

Machine Learning Framework for Quantum Sampling of Highly-Constrained, Continuous Optimization Problems

Blake A. Wilson^{1,3,*}, Zhaxylyk A. Kudyshev^{1,3,*}, Alexander V. Kildishev¹, Sabre Kais^{2,3}, Vladimir M. Shalaev^{1,3}, and Alexandra Boltasseva^{1,3}

¹*School of Electrical and Computer Engineering, Birck Nanotechnology Center and Purdue Quantum Science and Engineering Institute, Purdue University, West Lafayette, IN, USA*

²*School of Chemistry, Purdue University, West Lafayette, IN 47907, USA*

³*The Quantum Science Center (QSC), a National Quantum Information Science Research Center of the U.S. Department of Energy (DOE), Oak Ridge, TN 37931*

**authors with equal contribution*

(Received 19 November 2021;)

In recent years, there is a growing interest in using quantum computers for solving combinatorial optimization problems. In this work, we developed a generic, machine learning-based framework for mapping continuous-space inverse design problems into surrogate quadratic unconstrained binary optimization (QUBO) problems by employing a binary variational autoencoder and a factorization machine. The factorization machine is trained as a low-dimensional, binary surrogate model for the continuous design space and sampled using various QUBO samplers. Using the D-Wave Advantage hybrid sampler and simulated annealing, we demonstrate that by repeated resampling and retraining of the factorization machine, our framework finds designs that exhibit figures of merit exceeding those of its training set. We showcase the framework's performance on two inverse design problems by optimizing (i) thermal emitter topologies for thermophotovoltaic applications and (ii) diffractive meta-gratings for highly efficient beam steering. This technique can be further scaled to leverage future developments in quantum optimization to solve advanced inverse design problems for science and engineering applications.

I. INTRODUCTION

Combinatorial optimization has recently seen tremendous progress with new algorithms and heuristics, such as simulated annealing, genetic algorithms, and adiabatic optimization [1]. Specifically, the quadratic unconstrained binary optimization (QUBO) formalism of combinatorial optimization has attracted significant interest due to its applicability to a broad range of physical and NP-hard (i.e., non-deterministic polynomial-time hard) combinatorial optimization problems [2–4]. For example, it has been demonstrated that QUBO can be used for factoring integers [5], electronic structure calculations [6], capital budgeting [7], solving the maximum cut problem [8,9], graph coloring [10], traffic flow optimization [11], number partitioning [12], etc. Another key aspect that boosted interest in the QUBO formalism is its equivalence to Ising Hamiltonians, commonly used in physics and chemistry [13]. This equivalence goes both

ways, enabling the direct conversion of many physics/chemistry problems into the combinatorial optimization domain. Likewise, various physical platforms can perform highly efficient QUBO-based combinatorial optimization via physical processes involving Ising Hamiltonians [14–16].

Recent progress in the development of various near-term quantum computing platforms opens up more efficient ways for addressing the aforementioned optimization problems in terms of time and computational resource requirements by leveraging the power of physical mechanisms, specifically, quantum mechanics, in the processing unit. For example, the D-Wave's quantum annealers are actively used for solving QUBO problems via encoding the QUBO parameters into a system of coupled superconducting qubits and retrieving the lowest energy configuration via quantum annealing [17,18]. Surmounting evidence is showing that quantum annealing offers a so-called quantum

speedup over classical QUBO sampling methods [19,20]. The QUBO-based optimization consists of three main steps: (i) reformulating the optimization problem into a QUBO model; (ii) embedding the retrieved QUBO model parameters into the QUBO sampler; and (iii) retrieving the global solution of the problem. In most cases, the first step in QUBO-based optimization is realized by exploiting a one-to-one correspondence between the combinatorial problem under consideration and the architecture of the QUBO-solver [21,22]. On the one hand, this correspondence makes retrieving corresponding QUBO parameters of the problem trivial. On the other hand, it significantly reduces the types of problems considered, especially those outside the combinatorial domain.

Another important subfield of optimization problems is continuous optimization. Continuous optimization is built on a continuous domain, real-space parameters with differential, calculus-based relationships. Some simple continuous optimization problems can be solved exactly with simple functions. However, many problems do not have these analytic solutions and are solved numerically by employing search algorithms such as stochastic gradient descent and various heuristics, which cannot guarantee optimality. Many problems exist for which these search algorithms do not work well because of the vast continuous domain. Novel techniques that enable invertible mapping from continuous space optimization problems to QUBO problems may provide a way to take advantage of recent and future advancements in QUBO, including quantum optimization algorithms. Therefore, there is an apparent demand for a universal method of mapping continuous optimization problems into the QUBO formalism to generate better continuous space solutions.

Within this work, we developed a novel machine-learning assisted framework that maps a broad range of continuous optimization problems into QUBO problems and samples their optimized solutions using any available classical or quantum QUBO solver. Specifically, we demonstrated a binary variational autoencoder (bVAE) assisted QUBO framework (bVAE-QUBO) that encodes a continuous optimization problem into a binary, compressed space and samples this compressed space with quantum-assisted QUBO samplers. We showcase the performance of the developed framework on two practical, continuous

optimization problems of nanophotonics: Eq. (1) optimization of the topology of a thermal emitter for thermophotovoltaics (TPV) and Eq. (2) optimization of the topology of a dielectric, diffractive meta-grating for beam steering. Although the developed technique is showcased on inverse design problems of nanophotonics, it can be directly applied to a broad range of practical continuous optimization problems, e.g., in mechanical engineering, chemistry, material synthesis. By employing a quantum-assisted algorithm for continuous optimization, our framework provides a long-sought-after example of a quantum-assisted, machine learning algorithm that has potential for quantum-supremacy [23], uses noisy intermediate-scale quantum platforms for practical engineering problems [24], and could fully leverage future developments in quantum and classical QUBO sampling [25,26].

II. METHODS

A. bVAE-QUBO general framework

Motivated by the increasing number of qubits in D-Wave quantum annealers [14,17,18] and the recent work by Hastings in proving a relativized speedup for “stoquastic” adiabatic quantum computing [27], we developed a framework to map highly constrained continuous optimization problems into the QUBO model, which can be minimized by quantum annealers and other QUBO samplers. The first step of the developed technique is compressing a problem dataset from a discretized, continuous optimization problem onto a binary, compressed space by training a binary variational autoencoder (bVAE). We then map the binary, compressed space into an equivalent QUBO problem by training a second-order factorization machine and retrieving corresponding parameters of the QUBO model [28]. Finally, we optimize the retrieved QUBO problem to find an optimal binary vector via a QUBO sampler like the D-Wave quantum annealer. The factorization machine is retrained and resampled repeatedly using a QUBO sampler until it converges to produce good, continuous space solutions. Below we highlight each step of the process in more detail (Fig. 1.).

Step #1 of the developed framework is to train the bVAE network and construct a binary, compressed

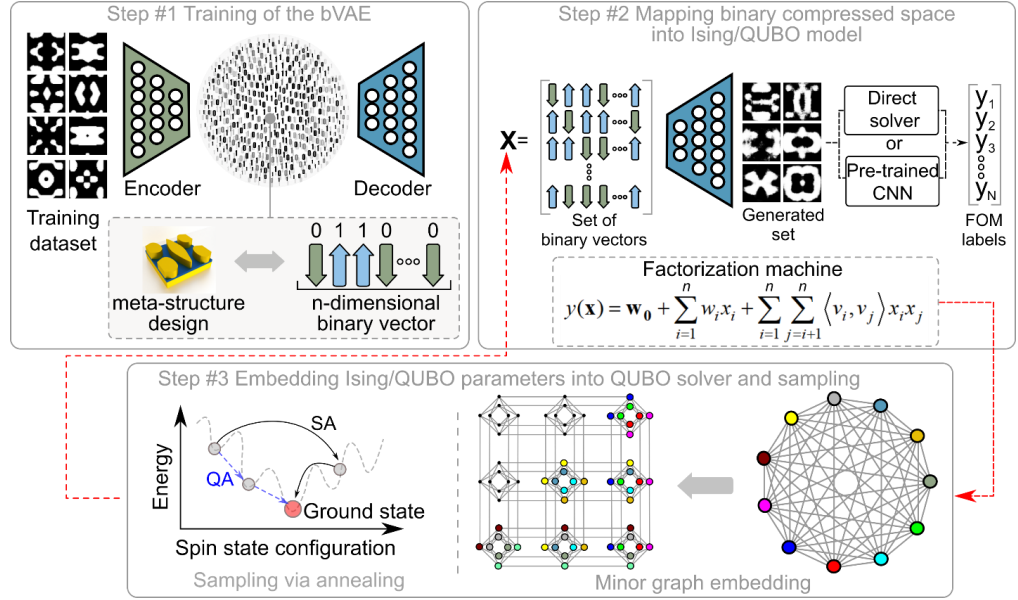


FIG. 1. The developed general bVAE-QUBO framework steps: (Step 1) training of the binary variational autoencoder (bVAE) and construction of binary compressed space representation of the dataset (inset); (Step 2) mapping the resulting binary compressed space into the QUBO/Ising model via training of the second order factorization machine; (Step 3) embedding the retrieved Ising/QUBO model parameters into the hardware and sampling the optimized data from the QUBO sampler.

representation of the optimization problem. This step maps the continuous optimization problem onto the binary domain and substantially reduces the dimension of the continuous space problem. Conventionally used variational autoencoders, which consist of two coupled neural networks (encoder and decoder), construct an invertible mapping between a lower-dimensional encoding of a dataset to the continuous-space solution [29,30]. Here, we use a limiting case of variational autoencoders – binary variational autoencoder, which constructs a binary, compressed space representation of the complex datasets [33]. More details on the bVAE structure and the training process can be found in Appendix A.

Step #2 of the bVAE-QUBO framework maps the bVAE's compressed space into a QUBO problem or Ising Hamiltonian via training a second-order factorization machine. This step exploits the fact that second-order factorization machines are equivalent to

QUBO objective functions and Ising Hamiltonians. Factorization machines, introduced by Rendle for learning sparse feature interactions [28], are low-capacity models that infer coupling coefficients, $\langle v_i, v_j \rangle$, by a factorization matrix, $V \in \mathbb{R}^{n \times k}$. The coupling coefficients, $\langle v_i, v_j \rangle$, are determined by taking the dot product of the i^{th} and j^{th} rows in V , which is equivalent to multiplying the factorization matrix by its transpose, VV^T . A factorization machine acts on an input binary vector, $x \in \{0,1\}^n$, and returns a figure of merit $y \in \mathbb{R}$.

$$y(x) = w_0 + \sum_{i=1}^n w_i x_i + \sum_{i=1}^n \sum_{j=i+1}^n \langle v_i, v_j \rangle x_i x_j, \quad (1)$$

$w_0 \in \mathbb{R}$ is a global bias, $w \in \mathbb{R}^n$ defines the weights for the discrete components of x . All free parameters, w_0, w , and $\langle v_i, v_j \rangle$ are optimized via supervised training of the factorization machine. Specifically, the factorization machine is trained on a randomly sampled dataset of binary vectors X and their corresponding

figure of merit labels Y from the binary, latent space of the bVAE. The figure of merit labels are calculated by passing a binary vector, $x \in X$, through the bVAE's decoder and calculating the figure of merit on its continuous-space solution.

