Graph-Theoretical Approaches for AI-Driven Discovery in Quantum Optics

Xuemei Gu
Max Planck Institute for the Science of Light
Staudtstrasse 2, 91058 Erlangen, Germany
xuemei.gu@mpl.mpg.de

Carlos Ruiz-Gonzalez
Max Planck Institute for the Science of Light
Staudtstrasse 2, 91058 Erlangen, Germany

Sören Arlt
Max Planck Institute for the Science of Light
Staudtstrasse 2, 91058 Erlangen, Germany

Tareq Jaouni
Nexus for Quantum Technologies, University of Ottawa
K1N 6N5, ON, Ottawa, Canada

Jan Petermann
Max Planck Institute for the Science of Light
Staudtstrasse 2, 91058 Erlangen, Germany

Sharareh Sayyad
Max Planck Institute for the Science of Light
Staudtstrasse 2, 91058 Erlangen, Germany

Ebrahim Karimi
Nexus for Quantum Technologies, University of Ottawa
K1N 6N5, ON, Ottawa, Canada

Nora Tischler
Centre for Quantum Dynamics, Griffith University
Brisbane, Queensland 4111, Australia

Mario Krenn
Max Planck Institute for the Science of Light
Staudtstrasse 2, 91058 Erlangen, Germany

Abstract

Emerging findings in the physical sciences frequently present new avenues for AI applications that can enhance its efficiency or broaden its scope. We show such a case in the field of quantum optics. Here, we present a method that can represent quantum optics experiments as abstract weighted graphs, converting problems that encompass both continuous and discrete elements into purely continuous optimization tasks. This allows efficient use of both gradient-based and neural network methods, circumventing the need for workarounds due to the discrete nature of the problems. The new representation not only simplifies the design

Machine Learning and the Physical Sciences Workshop, NeurIPS 2023.

process but also facilitates a deeper understanding and interpretation of strategies derived from neural networks.

1 Introduction

Designing experiments in quantum optics poses significant challenges due to its inherent complexity. The vast majority of experiments have historically been designed by experienced human researchers[1, 2]. However, recent years have seen a surge in using advanced computational and machine-learning techniques to assist in both designing and interpreting new quantum optics experiments [3, 4, 5, 6, 7, 8, 9, 10]. The design complexity stems from balancing discrete and continuous elements, involving tasks such as determining the discrete quantum optics network topology and optimizing continuous parameters like laser power or beam splitter reflectivity.

Here, we explore how new theoretical inventions from quantum optics and graph theory [11, 12, 13] can significantly enhance AI methods in quantum optics research. We show a way to transform a problem mixing both continuous and discrete elements into one that's purely continuous. This change not only simplifies the problem but also unlocks the potential for applying AI techniques directly. For instance, gradient-based designs, previously challenging, are now feasible. Furthermore, while neural networks once needed specialized architectures or workarounds like one-hot-encoding, they can now directly leverage the continuous representation. Our findings underscore the value of integrating domain knowledge into AI frameworks, opening a collaborative frontier for physical scientists.

2 Quantum Graph Photonics

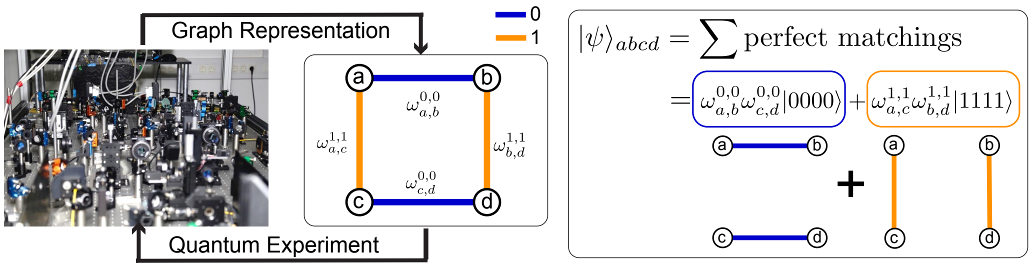

Figure 1: Bridge between quantum experiments and graph theory. In this example, the resulting quantum state is a coherent superposition of the two perfect matchings in the graph. By equalizing and normalizing the weights, we achieve a four-particle Greenberger-Horne-Zeilinger (GHZ) state.

