Augmenting High-dimensional Nonlinear Optimization with Conditional GANs

Pouya Rezazadeh Kalehbasti
pouyar@stanford.edu
Stanford University
Stanford, CA, USA

Michael D. Lepech
mlepech@stanford.edu
Stanford University
Stanford, CA, USA

Samarpreet Singh Pandher
samar89@stanford.edu
Stanford University
Stanford, CA, USA

ABSTRACT

Many mathematical optimization algorithms fail to sufficiently explore the solution space of high-dimensional nonlinear optimization problems due to the curse of dimensionality. This paper proposes generative models as a complement to optimization algorithms to improve performance in problems with high dimensionality. To demonstrate this method, a conditional generative adversarial network (C-GAN) is used to augment the solutions produced by a genetic algorithm (GA) for a 311-dimensional nonconvex multi-objective mixed-integer nonlinear optimization. The C-GAN, composed of two networks with three fully connected hidden layers, is trained on solutions generated by GA, and then given sets of desired labels (i.e., objective function values), generates complementary solutions corresponding to those labels. Six experiments are conducted to evaluate the capabilities of the proposed method. The generated complementary solutions are compared to the original solutions in terms of optimality and diversity. The generative model generates solutions with objective functions up to 79% better, and with hypervolumes up to 58% higher, than the original solutions. These findings show that a C-GAN with even a simple training approach and architecture can, with a much shorter runtime, highly improve the diversity and optimality of solutions found by an optimization algorithm for a high-dimensional nonlinear optimization problem. [GitHub repository: https://github.com/PouyaREZ/GAN_GA]

CCS CONCEPTS

• **Theory of computation → Mathematical optimization**.

KEYWORDS

C-GAN, GA, multiobjective optimization, sustainable urban system

ACM Reference Format:
Pouya Rezazadeh Kalehbasti, Michael D. Lepech, and Samarpreet Singh Pandher. 2021. Augmenting High-dimensional Nonlinear Optimization with Conditional GANs. In *2021 Genetic and Evolutionary Computation Conference Companion (GECCO '21 Companion), July 10–14, 2021, Lille, France*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3449726.3463675

GECCO '21 Companion, July 10–14, 2021, Lille, France

ACM ISBN 978-1-4503-8351-6/21/07…$15.00
https://doi.org/10.1145/3449726.3463675

1 INTRODUCTION

Mathematical optimization faces challenges when solving highly non-linear and high-dimensional optimization problems [4, 6], since the solution space for these problems is so vast that the optimization algorithm fails to properly explore the optimal solutions of the entire space in a reasonable amount of time [4, 6]. To solve this issue, this paper proposes using conditional generative adversarial networks (C-GANs) to learn the underlying distribution of the solutions generated by the optimization algorithm, and then generating unseen, more optimized solutions to the original optimization problem using the generative model (c.f. [3]). C-GAN consists of two adversarial models, a generator and a discriminator [7]. The generator learns the data distribution of the input solutions to generate new solutions, and the discriminator learns to detect if a solution belongs to the input data distribution or not. The adversarial training as well as conditioning on the data labels enable C-GAN to generate unseen solutions for given desired labels [7]. Research to date on applying GANs to optimization has focused on random generation rather than targeted generation of data points with desired labels [5, 6, 9].

The proposed method is tested on a nonconvex multi-objective mixed-integer nonlinear program (MINLP), which is solved using a genetic algorithm (GA). This optimization problem concerns the sustainable design of the buildings, energy plant, and energy distribution network in an urban district [1, 2] by minimizing the life-cycle cost (LCC) and greenhouse gas emissions (GHG) and maximizing the walkability (WLK) [8] of the district. A C-GAN generates complementary solutions for the optimization problem based on the solutions found by the GA. This paper handles the problem of training the C-GAN on the solutions of the optimization problem as one of multi-variate multiple regression, where the features (independent variables) of the training set are the 10 main integer inputs of the optimization problem, and the labels (dependent variables) are the 3 real-valued objective functions (OFs). The contributions of this work include, (i) a new method for augmenting traditional optimization for highly complex optimization problems, (ii) the first application in the literature of C-GANs to multivariate multiple regression, and (iii) the first direct application in the literature of C-GANs to mathematical optimization.

2 METHOD

This paper uses a C-GAN (Figure 1) to complement the performance of a genetic algorithm on a high-dimensional nonlinear optimization. The C-GAN trains on the results from the GA, then generates more diverse and more optimized solutions to the optimization problem than the GA has identified. Those generated solutions that satisfy the constraints of the original optimization problem

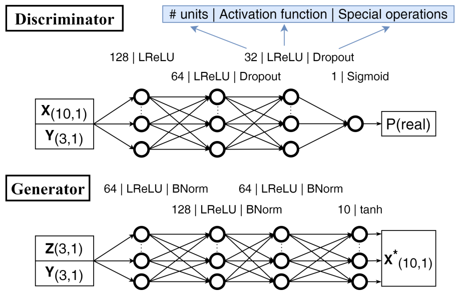

Figure 1: Architectures of C-GAN's generator and discriminator networks. [BNorm: Batch Norm., LReLU: Leaky ReLU]

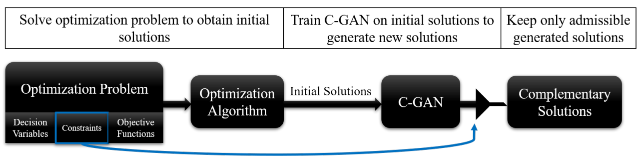

Figure 2: Overview of the proposed method

(i.e. admissible solutions) are kept as the complementary solutions. Figure 2 shows the proposed method.

