

# Tetra-AML: Automatic Machine Learning via Tensor Networks

Machine learning

# Tetra-AML: Automatic Machine Learning via Tensor Networks

A. Naumov, Ar. Melnikov, V. Abronin, F. Oxanichenko, K. Izmailov, M. Pflictsch, A. Melnikov, and M. Perelshtein\*  
*Terra Quantum AG, Kornhausstrasse 25, 9000 St. Gallen, Switzerland*

Neural networks have revolutionized many aspects of society but in the era of huge models with billions of parameters, optimizing and deploying them for commercial applications can require significant computational and financial resources. To address these challenges, we introduce the **Tetra-AML** toolbox, which automates neural architecture search and hyperparameter optimization via a custom-developed black-box **Tensor train Optimization** algorithm, **TetraOpt**. The toolbox also provides model compression through quantization and pruning, augmented by compression using tensor networks. Here, we analyze a unified benchmark for optimizing neural networks in computer vision tasks and show the superior performance of our approach compared to Bayesian optimization on the CIFAR-10 dataset. We also demonstrate the compression of ResNet-18 neural networks, where we use 14.5 times less memory while losing just 3.2% of accuracy. The presented framework is generic, not limited by computer vision problems, supports hardware acceleration (such as with GPUs and TPUs) and can be further extended to quantum hardware and to hybrid quantum machine learning models.

## INTRODUCTION

Over the past decade, neural networks have influenced practically every aspect of human society [1]. When the cutting-edge neural network AlexNet triumphed in the prestigious ImageNet challenge in the 2010s [2], the startling journey began. Today, AI systems like DALLE-2 that can produce realistic visuals and art from a description in natural language are still working toward the same overall goal [3]. While neural networks have grown in value for both people and enterprises, putting them into practice in the context of commercial operations is getting more and more difficult every year.

The process of deploying a model involves several steps, such as model training, architecture and hyperparameter optimization, and testing, which can be computationally intensive and require significant resources. Furthermore, deploying a large-size model can require significant storage and computation resources, which can increase costs [4]. Additionally, if the model needs to be implemented on a small device with limited memory [5], the model may need to be optimized for size, which can add additional complexity and time to the deployment process. At the same time, the model’s high accuracy must be maintained.

In addition, the practice of creating models with greater complexity in order to increase accuracy is still developing. There are about 60 million trainable parameters in AlexNet (2012); GPT-2 contains 1.5 billion parameters (2019); DALLE-2 and ChatGPT contain roughly 3.5 billion and 175 billion parameters, respectively, (2022), with the promise to increase the size by at least two orders of magnitude in the coming years. It is clear that the powerful and complicated models are getting more and more expensive in terms of optimization and deployment [6, 7].

To address these challenges, we develop an Automatic Machine Learning toolbox based on tensor networks, **Tetra-AML**. A tensor network is a powerful numerical tool that can advance the solution of the high-dimensional problems [8]. Here we develop the framework that allows us to apply tensor networks for automatic machine learning [9], including an automatic search for the best architectures of models – neural architecture search – with optimal parameters – hyperparameters optimization – with the help of our black-box optimization approach, **TetraOpt** [10]. Besides, it allows compression of the models via a combination of common quantization techniques and pruning approach [11, 12] augmented by compression using tensor networks [13, 14].

The general scheme of **Tetra-AML** is shown in Fig. 1. **Tetra-AML** offers the flexibility of bringing your own model or defining a use case to receive a suitable model. A user provides a dataset and specifies the search space for optimization. After that, the tool initiates parallel training of the models and applies post-training tensor network compression, pruning and quantization to create an optimal, compressed and accelerated model. Once the model is ready, users can download it for deployment.

In this work, we mainly focus on computer vision problems, as one of the most challenging ones, but the framework is generic and can be applied for any machine learning problem. We consider well-known CIFAR dataset and NATS benchmarking for neural architecture search [15]. Besides, **Tetra-AML** allows for hardware acceleration, such as with GPUs (Graphics Processing Units) and TPUs (Tensor Processing Units), and, moreover, can leverage the power of *quantum computers* for better optimization of classical networks and boosting the performance of quantum and hybrid quantum/classical machine learning models, which we discuss in this work.