A crucial benefit to restricting the factorization machine to a second-order model is that QUBO objective functions are only second-order polynomials. A QUBO sampler finds the minimum input binary vector to a second-order pseudo-boolean function via classical or quantum sampling algorithms.

$$\operatorname{argmin}_{x \in \{0,1\}^n} \sum_{i=0}^n Q_i x_i + \sum_{i_1=0}^n \sum_{i_2 > i_1}^n Q_{i_1, i_2} x_{i_1} x_{i_2}, \quad (2)$$

here $x \in \{0,1\}^n$ and Q is a $n \times n$ matrix containing local biases (diagonal terms) and coupling coefficients (off-diagonal terms). Comparing Eqs. (1) and (2), we can see a similarity between the parameters of the factorization machine and QUBO objective function:

$$w_i \leftrightarrow Q_i, \quad \langle v_i, v_j \rangle \leftrightarrow Q_{i,j} \quad (4)$$

Hence, by training the factorization machine regression model and retrieving w_0, w , and coupling matrix $\langle v_i, v_j \rangle$, we can create an equivalent QUBO model for a classical or quantum QUBO sampler. See Appendix D for a discussion on the relationship between QUBO and Ising Hamiltonians.

Step #3 embeds the retrieved QUBO/Ising problem from the factorization machine in Step #2 into a sampler and retrieves optimized solutions. While this step depends on the type and architecture of the QUBO sampler, we used the D-Wave quantum annealers and classical simulated annealing in our work. Nevertheless, we include more details on these topics in Appendix E.

Once the factorization machine's parameters are embedded into the QUBO sampler, the sampler will return a set of optimized binary vectors, x_{new} . The corresponding figure of merit labels y_{new} for the sampled vectors are assessed by generating their continuous-space solutions via the bVAE's decoder and retrieving the corresponding figure of merit values through a direct solver. Finally, the (X, Y) set initially used for training the factorization machine during the previous iteration is updated with (x_{new}, y_{new}) by appending the new vectors to the dataset. Step 3 concludes one iteration of the bVAE-QUBO, while the

next iteration starts with retraining the factorization machine on the updated set $(X \cup x_{new}, Y \cup y_{new})$. Steps #2 and #3 are looped until the figure of merit of the sampled solutions reaches saturation or the desired number of iterations is achieved. The main idea is that retraining the factorization machine on sampled vectors increases the variance of the model and forces it to be a better surrogate model for the continuous solution space. Sampling the newly trained factorization machine should not give any previous binary solution unless the sample has a high enough figure of merit, indicating a saturated factorization machine. In the next section, we showcase the performance of the developed bVAE-QUBO framework on the optimization of meta-structure designs for nanophotonic applications.

III. RESULTS

A. bVAE-QUBO for nanophotonic inverse design problems

Recently, new optimization frameworks, such as topology optimization [36–40] and metaheuristics [41,42], for nanophotonics have emerged as powerful design algorithms. However, these techniques require significant computational resources and have exponential asymptotic complexity with respect to problem constraints and the dimensions of the optimization parametric space. In response, various machine learning and deep learning algorithms have been adapted to address optimization problems in nanophotonics [43–52]. For example, generative adversarial networks [53,54] and adversarial autoencoders coupled with adjoint topology optimization techniques for optimizing meta-structures produced nonintuitive designs. It was also demonstrated that adversarial autoencoder-based optimization frameworks coupled with metaheuristic algorithms could perform global optimization searches within the compressed design spaces of a pre-trained adversarial autoencoder network [55,56]. However, due to the general complexity of the inverse design problems, such approaches may not be effective in the case of highly constrained problems, which demand multi-objective optimization within the high-dimensional latent spaces.

Within this section, we show empirical evidence that the developed bVAE-QUBO framework can address the aforementioned challenges. Specifically, we

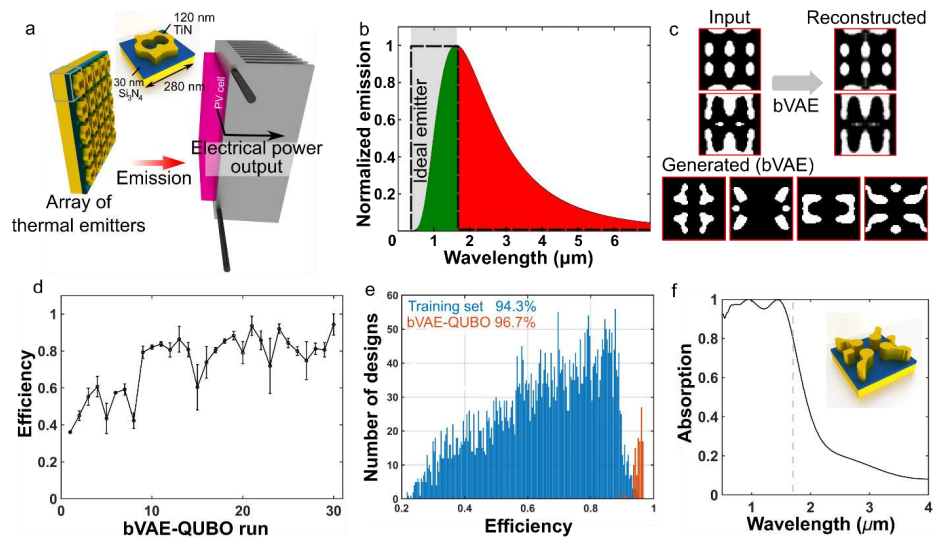


FIG. 2. Simulated annealing assisted bVAE-QUBO for thermal emitter design optimization. (a) Schematic of a thermophotovoltaic engine: a heater patterned with a thermal emitter array and a photovoltaic cell. Inset shows the base structure under consideration consists of a 300-nm-thick TiN back reflector, 30-nm-thick Si_3N_4 dielectric spacer, and 120-nm-thick top TiN patterned layer in $280 \times 280 \text{ nm}^2$ unit cell. The top TiN layer is set to be the optimization region. (b) Blackbody radiation of the bare heater (solid black curve) corresponding to emission of blackbody at 1800°C . The grey rectangular region highlights the GaSb photovoltaic cell working band. Only in-band radiation is converted into electrical power (green area), while out-of-band radiation causes heating of the photovoltaic cell (red area). Black dashed contour corresponds to an ideal thermal emitter's emissivity/absorption spectrum. (c) Two examples of input and reconstructed emitter designs by the trained bVAE network (top) and examples of randomly sampled thermal emitters (bottom). Patterns are top view of the optimization area, white color corresponds to TiN, while black corresponds to air. (d) Convergence plot of 30 iterations of the bVAE-QUBO framework. (e) Efficiency distribution of the training set used for bVAE training (5000 designs, blue histogram) and 100 designs generated via bVAE-QUBO (orange). (f) Emissivity/absorption spectra of the best design sampled with bVAE-QUBO. Vertical dashed line shows the upper bound of the GaSb photovoltaic's working band. Inset shows the schematic of the thermal emitter design.

demonstrate that by using the bVAE-QUBO framework, it is possible to construct a binary, compressed space representation of meta-devices with complex shapes and topologies and map it into a QUBO sampler for optimized, free-form design sampling. In the remainder of this section, we show the results from applying our framework to two case studies of optimizing (i) thermal emitters for thermophotovoltaics and (ii) dielectric, free-form diffractive grating for beam steering.

B. Thermal emitter for TPV application.

The TPV engine generates electrical power via radiative heat transfer between a heater and an array of

photovoltaic cells (Fig. 2(a)). High-efficiency power generation in a TPV system requires maximizing the portion of the emission that overlaps with the working band of the photovoltaic cell (in-band radiation, green area in Fig. 2(b)) and minimizing the rest of the spectra (out-of-band radiation, red area in Fig. 2(b)) [57–59]. There are three main requirements for implementing a high-efficiency TPV engine: (i) high temperature of the heater ($>1000^\circ\text{C}$), (ii) refractory material platform for the elements of the TPV system to ensure stable performance of the device at high temperatures, and (iii) pre-optimized emissivity properties of the heater.

While the first two constraints can be addressed by choosing a suitable refractory material platform [60–

62], the third requirement can be fulfilled by patterning the surface with a properly designed thermal emitting metasurface. In the ideal case, the emissivity should completely overlap the working band of the photovoltaic cell (black dotted step-function, Fig. 2(b)). Such surface emissivity ensures total cancelation of the parasitic out-of-band radiation, which leads to the reduction of the photovoltaic efficiency due to the heating while maintaining the maximum possible free-carrier generation rate. We consider TPV systems utilizing GaSb photovoltaic cells with a working band ranging from λ_{min} to λ_{max} (shaded area in Fig. 2(b)).

Within this work, we consider a gap plasmon metasurface [63,64] configuration consisting of an optically thick back titanium nitride (TiN) reflector, a 30-nm-thick silicon nitride (Si_3N_4) spacer, and a 120-nm-thick top layer (optimization region), with a fixed 280-nm lateral periodicity (inset, Fig. 2(a)). The main goal of the optimization is to determine the topological shape of the material distribution (TiN and air) within the optimization region, which ensures spectral emissivity matching the emissivity of the ideal emitter. For quantification of the device performance, we define the efficiency of the thermal emitter as a product of in-band (eff^{in}) and out-of-band efficiencies (eff^{out}). eff^{in} is an in-band radiance of the emitter normalized to the in-band radiance of ideal emitter at 1800 C, while out-of-band efficiency eff^{out} is defined as a ratio of the out-of-band radiance of the back reflector and radiance of the optimized design.

The first step of bVAE-QUBO is realized by training the bVAE network on topology optimized thermal emitter designs. The training set consists of 5000 topology optimized designs obtained via a previously developed adversarial autoencoder-based optimization framework. A VGGnet regression model trained on the same dataset for rapidly estimating the thermal emitter's efficiency based on its design. More information on training set generation is in Appendix F. During the training of the bVAE, the encoder takes 64×64 binary, greyscale topology images (top view of the antenna, Fig. 2(c)) as an input and trains to compress it into the 500-dimensional binary space. Likewise, the decoder trains to reconstruct the topology of the antenna design based on an inputted 500-dimensional binary vector. After training, the encoder can compress antenna designs into

the compressed latent space while the decoder can act as a generator that maps latent vectors to novel meta-structures. Two examples of the reconstructed antenna designs are shown in Fig. 2(c). Here one can see that the bVAE network ensures precise reconstruction of complicated antenna designs. Additionally, Fig. 2(c) shows some examples of the randomly sampled antenna designs using the bVAE's decoder. Note that the Gaussian filtering with 20 nm blur size is applied to the generated patterns to eliminate small features introduced by the bVAE noise.

