Increasing the Diversity of Deep Generative Models

Sebastian Berns

Queen Mary University of London, UK
s.berns@qmul.ac.uk

Abstract

Generative models are used in a variety of applications that require diverse output. Yet, models are primarily optimised for sample fidelity and mode coverage. My work aims to increase the output diversity of generative models for multi-solution tasks. Previously, we analysed the use of generative models in artistic settings and how its objective diverges from distribution fitting. For specific use cases, we quantified the limitations of generative models. Future work will focus on adapting generative modelling for downstream tasks that require a diverse set of high-quality artefacts.

Introduction

While generative models (GMs) were originally designed for representation learning and distribution fitting, a by-product of such models is the ability to produce artefacts of high fidelity. Particularly in the image domain, their generative and interpolation capabilities have been demonstrated. Artists have embraced the peculiar aesthetics of GMs, in particular generative adversarial networks (GANs), and co-opted them to fuel a re-ignition of digital art.

GMs are used in many other applications, e.g. for the generation of video game levels (Volz et al. 2018) and other assets. In robotics, GMs provide a compressed search space over behaviour repertoires which allows for the optimisation of sequences of movement (Cully and Demiris 2018). For the synthesis of biologically active small molecules, it is common to generate large libraries which are screened for promising candidates. Some GMs have been shown to yield a nanobody library that is 1000-fold smaller compared to conventional synthetic libraries, while providing an expression level almost two times as high (Shin et al. 2021). Architects and industrial designers benefit from GMs as creativity support tools to survey the possibility space of a design problem (Bradner, Iorio, and Davis 2014). Generated artefacts serve as a starting point for further design iterations.

As multi-solution tasks, these applications require not only one optimised solution, but several candidate artefacts from the full range of possibilities. What use would it have if we could generate the perfect video game level, but only with slight variations? The output of a GM in this context should therefore be as *diverse* as possible.

My work is motivated by the commonalities in the generative workflows of high-impact applications. The sole objective of GMs for distribution fitting does not take into account the requirements of some downstream tasks, like multi-solution optimisation. I argue for algorithmic adjustments to GMs towards a higher diversity of their output.

Active Divergence

GMs find application in a variety of downstream tasks, whose purpose may differ from the settings in which the modelling techniques are developed. For example, metrics of sample fidelity and mode coverage, in particular Fréchet inception distance (FID), have become the de facto standard performance measure, against which new contributions are scrutinised. Yet, I found that the purpose for GMs in an artistic setting is different from standard applications. Artist strive to tweak, hack or consciously break a GM to produce artefacts of high cultural value. While building on the feature learning capabilities of a model, a common goal is to *actively diverge* from the data distribution (Berns and Colton 2020). Such diversion can be achieved through a variety of techniques (Broad et al. 2021). Artists' goals are more complex than just accurately modelling the training data.

Automating Generative DL for Art

Creators use GMs in highly customised pipelines for art production. These workflows consist of many manual tasks, like the curation of data sets and definition of loss function, that influence the outcome of the generated artefacts. In our work, we propose a framework which integrates core concepts from automated machine learning (AutoML) with a central goal from computational creativity (CC) (Berns et al. 2021). AutoML provides techniques to facilitate the deployment of ML models through automation of optimisation tasks, such as hyper-parameter tuning. When applied to a system for art production, such automation, from the perspective of CC, can be understood as endowing a system with creative responsibilities. Combining these two concepts, thus, provides the opportunity for more creatively autonomous systems through the automation of key tasks. Our framework highlights targets for automation in artistic pipelines with the goal of increasing the creative autonomy of computational agents. We further analyse the affordances of the resulting human–computer co-creative systems.

Expressivity of Generative Models

While GMs are capable of providing a valuable lower-dimensional representation space, they are limited in their ability to represent and generate examples beyond the training data. In our study (Hagg et al. 2021), we consider a simple multi-solution 2D shape optimisation task, which we define manually to fully control its parameters. We compare the performance, as measured by pure diversity (PD), of a quality diversity (QD) algorithm (Pugh, Soros, and Stanley 2016) in two search spaces: 1) the manually-encoded parameter space, and 2) the learned latent space of a variational auto-encoder (VAE). In all tested configurations, latent space search yields a solution set of significantly lower diversity than QD search in parameter space. We conclude, that a model's expressivity, i.e. the coverage of a domain's range of parameter values, is bound by its training data.