The core insight is that many quantum optics experiments (using nonlinear pair-sources, single photon sources, liner optics, etc.) can be represented in an abstract way as colored, edge-weighted graphs, as illustrated in Fig.1. The vertices of the graph are photon paths to the detectors, and the edges stand for photon correlations with their amplitude given by the edge's weight. The color of the edge represents the mode such as the degree of freedom of a photon (or quantum state) of the correlations. The information of the experiment itself is then stored in the weights of the graphs. Likewise, each graph can be translated back to a concrete quantum experiment that can be implemented in the laboratory, which is very useful for real-world quantum optics experiments [14, 15, 16]. The graph contains information on both the experimental setup and the resulting quantum state, given by the weight function [8]:

$$\Phi(\boldsymbol{\omega}) = \sum_{m} \frac{1}{m!} \left(\sum_{e \in E(G)} \omega(e) x^{\dagger}(e) y^{\dagger}(e) + \text{h.c.} \right)^{m}, \tag{1}$$

where $E(G)$ is the set of edges of the graph. The term h.c. stands for hermitian conjugate, which includes annihilation operators. The quantum state is obtained by applying the weight function to the vacuum, i.e. $|\psi\rangle = \Phi(\boldsymbol{\omega})|vac\rangle$. We show an abstract graph which can produce GHZ states, as an example in Fig. 1. In this case, $\boldsymbol{\omega} = (\omega_{a,b}^{0,0}, \omega_{b,d}^{1,1}, \omega_{c,d}^{0,0}, \omega_{a,c}^{1,1})$ is a list of edge weight $\omega_{x,y}^{i,j} \in \mathbb{C}$ and

$|\omega^{i,j}_{x,y}| < 1$, the superscript and subscript represent the mode number and the optical path, respectively. Conditioned on one photon in each detector (a common technique in quantum optics), the resulting state is a coherent superposition of all the weighted perfect matchings in the graph. The weight of a perfect matching is the product of all its edge weights. A perfect matching is defined as a subset of edges that includes every vertex exactly once. In this scenario, each vertex inherits a color from the colored edge, defining the state of each photon. This representation extends beyond describing quantum states by employing the Choi–Jamiołkowski isomorphism [17, 18], where a graph can also characterize quantum networks and measurement protocols.

The **underlying physical reason** for this efficient representation is the insight that in large fraction of photonic quantum optics, the pure purpose of optical elements is to reshuffle fundamental pair-correlations between photons [19]. These pair-correlations (or higher order correlations [20]) are generated in nonlinear processes, and reshuffled by optical elements. A graph, therefore, depicts these correlations, enabling a direct description of the experiment's final state, which is implementable experimentally [21, 14, 15, 16] and is different from the graph states used for quantum computation [22]. With this representation, one can find new multi-photon quantum interference phenomenon [12] which has been experimentally demonstrated [14, 15] and inspired a very large-scale integrated optics experimental effort [16].

3 AI-Driven Discovery

3.1 Gradient-based discovery

For designing a new quantum optics experiment, we can directly use the graph-representation described above as outlined in Fig. 2 (a). The graph, parameterized by the weights ω_i, leads to the quantum state (or quantum network or quantum measurement protocol). From this, we can build loss functions in terms of the weights of the graph $L(\omega_i)$. For instance, if the goal is to design an experiment for a specific quantum state, the loss function could be the state's fidelity, parameterized by ω_i. Similar fidelity functions can be applied for quantum networks and quantum measurements.