The second step of the bVAE-QUBO framework starts by training the factorization machine on the binary vectors (X) and efficiency labels (Y) generated by the bVAE network. This data set is constructed by randomly sampling 11250 binary vectors from the binary, compressed space and calculating their corresponding efficiency labels. The efficiency labels are retrieved by generating their thermal emitter design using the decoder and calculating their efficiency using a pre-trained VGGnet. The (X, Y) set is constrained such that half of it corresponds to low-efficiency designs (70-80% efficiency), 30% of the designs in the set have moderate efficiencies (between 80% to 90%), and 20% of them have more than 90% efficiency. The supervised training of the factorization machine is done using the adaptive gradient descent optimization with the mean square error loss function. 70% of the (X, Y) dataset is used for training, while 10% for validation and 20% is used for testing. The trained factorization machine ensures $r^2 = 72\%$ and mean square error of 0.001. Additional information on the structure of the bVAE network and training the factorization machine can be found in Appendix A and C.

1. Simulated annealing-assisted bVAE-QUBO framework

Using simulated annealing as the QUBO sampler, the bVAE-QUBO framework executed 30 iterations. Thermal emitter design efficiencies sampled during each of the bVAE-QUBO runs are shown in Fig. 2(d). The data points and error bars show the mean efficiencies and corresponding standard deviations of 10 designs sampled during each of the bVAE-QUBO runs. We can see that updating QUBO parameters via retraining the factorization machine on the newly sampled vectors will significantly increase the quality of

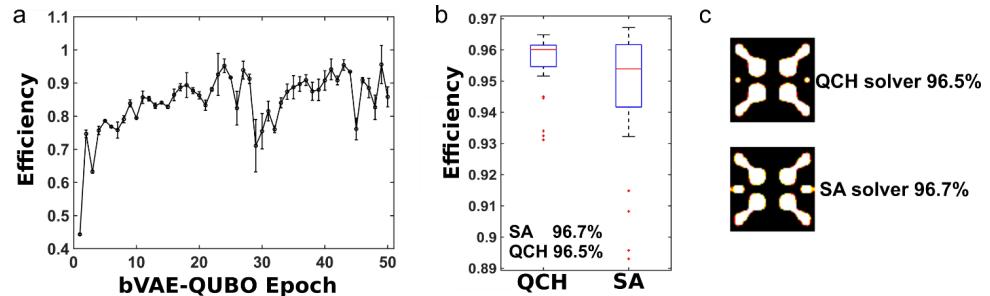


FIG. 3. bVAE-QUBO assisted with a hybrid (quantum-classical) sampler for optimizing high-efficiency thermal emitter designs. (a) Convergence plot of 50 iterations of quantum-classical hybrid sampler assisted bVAE-QUBO. (b) Efficiency distribution comparison of the top 100 thermal emitter designs sampled via the hybrid (left) and simulated annealing (right) samplers. Box plot shows the median (red line), interquartile range (box) and outliers (red markers). Here, labeling indicates maximum efficiencies within each set. (c) Top view of the best thermal emitter designs sampled via the hybrid sampler (top) and simulated annealing (bottom) assisted bVAE-QUBO (white color corresponds to TiN, while black corresponds to air).

sampled designs (from ~40% of the initially trained factorization machine up to >90%). Figure 2(e) shows the efficiency distributions of the dataset used for bVAE training (5000 designs, blue bars) and the best 100 designs sampled with bVAE-QUBO (orange bars). Finite difference time domain analysis (Lumerical FDTD) is used to assess the final efficiencies of each of the sets after running the framework. The best design in the training set ensures 94.3% efficiency, while one sampled via the simulated annealing-based bVAE-QUBO framework approach ensures 96.7%. The emissivity spectra of the best design in the bVAE-QUBO set are shown in Fig. 2(f), while the inset shows the corresponding design of the thermal emitter.

2. Quantum-classical hybrid-assisted bVAE-QUBO framework

Along with the simulated annealer, we tested the bVAE-QUBO framework based on the quantum-classical hybrid sampler. The hybrid sampler is a high-quality server-side sampler hosted in the D-Wave Leap ecosystem that uses a mix of quantum annealing and classical sampling to sample from large QUBO's. Fig. 3(a) shows the convergence plot of emitter efficiencies generated with hybrid sampling. As in the previous case, retraining the factorization machine with a refined dataset substantially increases the bVAE-QUBO framework's performance. We note that the main limitation of this approach is that the hybrid sampler returns one sample per bVAE-QUBO run. To augment

the sampled dataset during the bVAE-QUBO run, we copied the sample 10 times and flipped a single bit for each copy. This allows us to expand the number of samples per epoch while sacrificing variance in the resulting designs. The comparison of the efficiency distributions obtained via the hybrid and simulated annealing-assisted bVAE-QUBO framework is shown in Fig. 3(b). Here, we can see that using the hybrid sampler ensures narrower efficiency distribution with the median at 96% and interquartile range (25th to 75th percentile) between 95.5% and 96.2%. For the comparison, the efficiency distribution of the simulated annealing-based sampling has a 95.4% median and interquartile range between 94.2% and 96.2%. Both approaches provide almost identical maximum efficiencies, 96.5% (hybrid sampler) and 96.7% (simulated annealing). Corresponding thermal emitter designs are shown in Fig. 3(c). Interestingly, both samplers lead to the designs with identical topologies, with slightly different lateral dimensions of the antenna components. Such narrow distribution of the sampled design efficiencies in the hybrid case might be a consequence of a better optimization search provided by the quantum annealing part of the sampler, which ensures a higher probability of locating an optimum in comparison with the classical simulated annealing algorithm.

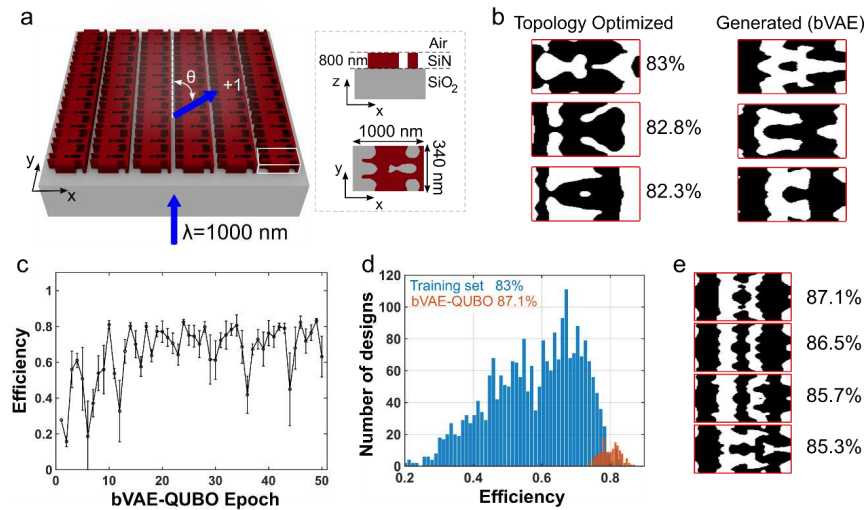


FIG. 4. bVAE-QUBO based diffractive meta-grating optimization. (a) Schematics of the meta-grating optimization domain. The main goal of topology optimization is to determine material distribution inside the optimization region (highlighted by white box) placed on the silicon dioxide substrate that ensures highest possible deflection efficiency at θ deflection angle. The inset shows the configuration of the unit cell. (b) Best designs in the topology optimized training set and examples of bVAE sampled metagrating designs (white color - SiN, black - air). (c) Convergence plot for 50 runs of the simulated annealing-assisted bVAE-QUBO framework. (d) Efficiency distributions of the training set used for bVAE training (2000 designs, blue histogram) and 100 designs generated via the bVAE-QUBO framework (orange). (e) Designs of best designs sampled by the bVAE-QUBO.

C. Inverse design of diffractive meta-gratings

In the second case study, we optimize dielectric, free-form diffractive gratings for beam steering. Different types of dielectric meta-structures, metasurfaces, and meta-gratings have been used for various imaging applications [65], spectroscopy [66], as well as integrated optics applications [67–69]. The development of the dielectric antenna designs is one of the major steps in the meta-structure design pipeline. It has recently been demonstrated that advanced optimization frameworks, such as genetic algorithms [70] and adjoint topology optimization [71], can be used to develop various types of dielectric meta-devices. We apply the bVAE-QUBO framework to optimize silicon nitride (SiN) meta-gratings for deflecting normally incident light to a pre-defined angle θ . The main goal of the optimization is to determine binary (SiN and air) material distribution within the optimization region, which maximizes the transmitted

energy of the normally incident plane wave into +1 diffraction order at a 60-degree deflection angle (Fig. 4(a)).

1. Training the bVAE network on topology optimized meta-grating designs

Adjoint topology optimization is used to generate 2000 SiN freeform meta-gratings (Appendix F) [71]. Figure 4(b) shows the best three designs in the training set. The bVAE network is trained on 100×100 pixelated patterns to compress the continuous-space representation of meta-grating into 500-dimensional binary compressed space. Some of the generated meta-grating designs by the trained bVAE network are shown in Fig. 4(b). The figure of merit labels are assessed with S4, a rigorous coupled-wave analysis (RCWA) solver [72,73]. The evolution of the efficiencies of the meta-grating sampled by the bVAE-QUBO framework is shown in Fig. 4(c). Similar to the previous example, gradual refinement of the QUBO parameters through retraining of the factorization machine significantly

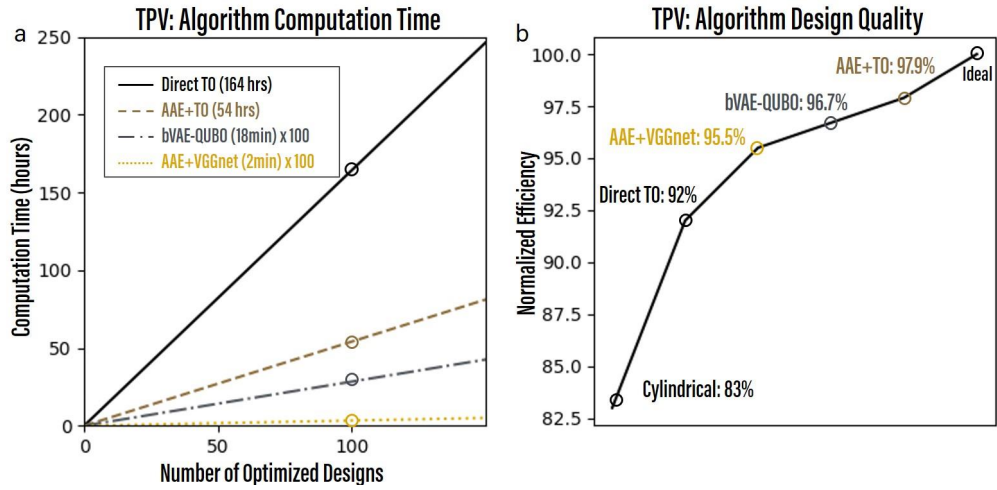


FIG. 5. Comparison of four topology optimization algorithms for generating optimal metasurfaces for the TPV application, see Section 3(b). Direct topology optimization (Direct TO) [36–40] is a gradient optimization technique that refines the topology of a metasurface by iteratively following the gradient vector. The machine learning based algorithms, such as Adversarial Autoencoder assisted TO (AAE+TO) [55,56], bVAE-QUBO, and Adversarial Autoencoder assisted TO with a VGGnet (AAE+VGGnet) [55,56], all employ an autoencoder to compress metasurface designs into a latent space and then perform a global optimization algorithm, e.g., adjoint optimization, quantum annealing, or simulated annealing, within the compressed latent space of the autoencoder. (a) Computation time (hours) for each of the aforementioned algorithms to produce 100 metasurface designs. (b) Highest efficiency metasurface generated from each algorithm for the TPV application.