Data vs Algorithm

Considering the statistical nature of GMs and their training process, we make two observations. On the one hand, a model recognises and models patterns which appear often in many different training examples. This is a useful process which finds commonalities across examples and builds shared features. On the other hand, the frequency of a pattern will determine its likelihood and can consequently make under-represented features improbable under the learned model. The most prominent features in the data will thus be assigned a high probability, while less frequent features will be deemed less likely and might even not be covered at all. It is necessary, therefore, to account for this by adjusting the weighting of samples and even out the probability under the model of all features. These statistical observations and the previously discussed evidence about the limitations in expressivity of GMs, justify a re-adjustment for the purpose of downstream applications that require diverse output.

There is an ongoing debate about the respective influence of algorithms and data on a resulting model, in particular in the context of fairness. Two opposing views place the importance on either the data or the algorithm. On one side, it is argued that a model is only as good as the data. The responsibility for its quality is placed onto data selection and preparation. Shortcomings of a model need to be addressed by improving the data set. The counter-argument states that data is inherently noisy and plagued with biases. Those biases that are deemed important could be mitigated by adaptations to the algorithms. The case for this is especially strong in domains where possibility for data collection is limited or particularly expensive. Both approaches pose potential avenues for improvements to be considered in future work.

Ongoing and Future Work

The core contribution of my thesis will consist in algorithmic proposals to increase the output diversity of GMs. My work is motivated by the usage of GMs in a variety of domains that require a diverse set of output, ranging from the generation of video game assets to small molecules. This is further supported by our analysis of the potentials and limitations of GMs w.r.t. diversity. Some results have already been presented and may be extended. An additional study on the correlation between measures for sample fidelity and measures for diversity is under way (to be finished by the end of September 2021). The core work on methods that increase the output diversity of GMs is scheduled for October 2021 to March 2022. This work is limited by the trade-off between sample fidelity and diversity. The main challenge will be to increase diversity while not reducing artefact quality too much. I seek to demonstrate the benefits of the developed methods in some of the aforementioned specialised applications. This will be achieved through two collaborations with domain experts (April–June and July–September 2022). The final year (October 2022 to September 2023) will be dedicated to thesis write-up.

Acknowledgements

Thanks to the reviewers for their helpful comments. Funding by the EPSRC Centre for Doctoral Training in Intelligent Games & Games Intelligence (IGGI) [EP/S022325/1].

References

Berns, S.; Broad, T.; Guckelsberger, C.; and Colton, S. 2021. Automating Generative Deep Learning for Artistic Purposes: Challenges and Opportunities. In *Proc. ICCC*. (I developed the framework, colleagues helped with write-up.).

Berns, S.; and Colton, S. 2020. Bridging Generative Deep Learning and Computational Creativity. In *Proc. ICCC*. (Own work, my advisor assisted with corrections.).

Bradner, E.; Iorio, F.; and Davis, M. 2014. Parameters Tell the Design Story: Ideation and Abstraction in Design Optimization. In *Proceedings of the Symposium on Simulation for Architecture & Urban Design*.

Broad, T.; Berns, S.; Colton, S.; and Grierson, M. 2021. Active Divergence with Generative Deep Learning - A Survey and Taxonomy. In *Proceedings of ICCC*. (Survey by my colleague, I assisted with the write-up.).

Cully, A.; and Demiris, Y. 2018. Hierarchical Behavioral Repertoires with Unsupervised Descriptors. In *Proceedings of GECCO*.

Hagg, A.; Berns, S.; Asteroth, A.; Colton, S.; and Bäck, T. 2021. Expressivity of Parameterized and Data-driven Representations in Quality Diversity Search. In *Proc. GECCO*. Nominated for best paper. (I substantially contributed to motivation and focus, formulation of recommendations derived from our findings, and the article's structure and write-up.).

Pugh, J. K.; Soros, L. B.; and Stanley, K. O. 2016. Quality Diversity: A New Frontier for Evolutionary Computation. *Frontiers in Robotics and AI*, 3.

Shin, J.-E.; Riesselman, A. J.; Kollasch, A. W.; McMahon, C.; Simon, E.; Sander, C.; Manglik, A.; Kruse, A. C.; and Marks, D. S. 2021. Protein Design and Variant Prediction Using Autoregressive Generative Models. *Nature Communications*, 12(1).

Volz, V.; Schrum, J.; Liu, J.; Lucas, S. M.; Smith, A.; and Risi, S. 2018. Evolving Mario Levels in the Latent Space of a Deep Convolutional Generative Adversarial Network. In *Proceedings of GECCO*.