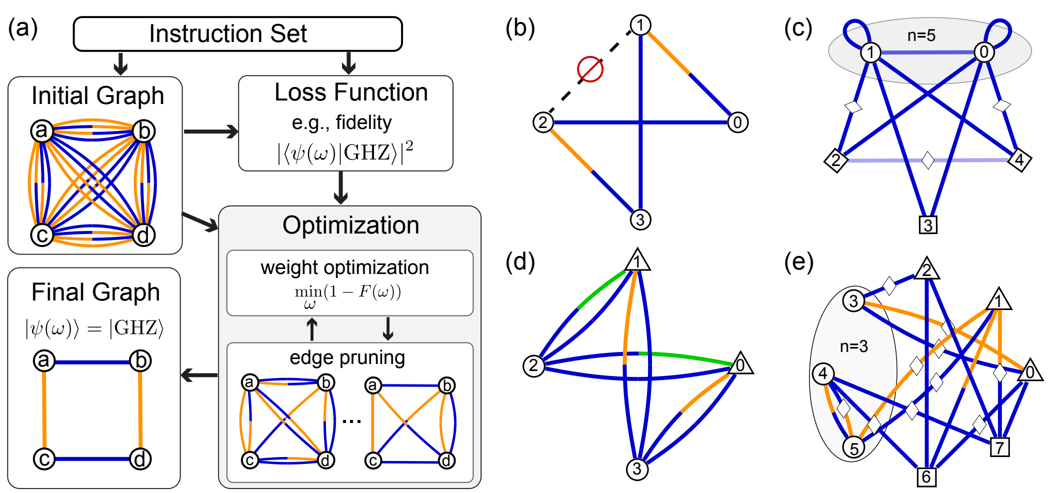

Figure 2: Gradient-based discovery. (a), The design workflow. The Instruction Set file details what we want (e.g., a quantum state) and instructions about the loss function and optimization details. Based on the task in the instruction file, one can impose topological constraints in the initial graph. (b), An entirely new way for performing quantum entanglement swapping (entangling photons that never interacted [23]), a core principle of quantum networks. (c), Efficient blueprint for the generation of 5-photon NOON states, which are at the core of highly sensitive quantum measurements [24, 25]. The square and diamond in the graph represent the ancillas and negative amplitude, respectively. (d), Quantum measurement for a quantum communication task with quantum advantage (e.g., Mean King's Problem [26]). The incoming refers to the triangle in the graph. (e), A heralded Fredkin gate acting on a target input where one qubit is set to mode zero.

The design of the experiment has become a continuous, high-dimensional optimization problem, which can be solved by gradient-based algorithms such as the Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm [27]. In reality, we are not only interested in a solution to the question, but in simple solutions. For that reason, our algorithm iteratively performs continuous optimization and edge pruning (i.e. removing edges), until the graph cannot be simplified any further. In many cases, the resulting solution contains only a small number of edges, which allows for physical interpretation and understanding of the underlying solution. We've successfully applied our algorithm on more than 100 previously unsolved open questions in quantum optics, creating a broad catalog of designs that can benefit areas such as quantum computing, communications, and sensing. Fig.2 (b)-(e) presents four diverse examples of these solutions (found with one minute using a CPU Xeon Gold 6130 with 187 GiB RAM). In tasks with higher dimensions and more particles, precisely estimating the required solution time is difficult due to greater complexity.