increases efficiencies of sampled designs. Within each iteration, the bVAE-QUBO samples ~10 meta-grating designs, generating 500 designs in total. Figure 4(d) shows a comparison of efficiencies of the bVAE training set (2000 topology optimized designs, blue bars) and the most efficient 100 designs sampled with the simulated annealing-assisted bVAE-QUBO framework (orange bars). The best designs in the topology optimized training set ensure 83% beam deflection efficiency, while the best design in the bVAE-QUBO set ensures 87.1 %. Figure 4(e) shows the meta-grating designs of the four highest efficiency designs sampled via the bVAE-QUBO framework. The figure indicates that similar to the previous case study, the bVAE-QUBO framework samples high-efficiency designs and produces significantly better designs than conventional topology optimization. Comparing the generated design quality from bVAE-QUBO to those in [71], we note that the 95% efficiency achieved at a similar diffraction angle (75°) is the “relative

efficiency”, defined as a deflected light normalized to the total power transmitted through the device. While in our case, we have considered diffraction efficiency, i.e., the power in the diffraction order normalized to the incident light. The equivalent metric from [71] is the “absolute efficiency”, which is the power of the deflected light beam normalized to the power of light transmitted through a bare SiO₂ substrate. In this case, the result achieved in Fig. 3(b) of [71] is 79%, while bVAE-QUBO achieved 90.1%. To convert diffractive efficiency into “absolute efficiency”, we use the transmission coefficient through a SiO₂-air boundary at normal incidence (96.6%) for a 1000 nm wavelength. However, it should be noted that there are some differences between our diffractive meta-grating designs and those in [71] which may account for the large gap in theoretical, absolute efficiencies, e.g., meta-grating material (Si vs. SiN), diffraction angle (75° vs 60°), and wavelength (1050nm vs 1000nm).

With regard to runtime performance, we highlight time requirements for each of the bVAE-QUBO's steps in our case studies. Training the bVAE network takes ~20 min using PyTorch, while training of the VGGnet regression model requires ~24 min on a standard desktop (Intel Core i7, 2.8 GHz CPU, 16 GB, Nvidia GeForce GTX 1050 GPU). Both the factorization machine training and execution of the bVAE-QUBO sampling are realized on the Google Colab platform. While the simulated annealing-assisted bVAE-QUBO requires only 3 s to sample at least one design, the hybrid sampler has a preset minimum annealing time of 10 s required to sample one design. The main bottlenecks come from the QUBO sampling time and determining the figure of merit of the device. More details are available in the Appendix E.

D. Topology Optimization Algorithm Comparison

The bVAE-QUBO algorithm can readily be compared to more tailored algorithms for solving the TPV problem from Section 3(b). The two main points of comparison are runtime, Fig. 5(a), and design quality, Fig. 5(b). The clear advantage of bVAE-QUBO over other methods is that it can generate high-quality designs at a fraction of the time without relying on gradient descent methods like adjoint optimization for deciding the next sample. When compared to the fastest gradient descent-based algorithm for TPV, AAE+VGGnet, bVAE-QUBO is able to generate higher-quality designs (+1.2%) while only taking slightly longer (+4min). When compared to the algorithm that generates higher-quality designs, AAE+TO, bVAE-QUBO is able to find a slightly lower efficiency design (-1.2%) while being 540 times faster. However, because bVAE-QUBO uses quantum optimization to perform global optimization, we predict the algorithm will outperform these classical algorithms in both computation time and result quality as the problem of interest becomes more difficult, and QUBO sampling technology improves. At present, it is unknown whether bVAE-QUBO has a quantum speedup over adjoint optimization methods. To show such a speedup experimentally, one would need to show an asymptotic advantage with bVAE-QUBO over classical sampling methods, such as quantum Monte-Carlo, on progressively larger instances of a tailored quantum problem. Similar work was done by King et

al. [20] to show an asymptotic speedup for quantum annealing on frustrated magnets. We will explore this potential speedup and design quality improvement in future work.

E. Variance in Sampled Solutions During bVAE-QUBO

Looking at Fig. 2(d) and Fig. 4(c), we note there is variance in the sampled efficiencies during bVAE-QUBO. For each epoch of bVAE-QUBO, the training dataset increases by up to twenty vectors (depending on the parameters of the QUBO sampler.) Retraining the factorization machine on these new vectors and the original training set will increase the variance of the factorization machine. As the variance of the factorization machine increases, the variance of the QUBO problem also increases. We believe that as the variance of the QUBO increases, several different minima may diverge from the global extremum of the original problem. We have tested this by plotting the convergence of bVAE-QUBO for 100 epochs with the diffractive meta-grating problem.

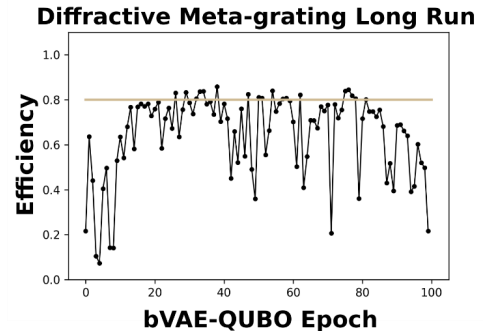


FIG. 6. Maximum efficiency sampled vector for each epoch during bVAE-QUBO. As the factorization machine begins to reach saturation, it first achieves its global maximum and keeps producing high efficiency solutions until the dataset size grows too large. The gold line shows the maximum efficiency vector in the original training set.

As shown in Fig. 6, when the training dataset grows too large compared to the factorization machine's capacity, it must sacrifice accuracy to compensate for learning so many vectors. This deteriorates the potential global

maximum of the factorization machine, so the sampled solutions begin to deteriorate in quality. The first considered approach to compensate for this is to limit the size of the dataset and replace the current dataset's minimum vectors with any newly sampled high-efficiency vectors that exceed these minimum vectors. This solution does not lead to improved results in practice. The factorization machine does not learn from the minimum vectors that are not added to the training set. So, it can become stuck and repeatedly produce low-efficiency vectors. We are interested in employing dataset manipulation in future work to overcome these limitations. Specifically, we will focus on reducing this variance in the dataset to obtain better convergence over longer bVAE-QUBO runs.

IV. Conclusion

Within this work, we developed a unique framework that maps a broad range of continuous optimization problems into QUBO problems, which can be optimized by any available QUBO sampler. The developed binary variational autoencoder-assisted QUBO framework reformulates a continuous-space problem into a QUBO problem and maps the constructed binary, compressed space into a QUBO sampler through a factorization machine model. The performance of the developed technique is demonstrated on two case studies of inverse design problems in nanophotonics, (i) thermal emitter topologies for TPV applications and (ii) diffractive meta-gratings for high-efficiency beam steering. This work is inspired by a recent factorization machine-assisted QUBO framework applied for optimization of "checkerboard" type multi-layer metamaterial structures by setting "one to one" mapping of material pixels of the structure into the system of qubits of the D-Wave machine [21]. In contrast to [21], the bVAE-QUBO framework can (i) map continuous-space optimization problems without imposing the one-to-one mapping from the QUBO solver to the problem solution space, and hence, (ii) reduce the dimension of the parametric space of the continuous domain by constructing a compressed binary space representation. By using a bVAE to map into the optimization space, our method can impose structural constraints such as manufacturability and symmetry without manipulating the cost functions. Additionally, this framework can easily be extended to continuous space datasets where optimization can be performed on classes of continuous

vectors with similar properties. Such generic formalism of the bVAE-QUBO framework opens up the possibility for mapping a broad range of highly constrained optimization problems of optics, chemistry, mechanics, finance, and computer science into any available QUBO sampler. While the current study showcased the quality of generated designs from the bVAE-QUBO-based framework on classical and quantum-classical hybrid samplers, we are not claiming an asymptotic speedup for the quantum-classical hybrid sampler or quantum annealing for inverse-design or "untailored" problems when compared to state-of-the-art heuristic solvers as discussed in [74]. However, we are planning to address this important scaling issue in our future work via a theoretical study in quantum machine learning and an experimental head-to-head competition between a quantum Monte Carlo sampler and a quantum annealer during an instance of bVAE-QUBO. Additionally, this work will be extended to similar frameworks with general, non-stoquastic Hamiltonians for adiabatic optimization.

Acknowledgments

This work is supported by the U.S. Department of Energy (DOE), Office of Science through the Quantum Science Center (QSC), a National Quantum Information Science Research Center, and Purdue's Elmore ECE Emerging Frontiers Center 'The Crossroads of Quantum and AI'.

The authors also acknowledge support from the National Science Foundation award 2029553-ECCS (thermophotovoltaics materials input parameters), DARPA/DSO Extreme Optics and Imaging (EXTREME) Program (HR00111720032) (AVK). We want to acknowledge and thank Oak-Ridge National Lab and the NASA Ames Center for allowing us to access their D-Wave 2000Q and D-Wave Advantage.

Author Contributions

Z.A.K. and B.A.W. conceived the main framework of the bVAE-QUBO. A.V.K., V.M.S., S.K., and A.B. supervised the project. Z.A.K. developed a code for the training of the bVAE network (step 1), a topology optimization framework for dataset generation, and performed FDTD and RCWA simulations. B.A.W. developed the Python libraries for implementing steps 2 and 3 of the bVAE-QUBO framework and collecting the

data for the bVAE-QUBO runs. All authors interpreted the data and wrote the manuscript.

Competing Interests

The authors declare no competing interests.

Data availability

All the data is available from the corresponding author upon reasonable request.