Different subsets of 65,610 solutions to the optimization problem, from a single run of GA, are used to perform six experiments to inspect the abilities of the proposed method (Table 1). In these experiments, the C-GAN is trained on different subsets of the original solutions which simulate situations where the optimization algorithm has produced solutions with desirable or undesirable objective function (OF) values. The solutions that C-GAN generates are then compared, in terms of different OF values and hypervolume, with their respective training sets.

In the first three experiments, WorstHalfGHG, WorstHalfLCC, and WorstHalfWLK, the C-GAN is trained on the worst half of the initial solutions in terms of each OF. These experiments measure how well C-GAN can generate solutions with, respectively, improved GHGs, LCCs, and WLKs compared to a training set composed of adversely selected solutions in terms of each OF. In the fourth experiment, WorstHalfAll, the same procedure is followed for a training set composed of solutions with all three OFs in the worst half of the OF values in the initial solutions. The generated solutions are then compared in terms of all OFs with the training set. In the fifth experiment, BestHalfAll, the C-GAN is trained on solutions with all three OFs in the best half of the OF values in the initial solutions. This experiment measures if the C-GAN can generate solutions with better OF values than a training set of high-quality solutions found by the GA. In the last experiment, FullData, the C-GAN is trained on the entire 65,610 initial solutions. This experiment measures the performance of C-GAN in creating solutions with better OF values than the entire initial solutions.

3 RESULTS AND CONCLUSIONS

Table 1 shows that the generator has created solutions with a minimum GHG of up to 21% lower, and a minimum LCC of up to 79% lower, than that of the training (input) sets. In the third and fourth experiments, the generator has produced solutions with maximum WLKs of 15.0 (highest possible value of WLK in the studied problem) from training sets with maximum WLKs of 0. However, the generative model has not generated solutions with better OF values than those of the entire original solutions in the FullData experiment. This probably indicates that the original optimization algorithm has discovered solutions with OF values close to their global optima. Nonetheless, in the third and fourth experiments, generative model has made solutions with hypervolumes of 1.0 from training sets with trivial hypervolumes. This shows that compared to the input dataset, the solutions generated by the C-GAN have higher spread and convergence to the optimal Pareto front.

Table 1: Improvements (in %) the C-GAN achieves in the objective functions and hypervolume in the six experiments

Experiment	Min_{GHG}	Min_{LCC}	Max_{WLK}	Hypervol.
WorstHalfGHG	9.0			0.1
WorstHalfLCC		78.7		0.7
WorstHalfWLK			NaN	NaN
WorstHalfAll	21.0	74.4	NaN	NaN
BestHalfAll	8.4	60.9	57.9	57.9
FullData	0.0	-0.8	0.0	0.1

The C-GAN has achieved these improvements in less than 3% of the time needed to run the original optimization method. These results speak to the promise of using generative models, specifically C-GANs, for improving the performance of optimization algorithms, like genetic algorithms, for high-dimensional optimization. This paper also demonstrates that C-GANs, even with simple architectures and small training iterations on low-quality solutions, can significantly improve the results of complex optimization problems.

ACKNOWLEDGMENTS

This material is based on work supported by the Leavell Fellowship on Sustainable Built Environment from Stanford University.

REFERENCES

[1] Robert E Best, Forest Flager, and Michael D Lepech. 2015. Modeling and optimization of building mix and energy supply technology for urban districts. *Applied energy* 159 (2015), 161–177.
[2] Robert E Best, P Rezazadeh Kalehbasti, and Michael D Lepech. 2020. A novel approach to district heating and cooling network design based on life cycle cost optimization. *Energy* 194 (2020), 116837.
[3] Paidamoyo Chapfuwa, Chenyang Tao, Chunyuan Li, Courtney Page, Benjamin Goldstein, Lawrence Carin Duke, and Ricardo Henao. 2018. Adversarial time-to-event modeling. In *International Conference on Machine Learning*. PMLR, 735–744.
[4] Stephen Chen, James Montgomery, and Antonio Bolufé-Röhler. 2015. Measuring the curse of dimensionality and its effects on particle swarm optimization and differential evolution. *Applied Intelligence* 42, 3 (2015), 514–526.
[5] Yi-nan Guo, Jianjiao Ji, Ying Tan, and Shi Cheng. 2020. Multi-objective Combinatorial Generative Adversarial Optimization and Its Application in Crowdsensing. In *International Conference on Swarm Intelligence*. Springer, 423–434.
[6] Cheng He, Shihua Huang, Ran Cheng, Kay Chen Tan, and Yaochu Jin. 2020. Evolutionary multiobjective optimization driven by generative adversarial networks (GANs). *IEEE transactions on cybernetics* (2020).
[7] Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014).
[8] Walk Score. 2011. *Walk score methodology*. Walk Score, Seattle, USA.
[9] Shipu Zhao and Fengqi You. 2020. Distributionally robust chance constrained programming with generative adversarial networks (GANs). *AIChE Journal* 66, 6 (2020), e16963.