3.2 Neural-network-based design and interpretability

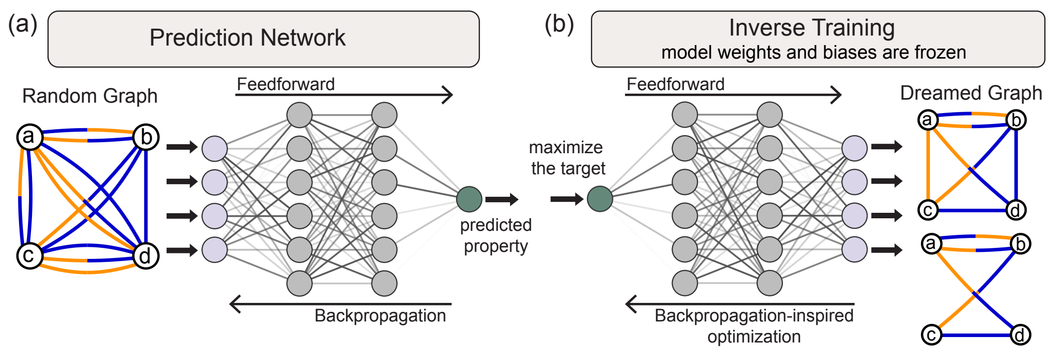

Figure 3: Quantum Graph Deep Dreaming. (a), A feed-forward neural network is trained to predict a property, e.g., fidelity, of a given random input setup represented by a graph. (b), In the deep dreaming phase (inverse training), the network's weights and biases are fixed, and the initial input graph's weights are updated iteratively to maximize the predictive output.

Here, we explore the application of neural-network-based design of quantum optics. We use a technique invented in the field of computer vision called *inception* or *deep dreaming* [28]. The idea is to train a neural network to predict the properties of a given structure (i.e. learns the *structure-property relation*), see Fig.3 (a). After the training, the neural network's weights and biases are frozen, and the training is inverted, see Fig.3(b). Instead of inputting a structure (in our case, a graph), we modify the structure such that the property is maximized. Deep dreaming has been applied in physics for various tasks such as quantum entanglement spectra [29] or the arrow of time [30]. Recently, it has also been used for design purposes in material science [31] and quantum circuits [32].

First, we demonstrate how we can design quantum states, such as the GHZ state, using a neural network that works on the photonic graph representation. The neural network, composed of a input layer of 24 neurons for all possible edge weights, four hidden layers (three with 400 neurons, one with 10), and a single-neuron output layer, is trained on 20 million low-fidelity (below 0.5) examples. After training the neural network to predict properties from the graphs, we then invert the network. As shown in Fig.4(a), the state's fidelity (the similarity between our target and the current value) increases continuously during the training, as the activation of the output neuron increases. In Fig.4(b), we observe how the distribution of 1000 random initial graphs (blue) is shifted towards high fidelity, clearly indicating the ability for design of quantum optics experiments. It took one day to get the trained neural network and results using a GPU (NVIDIA Volta V100 w/ 15 GB of HBM2 memory).

Having a neural network that can adapt arbitrary graphs towards a quantum system with specific properties allows us to extract strategies learned by the machine learning model. Specifically, we can start with particular graphs, and observe their evolution towards the final state. In Fig.4(c), it is evident that the network has learned to generate disjoint perfect matchings of opposite colors, which contribute the quantum state's missing terms. In Fig.4(d), we pose a more challenging task: instead of designing an experiment for a specific quantum state, we aim to find an experiment for a state with a specific property, even when we are uncertain of what the best state might be. In this case, our target