---

\* mpe@terraquantum.swiss

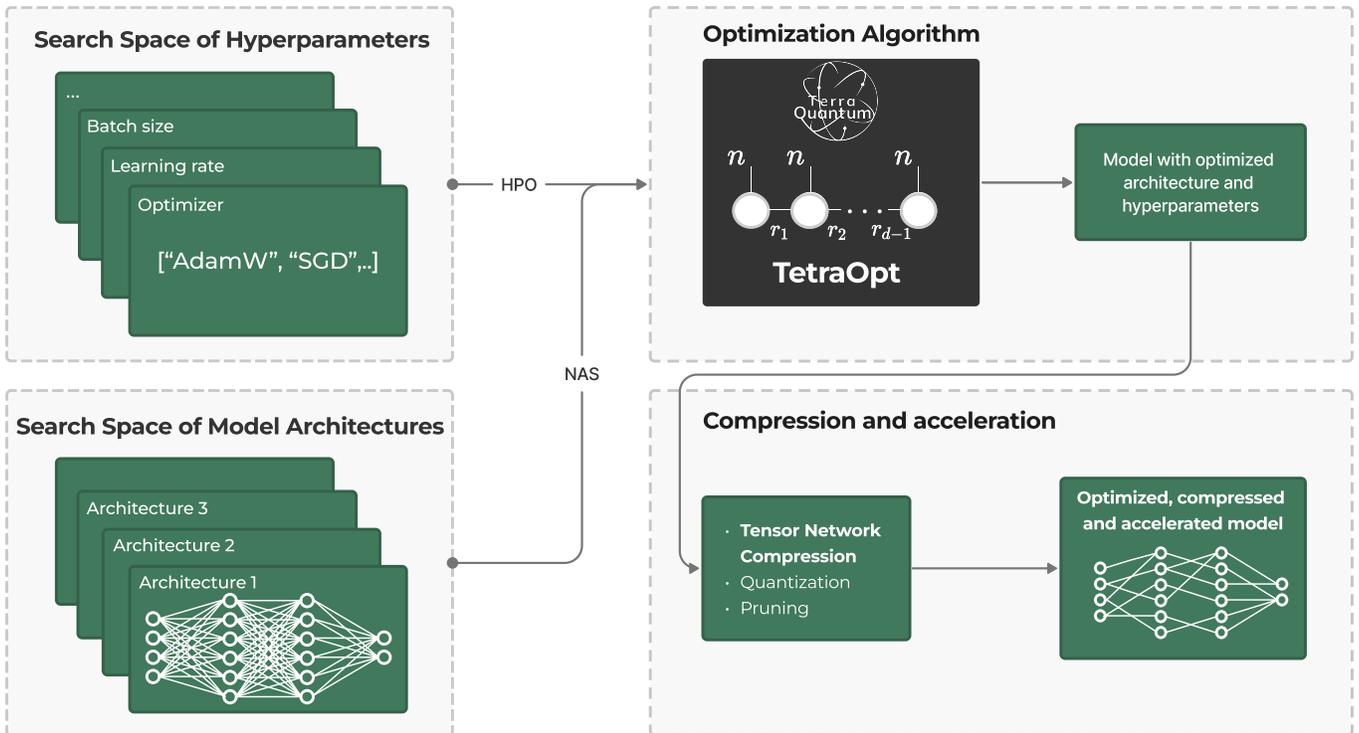


FIG. 1. General scheme of **Tetra-AML**. Both Neural Architecture Search (NAS) and Hyperparameter Optimization (HPO) are performed via **TetraOpt** (Terra Quantum’s black-box optimizer based on Tensor Trains). Then the model is compressed via tensor network methods, quantization and pruning.

## NEURAL ARCHITECTURE SEARCH AND HYPERPARAMETERS OPTIMIZATION

The first step in building a model is choosing the best model architecture for the task at hand. Neural Architecture Search (NAS) is an algorithm, or set of algorithms, that automates the process of finding the best architecture for a particular ML problem and dataset [16]. Such an approach explores the space of potential architectures and assesses their performance during an optimization method. Finding an architecture with the best performance while being computationally efficient is the ultimate aim of NAS. When compared to manually designing and fine-tuning architectures, this can **significantly save the development time and costs** and produce new, better-performing architectures.