APPENDIX A: STRUCTURE AND TRAINING OF A BINARY VARIATIONAL AUTOENCODER

Within this work, we map the continuous design space found in many physically constrained problems into a binary design space. Recently, it has been demonstrated that variational autoencoders can be

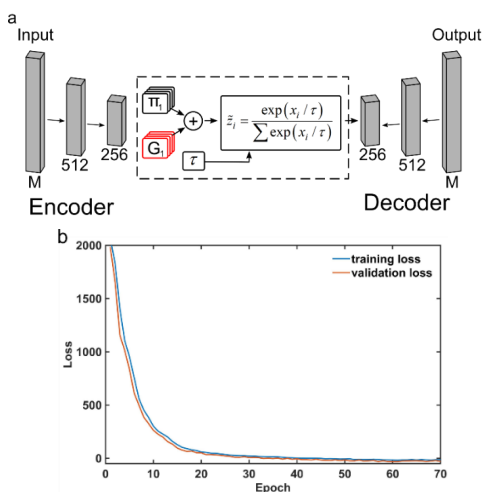


FIG. 7. Binary variational autoencoder training. (a) Structure of the bVAE network and Gumbel-softmax reparameterization of the binary latent variable. (b) Evolution of training and validation loss of bVAE network during training on thermal emitter design set.

generalized for categorical compression of complex 1D and 2D datasets, where the compressed space coordinates are discrete variable vectors [31,32]. Here, we use a limiting case of such categorical variational autoencoders – binary variational autoencoder, which allows the construction of binary, compressed space representation of the complex datasets [33]. A properly

trained bVAE can be considered an invertible map, $g : \{0,1\}^n \rightarrow \mathbb{R}^{m \times m}$, between the binary, compressed space vectors of size n and the discretized, continuous space solutions of dimensionality $m \times m$. This maps each binary vector, $x \in \{0,1\}^n$, in the domain to only one design for the problem of interest in a discretized space $\mathbb{R}^{m \times m}$. Constructing g can be done via training of a binary variational autoencoder (bVAE) on samples from the image of g . The bVAE is a deep neural network consisting of an encoder and a decoder. The encoder is a network with one input layer of $m \times m$ dimensions and two hidden dense layers with 512 and 256 neurons and ReLU activation functions. The decoder has an inverted structure to the encoder, two layers with 256 and 512 neurons, and one output layer. The structure of the network is shown in Fig. 6(a). Specifically, the bVAE network learns how to compress continuous space designs into a binary, latent space and then reconstruct them. Naturally, the bVAE decoder acts as g , and the bVAE encoder acts as the inverse of g . The bVAE is trained by minimizing both the reconstruction loss for a design and the Kullback-Leibler divergence loss, \mathcal{L}_{bVAE} [31]. The latter defines the deviation of the recognition distribution (obtained with the model data) from the pre-defined prior,

$$\mathcal{L}_{bVAE} = KL[q(z|y_m)|p(z)] - \log[p(y_m|z)] \quad (A1)$$

The main difference between the bVAE and the typical VAE network is that the priors, $p(z)$, and recognition model, $q(z|y)$, are under different distributions. The bVAE is under a Bernoulli distribution, while the typical VAE's distribution is assumed to be Gaussian. The main problem with the bVAE, like with the VAE, is that a latent variable x needs to be stochastically sampled with a pre-defined distribution to properly backpropagate the error for training the stochastic nodes in the network. This can be circumvented if we express the sample $x \sim p_\theta(z)$ such that the gradient can flow from the cost function to the set of the parameters θ (output of the encoder) without encountering stochastic nodes. For example, in a VAE network, the sampling of a latent variable with Gaussian distribution is realized by re-parameterization $x \sim N(\mu, \sigma)$ as $x = \mu + \sigma \cdot \epsilon$, where $\epsilon \sim N(0,1)$ and (μ, σ) are parameters of the encoder. This re-parameterization allows us to calculate their derivatives with respect to μ and σ and use ϵ as an additional input

parameter sampled during each training epoch. We used the Gumbel-softmax re-parameterization trick to backpropagate the error in the bVAE, which is a similar re-parameterization to the standard VAE [31,32]. Specifically, Gumbel-softmax is a re-parameterization trick for a distribution that we can smoothly deform into the categorical distribution during the training process. Gumbel-softmax samples the latent space vectors \tilde{x}_i based on the class probabilities $\pi_1, \pi_2 \dots \pi_n$ of the categorical representation as:

$$\tilde{x}_i = \frac{\exp\left[\frac{\log(\pi_i) + G_i}{\tau}\right]}{\sum_{j=1}^n \exp\left[\frac{\log(\pi_j) + G_j}{\tau}\right]}, \quad i = 1, \dots, n \quad (A2)$$

here G_i are independent and identically distributed variables sampled from Gumbel distribution $\text{Gumbel}(0,1)$. τ is a “temperature” parameter that controls how closely samples from Gumbel-softmax distribution approximates those from the true categorical distribution. During the training process, τ is gradually “annealed” from τ_{max} down to τ_{min} , which is a good approximation to a categorical latent space distribution. We swept the τ parameter from $\tau_{max} = 5$ to $\tau_{min} = 0.4$ with annealing rate $\gamma = 0.0003$. The evolution of the temperature follows an iterative form, $\tau_{\text{epoch}+1} = \tau_{\text{epoch}} \exp(-\gamma \cdot \text{epoch})$. We used a stochastic gradient descent optimization method, Adam (Adaptive Moment Estimation) [75], available through the Keras, and TensorFlow Python API during the training loop of the bVAE. The evolution of the training and validation losses of the bVAE network trained on 5000 thermal emitter designs are shown in Fig. 6(b). 85% of the design set is used for training and 15% for validation.

APPENDIX B: PSUEDO-BOOLEAN STRUCTURE OF THE FACTORIZATION MACHINE

Introduced by Rendle for learning sparse feature interactions, factorization machines are very useful, low-capacity models [28]. Consider a map, $h: \mathbb{R}^{m \times m} \rightarrow \mathbb{R}$, that calculates the figure of merit of a discretized, continuous space design. Then, $h(g): \{0,1\}^n \rightarrow \mathbb{R}$, where g is the bVAE decoder, maps a binary vector in the compressed space of the bVAE to its figure of merit. Let $\tilde{y}(x) = h[g(x)]$, then $\tilde{y}(x)$ is a

pseudo-boolean function. If we restrict the domain of the factorization machine to a binary space, then its model equation is isomorphic to an exhaustive pseudo-boolean function,

$$y(x) = \sum_{i=0}^n \langle v_i^{(1)} \rangle x_i + \sum_{i_1=0}^n \sum_{i_2>i_1} \langle v_{i_1}^{(2)}, v_{i_2}^{(2)} \rangle x_{i_1} x_{i_2} + \sum_{i_1=0}^n \sum_{i_2>i_1} \dots \sum_{i_n>i_{n-1}} \langle v_{i_1}^{(2)}, \dots, v_{i_n}^{(n)} \rangle x_{i_1} \dots x_{i_n} \quad (B1)$$

For each polynomial degree, d , there exists a factorization matrix $v^{(d)}$. Given a factorization matrix $v^{(d)} \in \mathbb{R}^{n \times k}$, we can define a coefficient for the polynomial $x_{i_1} \dots x_{i_d}$ of degree d as $\langle v_{i_1}^{(d)}, \dots, v_{i_d}^{(d)} \rangle$, which is a dot product between rows i_1, i_2, \dots, i_d in $v^{(d)}$. The advantage of factorization machines is that their polynomial coefficients are determined by row interactions in their factorization matrix, which couples a change in one coefficient to a change in all the rows associated with that coefficient. This makes them a lower capacity model than a model where each polynomial coefficient is unique and decoupled from every other coefficient. However, this means that the model can infer coefficients under sparse training sets.

If a factorization machine is trained to approximate $\tilde{y}(x)$, then we can treat it as a surrogate model to $\tilde{y}(x)$ and sample its global optimum in place of $\tilde{y}(x)$, thereby sampling globally optimal designs within its highly compressed space. Naturally, this model's space complexity can be exponentially large with respect to the size of x if $\tilde{y}(x)$ is a full pseudo-boolean function. We can make a calculated cut in the number of terms by noticing that given any polynomial of degree, d , the number of input strings where the coefficient of the polynomial contributes is 2^{n-d} . So, we argue that the highest priority coefficients are the low order polynomials where the probabilities of any coefficient from first-order terms or second-order terms contributing to the output value are $\frac{2^{n-1}}{2^n} = \frac{1}{2}$ and $\frac{2^{n-2}}{2^n} = \frac{1}{4}$ respectively. Additionally, sampling a second-order factorization machine as a surrogate model is much more feasible because QUBO solvers can minimize second-order/quadratic pseudo-boolean functions

without special quadratization transformations. By restricting the factorization machine to first and second-order terms, its model equation becomes,

$$y(x) = \sum_{i=0}^n \langle v_i^{(1)} \rangle x_i + \sum_{i_1=0}^n \sum_{i_2>i_1}^n \langle v_{i_1}^{(2)}, v_{i_2}^{(2)} \rangle x_{i_1} x_{i_2} \quad (B2)$$

A crucial benefit to restricting the factorization machine to a second-order model is that QUBO objective functions are only second-order polynomials. A QUBO sampler finds the minimum input string to a second-order pseudo-boolean function via classical or quantum sampling algorithms.

$$\underset{x \in \{0,1\}^n}{\operatorname{argmin}} \sum_{i=0}^n Q_i x_i + \sum_{i_1=0}^n \sum_{i_2>i_1}^n Q_{i_1, i_2} x_{i_1} x_{i_2} \quad (B3)$$

If we used a higher-order factorization machine, it would need to be quadratized to a second-order polynomial before being minimized by a QUBO sampler [76]. The scaling for this process can introduce an exponential number of variables or take an exponential amount of time with respect to the degree or input size of the QUBO. So, we restricted our framework to second-order factorization machines. Then we can directly map it to a QUBO without quadratization, where $Q_i = \langle v_i^{(1)} \rangle$ and $Q_{i_1, i_2} = \langle v_{i_1}^{(2)}, v_{i_2}^{(2)} \rangle$.

APPENDIX C: TRAINING THE FACTORIZATION MACHINE

For both nanophotonic applications, we found that binary vectors of size 500, i.e., $x \in \{0,1\}^{500}$, and factorization matrices of size 500 by 40, i.e., $v^{(2)} \in \mathbb{R}^{500 \times 40}$, were sufficient to find good designs. For factorization machine training, we constructed a training set of 11250 unique vectors by randomly sampling from the binary, compressed space of the bVAE and assessing the performance of the design via a pre-trained convolutional neural network (thermal emitters) or a rigorous coupled-wave analysis (diffraction gratings). The evolution of the loss function of the factorization machine for the thermal emitter showcase example is shown in Fig. 2. It is also important to note for training that QUBO samplers minimize an objective function while we want to

maximize the figure of merit for a problem. One can circumvent this issue by training the factorization machine on $c - \tilde{y}(x)$, where $\forall x \in \{0,1\}^n : c > \tilde{y}(x)$. Then, the minimization of the factorization machine corresponds to the maximization of $\tilde{y}(x)$.

APPENDIX D: MAPPING A QUBO INTO AN ISING MODEL

The restricted Ising Hamiltonian $H(\sigma)$ used by quantum annealers contains local biases h and a quadratic term proportional to a qubit coupling matrix J as:

$$H(\sigma) = \sum_{i=1}^N h_i \sigma_i^{(z)} + \sum_{\langle i,j \rangle} J_{ij} \sigma_i^{(z)} \sigma_j^{(z)} \quad (D1)$$

$\sigma_i^{(z)}$ is a Pauli-Z matrix acting on the i^{th} qubit's spin state. The spin state must collapse to an eigenvalue of -1 or +1 upon measurement. So, one can loosely consider this collapsed measurement value as a binary variable, $s_i \in \{-1, +1\}$.

$$H(s) = \sum_{i=1}^N h_i s_i + \sum_{\langle i,j \rangle} J_{ij} s_i s_j \quad (D2)$$

Comparing Eqs. (B2) and (D2), we can see the isomorphism between the parameters of the factorization machine and Ising model:

$$h_i \leftrightarrow w_i, \quad J_{ij} \leftrightarrow \langle v_i, v_j \rangle \quad (D3)$$

However, the domain of the Ising Hamiltonian is the spin vectors, $\{-1, +1\}^n$, and the domain of the factorization machine is Boolean, $\{0,1\}^n$. There does exist a trivial transformation between the two domains, namely the invertible substitution $s_i = 2x_i - 1$. Hence, by training the factorization machine regression model and retrieving w_0, w , and coupling matrix $\langle v_i, v_j \rangle$, it is possible to construct an equivalent Ising or QUBO model.