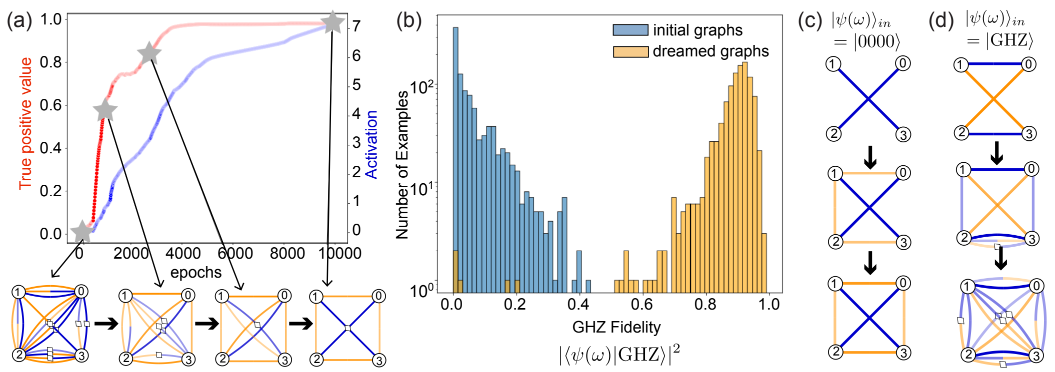

Figure 4: Results. (a), We show the evolution of an input graph's fidelity respect to the GHZ state, during the dreaming on the neural network. Intermediate steps of a random graph's evolution to its dreamed counterpart are displayed. (b), Distribution of initial vs. dreamed fidelity respect to GHZ state. No trained initial graph has fidelity above 0.5. (c)-(d), Strategies from the inverse training are explored, using specified inputs to optimize GHZ state fidelity (c) and the mean value of $\mathrm{tr}(\rho_M^2)$ (d). Disjoint yellow-cycle graphs are generated in (c), and in (d), edge weights are modified for the diagonal perfect matching, and new perfect matchings are created with specific weight contributions.

is the concurrence of the quantum state, a measure of multipartite quantum entanglement [33, 34]. Starting from a GHZ state (known to be highly entangled), we observe that the neural network creates more new terms as the evolution continues, each with specific weight contributions. Surprisingly, the concurrence of the resulting state is found to be significantly larger than that of the GHZ state. The network gradually learns to recognize increasingly complicated structures, which could help us to understand how the experiment works. Although it is a simulation at the moment, we expect to see experiments in the next step as many computer-designed experiments have been done [35, 36, 37].

4 Outlook

We demonstrate how the theoretical understanding of quantum optics and quantum information can enhance the use of AI in this field. Our work has led to the efficient design of more than 100 new quantum optics experiments through a direct, gradient-based approach. Furthermore, we show how neural-network-based optimization is directly applicable without modifying the network architectures. This allows us to observe the internal strategies that the neural network learns to optimize and design specific quantum states. One important follow-up question is to what extent we can expand the physical graph-based representation to include even more experimental technologies, enabling highly efficient, purely gradient-based, and neural-network-based optimization. For example, the frequency of photons [38, 39] can be mapped by discretization or a more continuous function on the edges. Complex photon number distributions [24, 40] will allow for AI-based design of experiments for advanced sensing experiments. Recent insights into quantum computing circuits have demonstrated a very related representation based on information flow graphs in circuits [41], which could inspire an extension to AI-based applications in quantum computing.

In conclusion, our findings emphasize the crucial role of integrating physical science insights into AI applications, whether it's through recognizing physical symmetries or understanding the underlying representation of physical data.

References

[1] Jian-Wei Pan, Zeng-Bing Chen, Chao-Yang Lu, Harald Weinfurter, Anton Zeilinger, and Marek Żukowski. Multiphoton entanglement and interferometry. *Reviews of Modern Physics*, 84(2):777, 2012.

[2] Emanuele Pelucchi, Giorgos Fagas, Igor Aharonovich, Dirk Englund, Eden Figueroa, Qihuang Gong, Hübel Hannes, Jin Liu, Chao-Yang Lu, Nobuyuki Matsuda, Jian-Wei Pan, Florian Schreck, Fabio Sciarrino, Christine Silberhorn, Jianwei Wang, and Klaus D. Jöns. The potential and global outlook of integrated photonics for quantum technologies. *Nature Reviews Physics*, 4(3):194–208, 2022.

[3] Mario Krenn, Mehul Malik, Robert Fickler, Radek Lapkiewicz, and Anton Zeilinger. Automated search for new quantum experiments. *Physical Review Letters*, 116(9):090405, 2016.

[4] PA Knott. A search algorithm for quantum state engineering and metrology. *New Journal of Physics*, 18(7):073033, 2016.

[5] L O'Driscoll, Rosanna Nichols, and Paul A Knott. A hybrid machine learning algorithm for designing quantum experiments. *Quantum Machine Intelligence*, 1:5–15, 2019.

[6] Alexey A Melnikov, Pavel Sekatski, and Nicolas Sangouard. Setting up experimental bell tests with reinforcement learning. *Physical Review Letters*, 125(16):160401, 2020.

[7] Mario Krenn, Manuel Erhard, and Anton Zeilinger. Computer-inspired quantum experiments. *Nature Reviews Physics*, 2(11):649–661, 2020.