In addition to NAS, it is crucial to conduct hyperparameters optimization (HPO) for the model to maximize its accuracy within the specific problem and dataset. Hyperparameters can be considered values that are set for the model and do not change during the training regime, and may include variables such as learning rate, decay rates, choice of optimizer for the model, etc. This tuning can be also done in an automatic manner, e.g., using black-box optimization methods [17].

Unlike other approaches, such as Random Search, Bayesian optimization methods, Reinforcement Learning or Genetic algorithms, our method utilizes a black-

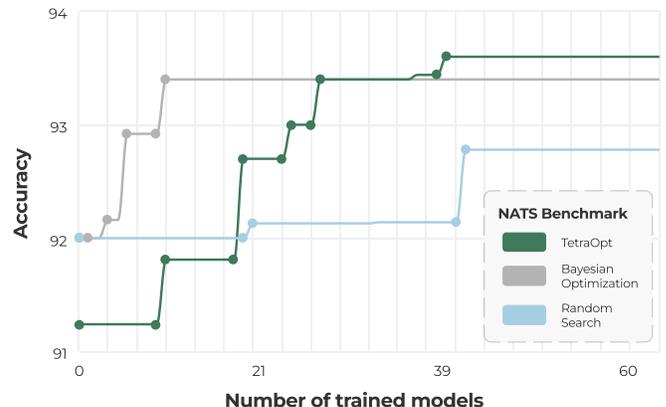


FIG. 2. Validation accuracy dependence on the number of neural network models (architectures) runs for **TetraOpt**, Bayesian, and Random Search algorithms. **TetraOpt** achieves 93.7% accuracy, while Bayesian optimization and Random Search find architectures only with 93.5% and 92.8% accuracy, respectively, with the same number of model runs. The experiments were carried out on a standard NAS benchmark [15], where the CIFAR-10 dataset is used.

box optimization approach based on tensor networks, **TetraOpt**. **TetraOpt** is a versatile global optimization algorithm that can handle various types of variables, including integers and categorical data, and is well-suited