APPENDIX E: RETRIEVING OPTIMAL VECTORS FROM A QUBO SAMPLER

A. QUBO Connectivity Graph

Once the factorization machine equation is mapped into an equivalent QUBO or Ising form, we employ a

QUBO sampler to find an input vector that minimizes the model's output. One thing to keep in mind is that the coupling matrix in a QUBO, $Q_{i,j}$, forms an undirected graph. Let $G = \{V, E\}$ be the connectivity graph for a QUBO coupling matrix, Q , where $V = \{1, 2, \dots, n\}$ is the set of nodes in the graph, i.e. input bits, and $E = \{\{i, j\} \mid i, j \in V \text{ and } Q_{i,j} \neq 0 \text{ and } i \neq j\}$ is the set of edges in the graph. One would assume that the connectivity graph of a QUBO and its sampler must match to be compatible. However, this limitation only exists for QUBO samplers that use physical processes for sampling, such as quantum annealers. Additionally, minor-embedding techniques can convert Q to another graph, Q' , that is compatible with the QUBO sampler and has the same minimum input vectors as Q . Unfortunately, this process introduces many more nodes and edges, and it is possible that the number of nodes and edges in Q' exceeds that of the physical QUBO sampler.

A special case for this embedding is when considering fully connected graphs. If Q is fully connected, such as with our factorization machine's QUBO, then the maximum number of nodes for which a Q' exists for a given physical sampler is known as the "clique" size. The clique size for a QUBO sampler depends on the topology of the QUBO sampler and the number of available nodes. There are two available topologies for the D-Wave quantum annealer connectivity: Pegasus (D-Wave Advantage) [34] and Chimera (D-Wave 2000Q) [35]. The D-Wave Advantage quantum annealer has a Pegasus connectivity graph with a maximum clique size of 180 qubits, while the D-Wave 2000Q's Chimera graph clique size is 64 qubits. Along with quantum annealing, the D-Wave Leap ecosystem supports simulated annealing and quantum-classical hybrid samplers, which can handle any clique size. Naturally, the quantum annealer may ensure asymptotic speedup over classical computers with respect to the QUBO size.

B. Handling Larger QUBO Sizes

Unfortunately, we found in practice that our factorization machines required 500-dimensional input vectors to be good models, which exceeds the maximum clique size of the D-Wave Advantage. To address this, we considered two main options which extend to any optimization problem whose compressed solution space

exceeds the clique size of the QUBO solver: i) classical sampling techniques, and ii) hybrid quantum-classical decomposition sampling techniques [77]. The benefit of using classical sampling techniques, such as simulated annealing or quantum Monte-Carlo, is that the QUBO size can be as large as the computer's memory will allow and the problem can often be parallelized. The tradeoff is that the algorithm loses any advantage given by a quantum sampling scheme. Alternatively, one can use a hybrid quantum-classical decomposition solver that decomposes the problem into sub-problems for the quantum annealer to solve. These algorithms have been shown to provide better solutions in faster time than popular classical sampling methods. It is also worth noting that it may be possible to reduce the size of the final QUBO by additional regularization of the bVAE training process, which can adapt the distribution of the binary, compressed space to better match the factorization machine model. Due to the aforementioned restriction of maximum clique size, we resorted to using D-Wave's simulated annealer and quantum-classical hybrid sampler for our problems.

APPENDIX F: GENERATING A TOPOLOGY OPTIMIZED TRAINING SET

A. Thermal emitter

Within this work, we have used a previously developed adversarial autoencoder-based optimization framework to generate the training set for the bVAE-QUBO algorithm. The adversarial autoencoder network is initially trained on 200 topology optimized designs of a three-layered gap-plasmon structure. 200 topology optimized designs have been enlarged via a data argumentation technique developed in [55].

The adversarial autoencoder consists of three coupled neural networks: the encoder, the decoder/generator, and the discriminator [78]. **The encoder** takes a 4096-dimensional input vector (that corresponds to a 64×64 binary design pattern). The input is fed into the first of two fully connected, hidden layers with 512 neurons each and a ReLU activation function on the output of both layers. One batch normalization layer is coupled to the second hidden layer. The output is a 15-neuron layer. **The decoder** has the same architecture as the encoder but with the reversed sequence. The decoder generates a 4096-element output vector based on 15-dimensional, binary

input. For the output layer, we use the hyperbolic tangent activation function. **The discriminator** takes a 15-dimensional latent vector as an input and has one output neuron for binary classification (fake/real). Here we use two hidden linear layers with 512 and 256 neurons. The two hidden layers use a ReLU activation function, and the output layer uses a sigmoid function.

Once the adversarial autoencoder network is trained, the decoder generates an additional 5000 designs. Moreover, to avoid time-consuming full-wave analysis during the execution of the bVAE-QUBO, a VGGnet convolutional neural network is trained for rapid estimation of a thermal emitter's efficiency based on its design. The VGGnet takes 64 by 64 images of the design as an input and passes it through three hidden layers, consisting of convolutional layers with ReLU activation functions. Each hidden layer is followed by a max-pooling layer, which ensures the down-sampling of the feature maps. The stack of convolutional layers is followed by one fully connected layer. The final layer has a linear activation function with a mean squared error loss function for efficiency prediction. The supervised training of the VGGnet is realized on the same 5000 designs. The VGGnet regression model ensures ~93% accuracy (r^2 coefficient) of predicting the design efficiency. Training the adversarial autoencoder and VGG networks is done similarly to [55].

B. Diffractive meta-grating

For the diffractive meta-grating example, we developed a topology optimized dataset by the adjoint topology optimization method. Here we follow the optimization framework developed in [71]. The main goal of the optimization is to determine binary material distribution (SiN in air) within the optimization region, which maximizes transmission into a +1-diffraction order at a 60-degree angle of a normally incident plane wave. The optimization region is set to be a $1\text{ }\mu\text{m} \times 0.34\text{ }\mu\text{m}$ region with a thickness of $0.8\text{ }\mu\text{m}$ placed on top of the SiO_2 substrate. A $1\text{ }\mu\text{m}$ wavelength plane wave excitation occurs from the substrate side. Topology optimization attempts to maximize the transmission efficiency of the incident light into a pre-defined diffraction order via maximization of the overlap integral between total field induced by the incident $[E^{\text{fwd}}, H^{\text{fwd}}]$ and the field induced by

backward propagating adjoint field $[E^{\text{bwd}}, H^{\text{bwd}}]$. The overlap integral is calculated above the optimization region ($z = z_1$).

$$F = \left| \int_{y_{\min}}^{y_{\max}} \int_{x_{\min}}^{x_{\max}} (E^{\text{fwd}}(x, y, z_1) \times H^{\text{bwd}}(x, y, z_1) - E^{\text{bwd}}(x, y, z_1) \times H^{\text{fwd}}(x, y, z_1)) \cdot n_z dx dy \right|^2 \quad (F1)$$

The main goal of the adjoint formalism is to express the gradient $\partial F(x, y)/\partial \epsilon$ as a function of field distributions inside the optimization region induced by forward and backward (adjoint) excitation. Such formalism obtains gradients at each location of the optimization region via only two full-wave analyses. More details on the adjoint topology optimization formalism for dielectric meta-grating optimization can be found in [71].

References

- [1] K. L. Hoffman, J. Comput. Appl. Math. **124**, 341 (2000).
- [2] S. A. Cook, in *Proc. Third Annu. ACM Symp. Theory Comput. - STOC '71* (ACM Press, New York, New York, USA, 1971), pp. 151–158.
- [3] G. Kochenberger, J.-K. Hao, Z. Lü, H. Wang, and Y. Wang, J. Comb. Optim. **28**, 58 (2014).
- [4] F. Glover, G. Kochenberger, and Y. Du, 4OR **17**, 335 (2019).
- [5] S. Jiang, K. A. Britt, A. J. McCaskey, T. S. Humble, and S. Kais, Sci. Rep. **8**, 17667 (2018).
- [6] R. Xia and S. Kais, Nat. Commun. **9**, 4195 (2018).
- [7] D. J. Laughunn, Oper. Res. **18**, 454 (1970).
- [8] E. Boros and P. L. Hammer, Ann. Oper. Res. **33**, 151 (1991).
- [9] G. A. Kochenberger, J.-K. Hao, Z. Lü, H. Wang, and F. Glover, J. Heuristics **19**, 565 (2013).
- [10] F. Wang and Z. Xu, J. Heuristics **19**, 529 (2013).
- [11] F. Neukart, G. Compostella, C. Seidel, D. von Dollen, S. Yarkoni, and B. Parney, Front. ICT **4**, (2017).
- [12] B. Alidaee, F. Glover, G. A. Kochenberger, and C. Rego, J. Appl. Math. Decis. Sci. **2005**, 113 (2005).
- [13] A. Lucas, Front. Phys. **2**, (2013).
- [14] A. B. Finnila, M. A. Gomez, C. Sebenik, C.