[8] Mario Krenn, Jakob S Kottmann, Nora Tischler, and Alán Aspuru-Guzik. Conceptual understanding through efficient automated design of quantum optical experiments. *Physical Review X*, 11(3):031044, 2021.

[9] Xavier Valcarce, Pavel Sekatski, Elie Gouzien, Alexey Melnikov, and Nicolas Sangouard. Automated design of quantum-optical experiments for device-independent quantum key distribution. *Physical Review A*, 107(6):062607, 2023.

[10] Daniel Gómez Aguado, Vicent Gimeno, Julio José Moyano-Fernández, and Juan Carlos Garcia-Escartin. Qoptcraft: A python package for the design and study of linear optical quantum systems. *Computer Physics Communications*, 282:108511, 2023.

[11] Mario Krenn, Xuemei Gu, and Anton Zeilinger. Quantum experiments and graphs: Multiparty states as coherent superpositions of perfect matchings. *Physical Review Letters*, 119(24):240403, 2017.

[12] Xuemei Gu, Manuel Erhard, Anton Zeilinger, and Mario Krenn. Quantum experiments and graphs ii: Quantum interference, computation, and state generation. *Proceedings of the National Academy of Sciences*, 116(10):4147–4155, 2019.

[13] Xuemei Gu, Lijun Chen, Anton Zeilinger, and Mario Krenn. Quantum experiments and graphs. iii. high-dimensional and multiparticle entanglement. *Physical Review A*, 99(3):032338, 2019.

[14] Lan-Tian Feng, Ming Zhang, Di Liu, Yu-Jie Cheng, Guo-Ping Guo, Dao-Xin Dai, Guang-Can Guo, Mario Krenn, and Xi-Feng Ren. On-chip quantum interference between the origins of a multi-photon state. *Optica*, 10(1):105–109, 2023.

[15] Kaiyi Qian, Kai Wang, Leizhen Chen, Zhaohua Hou, Mario Krenn, Shining Zhu, and Xiao-song Ma. Multiphoton non-local quantum interference controlled by an undetected photon. *Nature Communications*, 14(1):1480, 2023.

[16] Jueming Bao, Zhaorong Fu, Tanumoy Pramanik, Jun Mao, Yulin Chi, Yingkang Cao, Chonghao Zhai, Yifei Mao, Tianxiang Dai, Xiaojiong Chen, et al. Very-large-scale integrated quantum graph photonics. *Nature Photonics*, pages 1–9, 2023.

[17] Man-Duen Choi. Completely positive linear maps on complex matrices. *Linear algebra and its applications*, 10(3):285–290, 1975.

[18] Andrzej Jamiołkowski. Linear transformations which preserve trace and positive semidefiniteness of operators. *Reports on Mathematical Physics*, 3(4):275–278, 1972.

[19] Christophe Couteau. Spontaneous parametric down-conversion. *Contemporary Physics*, 59(3):291–304, 2018.

[20] Xuemei Gu, Lijun Chen, and Mario Krenn. Quantum experiments and hypergraphs: Multiphoton sources for quantum interference, quantum computation, and quantum entanglement. *Physical Review A*, 101(3):033816, 2020.

[21] Armin Hochrainer, Mayukh Lahiri, Manuel Erhard, Mario Krenn, and Anton Zeilinger. Quantum indistinguishability by path identity and with undetected photons. *Reviews of Modern Physics*, 94(2):025007, 2022.

[22] Robert Raussendorf and Hans J Briegel. A one-way quantum computer. *Physical Review Letters*, 86(22):5188, 2001.

[23] Jian-Wei Pan, Dik Bouwmeester, Harald Weinfurter, and Anton Zeilinger. Experimental entanglement swapping: entangling photons that never interacted. *Physical Review Letters*, 80(18):3891, 1998.

[24] Itai Afek, Oron Ambar, and Yaron Silberberg. High-noon states by mixing quantum and classical light. *Science*, 328(5980):879–881, 2010.