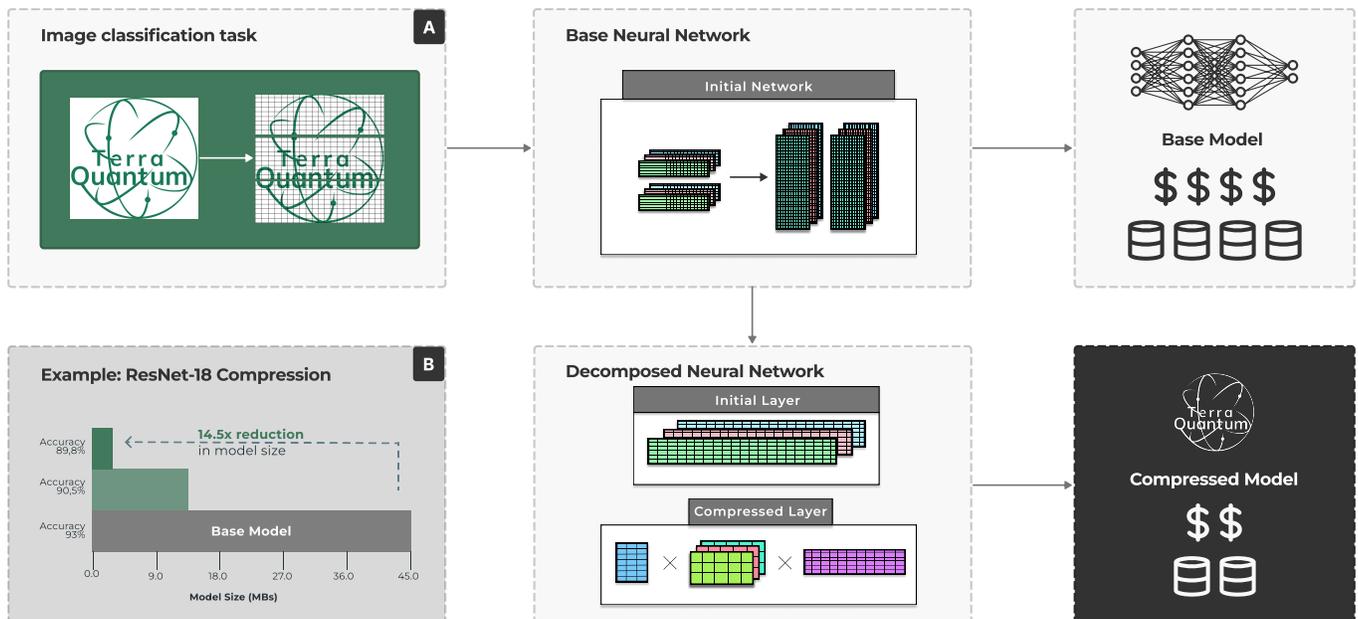


FIG. 3. (A) General scheme of Neural Network compression. (top) General Neural network scheme for image recognition. (bottom) Compressed Convolution layer - 4D convolution kernel is represented as a sum of tensor product of small tensors (Canonical Decomposition). Initial layer has  $C_{in} \times C_{out} \times D^2$  parameters, while after compression remains only  $C_{in} \times R + D^2 \times R^2 + R \times C_{out}$ . For small  $R$ , it provides significant compression in the occupied memory. (B) Compression of state-of-the-art ResNet-18 on the CIFAR-10 dataset via tensor networks. The diagram shows the achieved accuracy depending on the compression of the model. Bottom bar: Uncompressed Base Model. Middle bar: TN Compressed Model (Compression coefficient: 4.5). Top bar: TN Compressed Model (Compression coefficient: 14.5).

for parallel computing and hardware-acceleration via, e.g., GPUs or TPUs. Such an optimization method is also extended to quantum processing units and can advance classical and hybrid quantum-classical machine learning models; see more discussion below. A deep understanding of tensor networks and the ability to develop a custom optimizer advances our solution with a unique advantage in the automatic machine learning market and enables us to deliver faster and more accurate results for innovative business that actively utilizes machine learning models.

In Fig. 2, we present the results of a comparison between **TetraOpt**, Bayesian optimization and Random Search on a standard NAS benchmark, NATS [15]. NATS is a well-known benchmark used to evaluate neural architecture search algorithms – it uses the CIFAR-10 dataset. The search space of the NATS benchmark contains a wide variety of neural network architectures, including different numbers of layers, types of layers, and connections between layers. It includes a search space of 15,625 neural architectures, which makes NATS a comprehensive and universal benchmark that enables fair comparisons between different NAS algorithms. We observe that **TetraOpt** achieves the best validation accuracy for the same number of model runs.

## COMPRESSION OF NEURAL NETWORKS USING TENSOR NETWORKS

Besides complicated tuning, the significant issue with the large models lies in the huge size of a model, which creates additional computational and financial costs that limits the efficient deployment and support. It is even more relevant for devices with limited memory, such as local machines, mobile devices, and Internet of Things in general [5]. In this case, the model compression is a key technology that allows for keeping the accuracy at a decent level while substantially reducing the costs. While many existing approaches mostly use pruning and quantization [18], we focus on the development of new compression algorithms based on tensor networks [13, 19].

Tensor network compression is a cutting-edge technology that potentially offers significant advantages for businesses by compressing deep neural networks while maintaining their accuracy. This approach reduces the number of model parameters and, consequently, the size of the model. The main idea of this method is to decompose a huge matrix (tensor) used in a neural network layer into smaller tensors with an exponential reduction in the memory and a substantial reduction in the runtime [8].

The crucial fact is that this technique can be applied to compress any machine learning model, including feed-forward layers [13], recurrent neural networks [20], convo-

lutional neural networks [19], and potentially, state-of-art transformers [21]. Moreover, by combining the powerful techniques of pruning, quantization and tensor network compression, we gain even more efficiency in resource savings [14].