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0060481

- Stenson, and J. D. Doll, Chem. Phys. Lett. **219**, 343 (1994).
- [15] A. Das and B. K. Chakrabarti, Rev. Mod. Phys. **80**, 1061 (2008).
- [16] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. S. Truncik, S. Uchaikin, J. Wang, B. Wilson, and G. Rose, Nature **473**, 194 (2011).
- [17] B. Apolloni, C. Carvalho, and D. de Falco, Stoch. Process. Their Appl. **33**, 233 (1989).
- [18] T. Kadowaki and H. Nishimori, Phys. Rev. E **58**, 5355 (1998).
- [19] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, SIAM Rev. **50**, 755 (2008).
- [20] A. D. King, J. Raymond, T. Lanting, S. V. Isakov, M. Mohseni, G. Poulin-Lamarre, S. Ejtemaee, W. Bernoudy, I. Ozfidan, A. Y. Smirnov, M. Reis, F. Altomare, M. Babcock, C. Baron, A. J. Berkley, K. Boothby, P. I. Bunyk, H. Christiani, C. Enderud, B. Evert, R. Harris, E. Hoskinson, S. Huang, K. Jooya, A. Khodabandelou, N. Ladizinsky, R. Li, P. A. Lott, A. J. R. MacDonald, D. Marsden, G. Marsden, T. Medina, R. Molavi, R. Neufeld, M. Norouzzpour, T. Oh, I. Pavlov, I. Perminov, T. Prescott, C. Rich, Y. Sato, B. Sheldan, G. Sterling, L. J. Swenson, N. Tsai, M. H. Volkmann, J. D. Whittaker, W. Wilkinson, J. Yao, H. Neven, J. P. Hilton, E. Ladizinsky, M. W. Johnson, and M. H. Amin, Nat. Commun. **12**, 1113 (2021).
- [21] K. Kitai, J. Guo, S. Ju, S. Tanaka, K. Tsuda, J. Shiomi, and R. Tamura, Phys. Rev. Res. **2**, 013319 (2020).
- [22] P. Kairys, A. D. King, I. Ozfidan, K. Boothby, J. Raymond, A. Banerjee, and T. S. Humble, PRX Quantum **1**, 020320 (2020).
- [23] A. Perdomo-Ortiz, M. Benedetti, J. Realpe-Gómez, and R. Biswas, Quantum Sci. Technol. **3**, 030502 (2018).
- [24] J. Preskill, Quantum **2**, 79 (2018).
- [25] S. Matsubara, H. Tamura, M. Takatsu, D. Yoo, B. Vatakhahghadim, H. Yamasaki, T. Miyazawa, S. Tsukamoto, Y. Watanabe, K. Takemoto, and A. Sheikholeslami, in *Adv. Intell. Syst. Comput.* (2018), pp. 432–438.
- [26] S. Tsukamoto, M. Takatsu, S. Matsubara, and H. Tamura, Fujitsu Sci. Tech. J. **53**, 8 (2017).
- [27] M. B. Hastings, ArXiv (2020).
- [28] S. Rendle, in *Proc. The2010 IEEE Int. Conf. Data Min.* (2010), pp. 995–1000.
- [29] Y. Pu, Z. Gan, R. Henao, X. Yuan, C. Li, A. Stevens, and L. Carin, Adv. Neural Inf. Process. Syst. **2352** (2016).
- [30] D. P. Kingma and M. Welling, 2nd Int. Conf. Learn. Represent. ICLR 2014 - Conf. Track Proc. (2013).
- [31] E. Jang, S. Gu, and B. Poole, 5th Int. Conf. Learn. Represent. ICLR 2017 - Conf. Track Proc. (2016).
- [32] C. J. Maddison, A. Mnih, and Y. W. Teh, 5th Int. Conf. Learn. Represent. ICLR 2017 - Conf. Track Proc. (2016).
- [33] R. Sicks, R. Korn, and S. Schwaar, ArXiv (2020).
- [34] N. Dattani, S. Szalay, and N. Chancellor, ArXiv (2019).
- [35] J. King, S. Yarkoni, M. M. Nevisi, J. P. Hilton, and C. C. McGeoch, (2015).
- [36] R. E. Christiansen and O. Sigmund, ArXiv (2020).
- [37] J. S. Jensen and O. Sigmund, Appl. Phys. Lett. **84**, 2022 (2004).
- [38] S. Molesky, Z. Lin, A. Y. Piggott, W. Jin, J. Vucković, and A. W. Rodriguez, Nat. Photonics **12**, 659 (2018).
- [39] C. M. Lalau-Keraly, S. Bhargava, O. D. Miller, and E. Yablonovitch, Opt. Express **21**, 21693 (2013).
- [40] Z. Lin, V. Liu, R. Pestourie, and S. G. Johnson, Opt. Express **27**, 15765 (2019).
- [41] S. Jafar-Zanjani, S. Inampudi, and H. Mosallaei, Sci. Rep. **8**, 11040 (2018).
- [42] D. Z. Zhu, E. B. Whiting, S. D. Campbell, D. B. Burckel, and D. H. Werner, ACS Photonics **6**, 2741 (2019).
- [43] W. Ma, F. Cheng, and Y. Liu, ACS Nano **12**, 6326 (2018).
- [44] J. Peurifoy, Y. Shen, L. Jing, Y. Yang, F. Cano-

- Renteria, B. G. DeLacy, J. D. Joannopoulos, M. Tegmark, and M. Soljačić, *Sci. Adv.* **4**, eaar4206 (2018).
- [45] L. Jin, Y.-W. Huang, Z. Jin, R. C. Devlin, Z. Dong, S. Mei, M. Jiang, W. T. Chen, Z. Wei, H. Liu, J. Teng, A. Danner, X. Li, S. Xiao, S. Zhang, C. Yu, J. K. W. Yang, F. Capasso, and C.-W. Qiu, *Nat. Commun.* **10**, 4789 (2019).
- [46] C. C. Nadell, B. Huang, J. M. Malof, and W. J. Padilla, *Opt. Express* **27**, 27523 (2019).
- [47] Y. Kiarashinejad, M. Zandehshahvar, S. Abdollahramezani, O. Hemmatyar, R. Pourabolghasem, and A. Adibi, *Adv. Intell. Syst.* **2**, 1900132 (2020).
- [48] Y. Kiarashinejad, S. Abdollahramezani, and A. Adibi, *Npj Comput. Mater.* **6**, 12 (2020).
- [49] I. Sajedian, J. Kim, and J. Rho, *Microsystems Nanoeng.* **5**, 27 (2019).
- [50] Z. Liu, D. Zhu, K. Lee, A. S. Kim, L. Raju, and W. Cai, *Adv. Mater.* **32**, 1904790 (2020).
- [51] W. Ma, Z. Liu, Z. A. Kudyshev, A. Boltasseva, W. Cai, and Y. Liu, *Nat. Photonics* **15**, 77 (2021).
- [52] J. Jiang, M. Chen, and J. A. Fan, *Nat. Rev. Mater.* (2020).
- [53] J. Jiang, D. Sell, S. Hoyer, J. Hickey, J. Yang, and J. A. Fan, *ACS Nano* **13**, 8872 (2019).
- [54] J. Jiang and J. A. Fan, *Nano Lett.* **19**, 5366 (2019).
- [55] Z. A. Kudyshev, A. V. Kildishev, V. M. Shalaev, and A. Boltasseva, *Appl. Phys. Rev.* **7**, 021407 (2020).
- [56] Z. A. Kudyshev, A. V. Kildishev, V. M. Shalaev, and A. Boltasseva, *Nanophotonics* **10**, 371 (2020).
- [57] A. Lenert, D. M. Bierman, Y. Nam, W. R. Chan, I. Celanović, M. Soljačić, and E. N. Wang, *Nat. Nanotechnol.* **9**, 126 (2014).
- [58] D. M. Bierman, A. Lenert, W. R. Chan, B. Bhatia, I. Celanović, M. Soljačić, and E. N. Wang, *Nat. Energy* **1**, (2016).
- [59] Y. X. Yeng, M. Ghebrebrhan, P. Bermel, W. R. Chan, J. D. Joannopoulos, M. Soljacic, and I. Celanovic, *Proc. Natl. Acad. Sci.* **109**, 2280 (2012).
- [60] H. Reddy, U. Guler, Z. Kudyshev, A. V. Kildishev, V. M. Shalaev, and A. Boltasseva, *ACS Photonics* **4**, 1413 (2017).
- [61] M. Chirumamilla, A. Chirumamilla, Y. Yang, A. S. Roberts, P. K. Kristensen, K. Chaudhuri, A. Boltasseva, D. S. Sutherland, S. I. Bozhevolnyi, and K. Pedersen, *Adv. Opt. Mater.* **5**, 1700552 (2017).
- [62] L. Gui, S. Bagheri, N. Strohfeldt, M. Hentschel, C. M. Zgrabik, B. Metzger, H. Linnenbank, E. L. Hu, and H. Giessen, *Nano Lett.* **16**, 5708 (2016).
- [63] F. Ding, Y. Yang, R. A. Deshpande, and S. I. Bozhevolnyi, *Nanophotonics* **7**, 1129 (2018).
- [64] H.-H. Hsiao, C. H. Chu, and D. P. Tsai, *Small Methods* **1**, 1600064 (2017).
- [65] P. Genevet and F. Capasso, *Reports Prog. Phys.* **78**, 024401 (2015).
- [66] A. Pors, M. G. Nielsen, and S. I. Bozhevolnyi, *Optica* **2**, 716 (2015).
- [67] C. J. Chang-Hasnain and W. Yang, *Adv. Opt. Photonics* **4**, 379 (2012).
- [68] O. Quevedo-Teruel, H. Chen, A. Díaz-Rubio, G. Gok, A. Grbic, G. Minatti, E. Martini, S. Maci, G. V. Eleftheriades, M. Chen, N. I. Zheludev, N. Papasimakis, S. Choudhury, Z. A. Kudyshev, S. Saha, H. Reddy, A. Boltasseva, V. M. Shalaev, A. V. Kildishev, D. Sievenpiper, C. Caloz, A. Alù, Q. He, L. Zhou, G. Valerio, E. Rajo-Iglesias, Z. Sipus, F. Mesa, R. Rodríguez-Berral, F. Medina, V. Asadchy, S. Tretyakov, and C. Craeye, *J. Opt.* **21**, 073002 (2019).
- [69] N. Yu and F. Capasso, *Flat Optics with Designer Metasurfaces*. (Nature Publishing Group, 2014), pp. 139–50.
- [70] S. D. Campbell, D. Sell, R. P. Jenkins, E. B. Whiting, J. A. Fan, and D. H. Werner, *Opt. Mater. Express* **9**, 1842 (2019).
- [71] D. Sell, J. Yang, S. Doshay, R. Yang, and J. A. Fan, *Nano Lett.* **17**, 3752 (2017).
- [72] V. Liu and S. Fan, *Comput. Phys. Commun.* **183**, 2233 (2012).
- [73] C. M. Roberts, S. Inampudi, and V. A. Podolskiy, *Opt. Express* **23**, 2764 (2015).
- [74] S. Mandrà, Z. Zhu, W. Wang, A. Perdomo-Ortiz, and H. G. Katzgraber, *Phys. Rev. A* **94**, 022337 (2016).
- [75] F. Chollet, *Keras.Io* (2015).
- [76] N. Dattani, *ArXiv* (2019).
- [77] M. Booth, S. P. Reinhardt, and A. Roy, *Quantum*

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

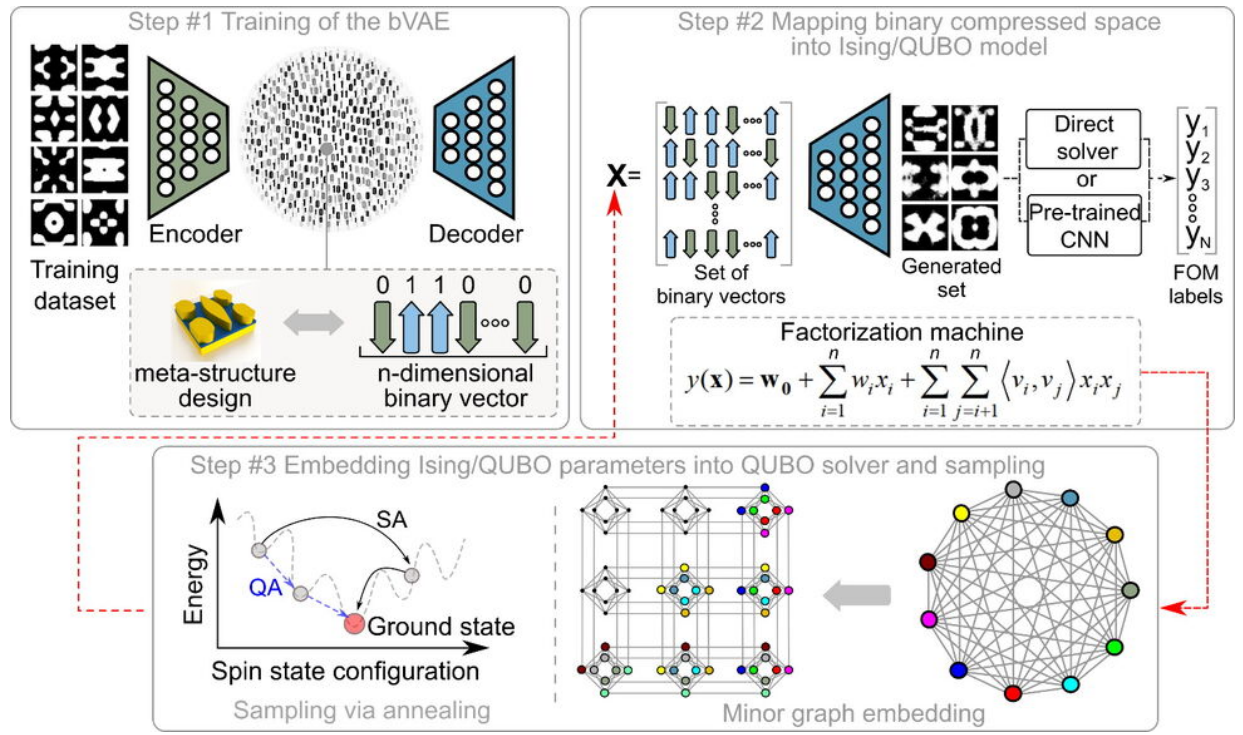
PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0060481

Exec. Tech. Rep. (2017).

