a depth prediction model for a self-driving car, the generative model, such as a MTDM, may initially be configured to predict depth for standard urban scenes ("An urban scene with several cars on the street"). When training the DSTM, text prompts are used to generate more complex data with the MTDM ("A rainy urban scene with cars on the street", "A night-time urban scene with complex lighting", "An intersection with several pedestrians"). The goal is to define a training curriculum that progressively shapes the data in a way that the variance of the data distribution is increased and rare observations (long-tail data distribution) are generated. To enable dataset shaping we plan to address the following open research challenges:

1. We aim to explore the use of existing LLMs for text prompt engineering to automatically produce varied distributions of data. Whether existing LLMs can be controlled in a way to generate distributions of data with specific features is an open research question.
2. To use generative models as data generators for agent-based DSTMs we need to extend existing approaches to not only predict synthetic data, but also the required label for a downstream task, such as the training of a self-learning agent. Some computer vision approaches show initial success in multi-task image generation, but reliably predicting different types of labels (e.g., depth, surface normal, semantic segmentation) with the quality required for downstream training remains a challenging problem.
3. To automatically shape datasets, we aim to define a feedback loop that guides the generation of text prompts according to the performance of an agent-based DSTM. To define this loop, we need to identify how the validation loss of a DSTM can be used to adjust text prompts in a meaningful manner.
4. As generative models produce new data-label tuples, for which no ground truth exist, validating the quality of the synthesized data and labels is challenging. Our goal is to identify new ways to assess the quality of the generated data and labels.

More generally, the goal of this project is to explore generating synthetic data for training neural networks. The use of synthetic data has recently gained momentum as automatically generated images are a cost-efficient alternative to manually annotated datasets. Most commonly computer graphics algorithms are used to model and render scenes to generate photorealistic imagery. Once defined, the advantage of this process is that the required labels can be generated at low computational costs. Generative neural networks (e.g., GANs or diffusion models) on the other hand can be leveraged to generate photorealistic images of objects and scenes, however, only a few methods exist to generate images jointly with the required labels. By defining generative models, such as MTDMs, we aim to contribute to the landscape of methods for automatically generating synthetic data for training neural networks and particularly self-learning agents for different tasks, such as computer vision.

4 Conclusion

As outlined, our goal is to address the challenge of data scarcity in open-world scenarios and propose dataset shaping – the adaptive generation of data – as a solution to improve the training of self-learning agents. By leveraging generative models, such as MTDMs, we want to investigate generating diverse and visually complex synthetic datasets that dynamically adapt based on the performance of a given downstream task model. Our approach focuses on defining a feedback loop to progressively refine training data distributions that allow us to enhance generalization and robustness in real-world applications such as autonomous driving and robotics. By addressing key research questions related to prompt engineering, multi-modal image generation, dataset validation, and feedback-driven optimization, this research will contribute to fields of generative AI, active learning, self-learning agents and synthetic data generation. With automated dataset generation methods, we enable self-learning agents to better navigate in complex environments with minimal manual intervention.

Acknowledgements

This work was supported by the IDIR-Project (Digital Implant Research), a cooperation financed by Kiel University, University Hospital Schleswig-Holstein and Helmholtz Zentrum Hereon.

References

1. Borkman, S., Crespi, A., Dhakad, S., Ganguly, S., Hogins, J., Jhang, Y., Kamalzadeh, M., Li, B., Leal, S., Parisi, P., Romero, C., Smith, W., Thaman, A., Warren, S., and Yadav, N. Unity perception: Generate synthetic data for computer vision. *CoRR abs/2107.04259* (2021).
2. Brock, A., Donahue, J., and Simonyan, K. Large scale GAN training for high fidelity natural image synthesis. In *ICLR* (2019).
3. Chen, Q., and Koltun, V. Photographic image synthesis with cascaded refinement networks. In *ICCV* (10 2017), pp. 1520–1529.
4. Huang, T., Liu, J., You, S., and Xu, C. Active generation for image classification. In *ECCV* (2024), p. 270–286.
5. Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. Image-to-image translation with conditional adversarial networks. pp. 5967–5976.
6. Li, D., Ling, H., Kim, S. W., Kreis, K., Barriuso, A., Fidler, S., and Torralba, A. Bigdatasetgan: Synthesizing imagenet with pixel-wise annotations, 2022.
7. Müller-Schloer, C., and Tomforde, S. *Organic Computing - Technical Systems for Survival in the Real World*. Birkhäuser, 2017.
8. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models. *CVPR* (2021).
9. Settles, B. Active learning literature survey.
10. Song, Y., and Ermon, S. Generative modeling by estimating gradients of the data distribution, 07 2019.
11. Wu, X., Xu, K., and Hall, P. A survey of image synthesis and editing with generative adversarial networks. *Tsinghua Science and Technology 22*, 6 (2017).
12. Zhang, Y., Ling, H., Gao, J., Yin, K., Lafleche, J.-F., Barriuso, A., Torralba, A., and Fidler, S. Datasetgan: Efficient labeled data factory with minimal human effort. *CVPR* (2021), 10140–10150.
13. Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. *ICCV* (2017), 2242–2251.