Here, to illustrate the capabilities of our framework we focus on the image classification problem and ResNet neural networks, which we previously used in Ref. 22. We show the results of the compression of ResNet-18 on the CIFAR-10 dataset in Fig. 3. During the experiment, we compress only CNN layers since they occupy an overwhelming amount of memory in the ResNet-18 architecture. To measure the occupied memory, we estimated the total number of parameters in the compressed and uncompressed models, taking into account fully-connected layers. As one can see, with a higher compression ratio, the accuracy slightly worsens, so the desired accuracy and compression might be defined by a user according to a specific problem.

### PATHWAY TO QUANTUM COMPUTING

We develop the **Tetra-AML** framework keeping in mind the quantum performance enhancement that we can obtain from actively developing quantum hardware. For instance, **TetraOpt** is well extended to a quantum version, which is theoretically capable of providing more optimal points for an optimized objective – a new optimal set of hyperparameters and neural architectures.

On the other hand, when building hybrid quantum-classical models that combine parametrized quantum circuits with classical neural networks, the process of searching for hyperparameter configurations that result in im-

proved model accuracy and training presents a significant challenge. Hyperparameter Optimization and Neural Architecture Search techniques can be applied to searching the optimal parameters in Quantum Machine Learning [22] and optimal balance between classical and quantum contribution in a hybrid model. Such methods can be used for pure quantum models and help to find the best quantum ansatz for a given dataset or a problem.

Overall, NAS provides a promising avenue for developing optimal model architectures for quantum and hybrid quantum computing [23], which could lead to significant advancements in the quantum machine learning field [24–27].

### CONCLUSION

Deep Neural Networks are becoming increasingly prevalent in business as they are able to handle large and complex data sets, leading to improved performance in various tasks, such as image recognition, natural language processing and prediction. However, as networks become larger and more complex, the costs for development and deployment also increase. Here, we propose an off-the-shelf solution that addresses these challenges by using tensor network techniques to find the optimal neural network architecture and hyperparameter optimization for a particular task and dataset and then compress the model to reduce its size while maintaining desired accuracy. We develop a tool for optimizing neural networks that will significantly enhance computational efficiency, reduce the number of parameters, and ultimately improve the overall performance of deep learning models including hybrid quantum neural networks.

- 
- [1] M. I. Jordan and T. M. Mitchell, “Machine learning: Trends, perspectives, and prospects,” *Science*, vol. 349, no. 6245, pp. 255–260, 2015.
  - [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” *Advances in Neural Information Processing Systems*, vol. 25, pp. 1097–1105, 2012.
  - [3] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, “Hierarchical text-conditional image generation with clip latents,” 2022.
  - [4] C. Sun, A. Shrivastava, S. Singh, and A. Gupta, “Revisiting unreasonable effectiveness of data in deep learning era,” 2017.
  - [5] S. Kumar, P. Tiwari, and M. Zymbler, “Internet of things is a revolutionary approach for future technology enhancement: a review,” *Journal of Big Data*, vol. 6, no. 1, p. 111, 2019.
  - [6] J. Sevilla, L. Heim, A. Ho, T. Besiroglu, M. Hobbhahn, and P. Villalobos, “Compute trends across three eras of machine learning,” in *2022 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, IEEE, 2022.
  - [7] D. Patterson, J. Gonzalez, Q. Le, C. Liang, L.-M. Munguia, D. Rothchild, D. So, M. Texier, and J. Dean, “Carbon emissions and large neural network training,” *arXiv preprint arXiv:2104.10350*, 2021.
  - [8] I. V. Oseledets, “Tensor-train decomposition,” *SIAM Journal on Scientific Computing*, vol. 33, no. 5, pp. 2295–2317, 2011.
  - [9] X. He, K. Zhao, and X. Chu, “AutoML: A survey of the state-of-the-art,” *Knowledge-Based Systems*, vol. 212, p. 106622, 2021.
  - [10] D. Morozov, A. Melnikov, V. Shete, and M. Perelshtein, “Protein-protein docking using a tensor train black-box optimization method,” *arXiv preprint arXiv:2302.03410*, 2023.
  - [11] M. Nagel, M. Fournarakis, R. A. Amjad, Y. Bondarenko, M. Van Baalen, and T. Blankevoort, “A white paper on neural network quantization,” *arXiv preprint arXiv:2106.08295*, 2021.
  - [12] S. Anwar, K. Hwang, and W. Sung, “Structured pruning of deep convolutional neural networks,” *J. Emerg. Technol. Comput. Syst.*, vol. 13, no. 3, 2017.