[25] Emanuele Polino, Mauro Valeri, Nicolò Spagnolo, and Fabio Sciarrino. Photonic quantum metrology. *AVS Quantum Science*, 2(2), 2020.

[26] Lev Vaidman, Yakir Aharonov, and David Z. Albert. How to ascertain the values of $\mathrm{sigma}_{\mathrm{x}}$, σ_{y}, and σ_{z} of a spin-1/2 particle. *Phys. Rev. Lett.*, 58:1385–1387, Apr 1987.

[27] Andrzej Ruszczynski. *Nonlinear optimization*. Princeton university press, 2011.

[28] Alexander Mordvintsev, Christopher Olah, and Mike Tyka. Inceptionism: Going deeper into neural networks, 2015. https://ai.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html.

[29] Frank Schindler, Nicolas Regnault, and Titus Neupert. Probing many-body localization with neural networks. *Physical Review B*, 95(24):245134, 2017.

[30] Alireza Seif, Mohammad Hafezi, and Christopher Jarzynski. Machine learning the thermodynamic arrow of time. *Nature Physics*, 17(1):105–113, 2021.

[31] Cynthia Shen, Mario Krenn, Sagi Eppel, and Alan Aspuru-Guzik. Deep molecular dreaming: Inverse machine learning for de-novo molecular design and interpretability with surjective representations. *Machine Learning: Science and Technology*, 2(3):03LT02, 2021.

[32] Romi Lifshitz. Quantum deep dreaming: A novel approach for quantum circuit design. *arXiv:2211.04343*, 2022.

[33] William K Wootters. Entanglement of formation of an arbitrary state of two qubits. *Physical Review Letters*, 80(10):2245, 1998.

[34] Zhi-Hao Ma, Zhi-Hua Chen, Jing-Ling Chen, Christoph Spengler, Andreas Gabriel, and Marcus Huber. Measure of genuine multipartite entanglement with computable lower bounds. *Physical Review A*, 83(6):062325, 2011.

[35] Mehul Malik, Manuel Erhard, Marcus Huber, Mario Krenn, Robert Fickler, and Anton Zeilinger. Multi-photon entanglement in high dimensions. *Nature Photonics*, 10(4):248–252, 2016.

[36] Amin Babazadeh, Manuel Erhard, Feiran Wang, Mehul Malik, Rahman Nouroozi, Mario Krenn, and Anton Zeilinger. High-dimensional single-photon quantum gates: concepts and experiments. *Physical review letters*, 119(18):180510, 2017.

[37] Manuel Erhard, Mehul Malik, Mario Krenn, and Anton Zeilinger. Experimental greenberger–horne–zeilinger entanglement beyond qubits. *Nature Photonics*, 12(12):759–764, 2018.

[38] Christian Reimer, Michael Kues, Piotr Roztocki, Benjamin Wetzel, Fabio Grazioso, Brent E Little, Sai T Chu, Tudor Johnston, Yaron Bromberg, Lucia Caspani, et al. Generation of multiphoton entangled quantum states by means of integrated frequency combs. *Science*, 351(6278):1176–1180, 2016.

[39] Pingyu Zhu, Qilin Zheng, Shichuan Xue, Chao Wu, Xinyao Yu, Yang Wang, Yingwen Liu, Xiaogang Qiang, Junjie Wu, and Ping Xu. On-chip multiphoton greenberger—horne—zeilinger state based on integrated frequency combs. *Frontiers of Physics*, 15:1–9, 2020.

[40] Natalia Bruno, Anthony Martin, Pavel Sekatski, Nicolas Sangouard, RT Thew, and Nicolas Gisin. Displacement of entanglement back and forth between the micro and macro domains. *Nature Physics*, 9(9):545–548, 2013.

[41] Abhinav Anand, Lasse Bjørn Kristensen, Felix Frohnert, Sukin Sim, and Alán Aspuru-Guzik. Information flow in parameterized quantum circuits. *arXiv:2207.05149*, 2022.