- [13] A. Novikov, D. Podoprikin, A. Osokin, and D. P. Vetrov, "Tensorizing neural networks," *Advances in Neural Information Processing Systems*, vol. 28, 2015.
- [14] D. Lee, D. Wang, Y. Yang, L. Deng, G. Zhao, and G. Li, "QTTNet: Quantized tensor train neural networks for 3D object and video recognition," *Neural Networks*, vol. 141, pp. 420–432, 2021.
- [15] X. Dong, L. Liu, K. Musial, and B. Gabrys, "NATS-bench: Benchmarking NAS algorithms for architecture topology and size," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021.
- [16] G. Kyriakides and K. Margaritis, "An introduction to neural architecture search for convolutional networks," *arXiv preprint arXiv:2005.11074*, 2020.
- [17] S. Falkner, A. Klein, and F. Hutter, "Bohb: Robust and efficient hyperparameter optimization at scale," in *International Conference on Machine Learning*, pp. 1437–1446, PMLR, 2018.
- [18] Y. Cheng, D. Wang, P. Zhou, and T. Zhang, "A survey of model compression and acceleration for deep neural networks," *arXiv preprint arXiv:1710.09282*, 2017.
- [19] V. Lebedev, Y. Ganin, M. Rakhuba, I. Oseledets, and V. Lempitsky, "Speeding-up convolutional neural networks using fine-tuned cp-decomposition," *arXiv preprint arXiv:1412.6553*, 2014.
- [20] A. Tjandra, S. Sakti, and S. Nakamura, "Compressing recurrent neural network with tensor train," in *2017 International Joint Conference on Neural Networks (IJCNN)*, pp. 4451–4458, IEEE, 2017.
- [21] H. Pham Minh, N. Nguyen Xuan, and S. Tran Thai, "TT-ViT: Vision transformer compression using tensor-train decomposition," in *Computational Collective Intelligence: 14th International Conference, ICCCI 2022, Hammamet, Tunisia, September 28–30, 2022, Proceedings*, pp. 755–767, Springer, 2022.
- [22] A. Sagingalieva, A. Kurkin, A. Melnikov, D. Kholmistrov, M. Perelshtein, A. Melnikov, A. Skolik, and D. Von Dollen, "Hyperparameter optimization of hybrid quantum neural networks for car classification," *arXiv preprint arXiv:2205.04878*, 2022.
- [23] M. Perelshtein, A. Sagingalieva, K. Pinto, V. Shete, A. Pakhomchik, A. Melnikov, F. Neukart, G. Gesek, A. Melnikov, and V. Vinokur, "Practical application-specific advantage through hybrid quantum computing," *arXiv preprint arXiv:2205.04858*, 2022.
- [24] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.
- [25] A. Sagingalieva, M. Kordzanganeh, N. Kenbayev, D. Kosichkina, T. Tomashuk, and A. Melnikov, "Hybrid quantum neural network for drug response prediction," *arXiv preprint arXiv:2211.05777*, 2022.
- [26] S. Rainjonneau, I. Tokarev, S. Iudin, S. Rayaprolu, K. Pinto, D. Lemtiuzhnikova, M. Koblan, E. Barashov, M. Kordzanganeh, M. Pflitsch, *et al.*, "Quantum algorithms applied to satellite mission planning for Earth observation," *arXiv preprint arXiv:2302.07181*, 2023.
- [27] A. Melnikov, M. Kordzanganeh, A. Alodjants, and R.-K. Lee, "Quantum machine learning: from physics to software engineering," *Advances in Physics: X*, vol. 8, no. 1, p. 2165452, 2023.