[78] A. Makhzani, J. Shlens, N. Jaitly, I. Goodfellow, and B. Frey, arXiv:1511.05644 (2015).

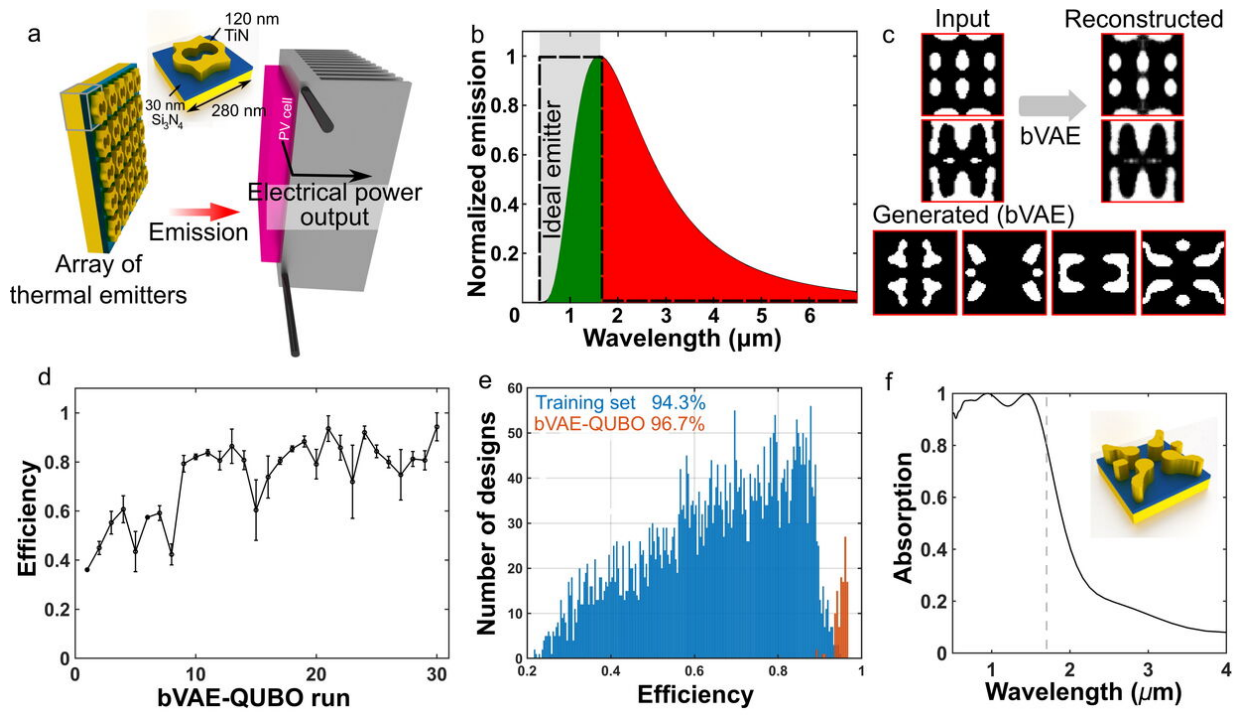
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0060481



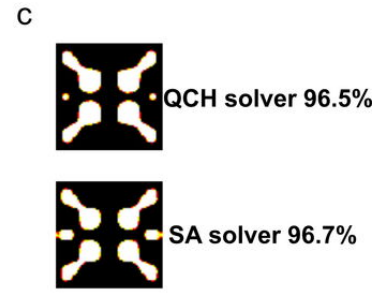
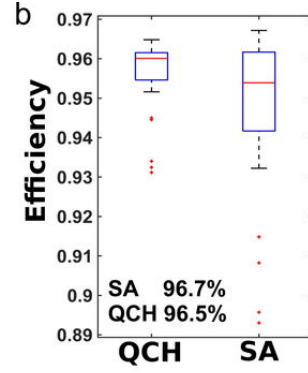
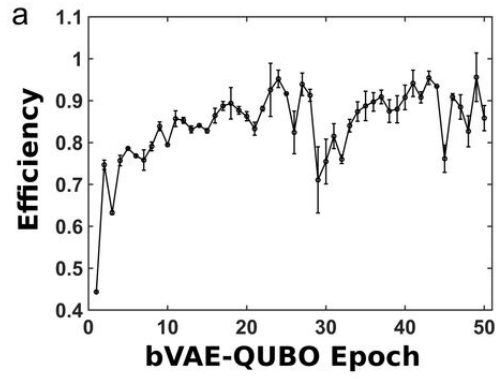
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0060481



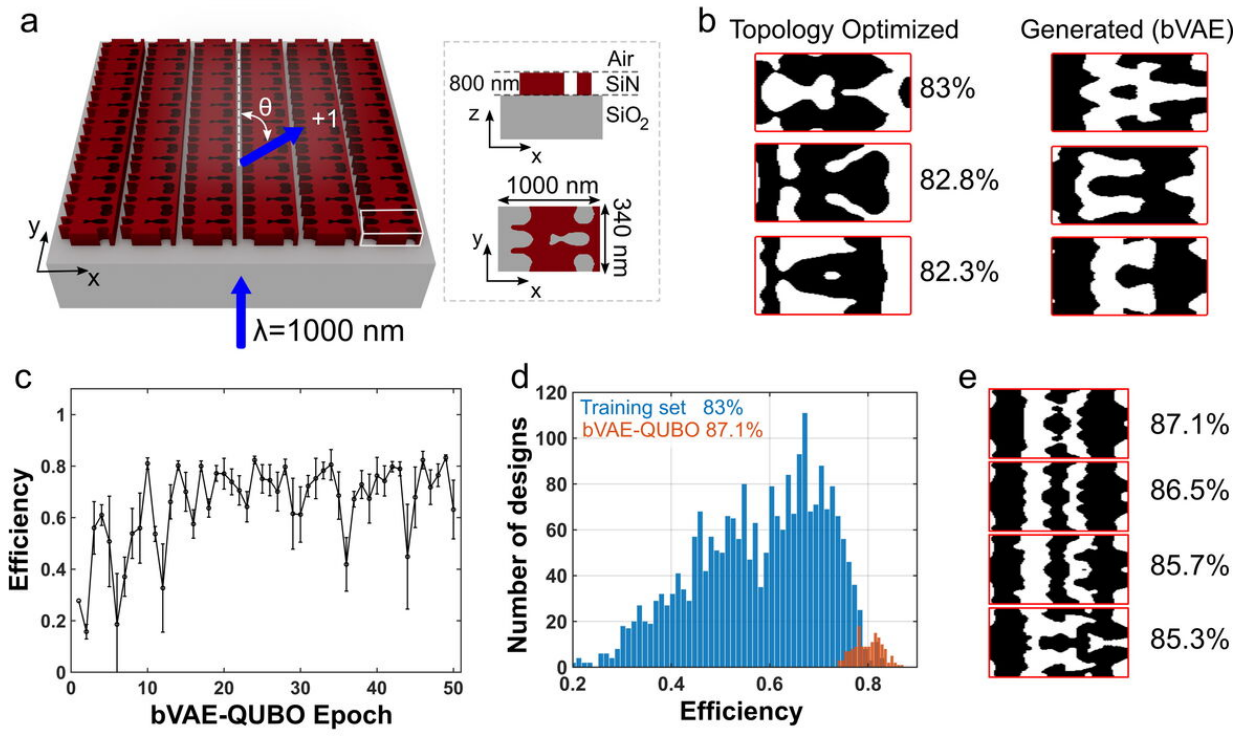
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0060481



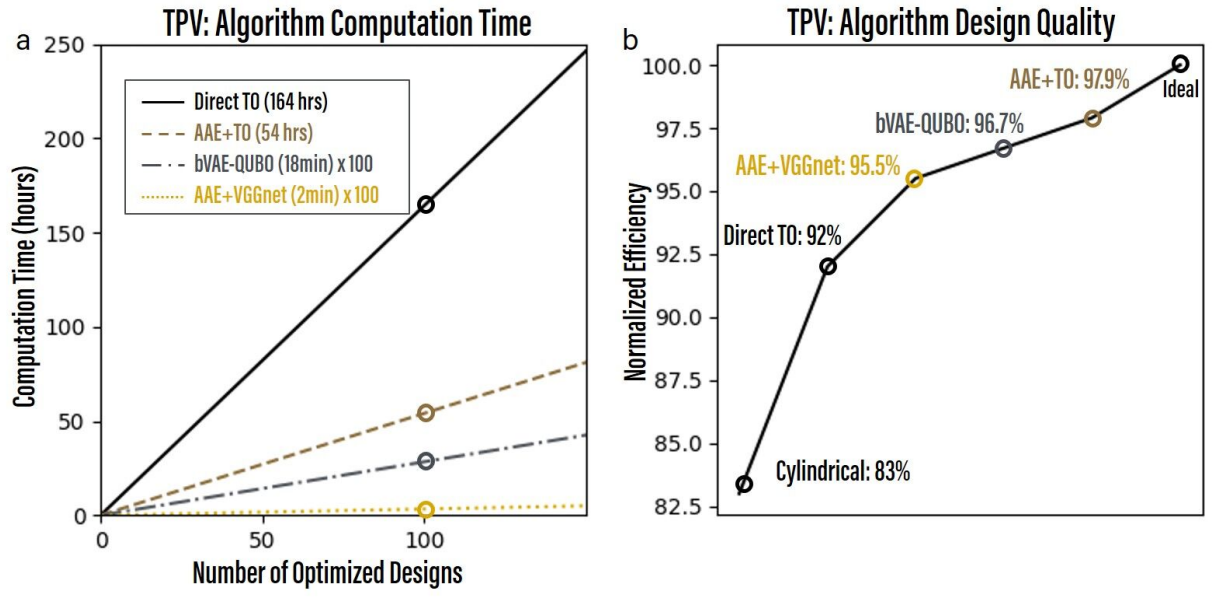
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0060481



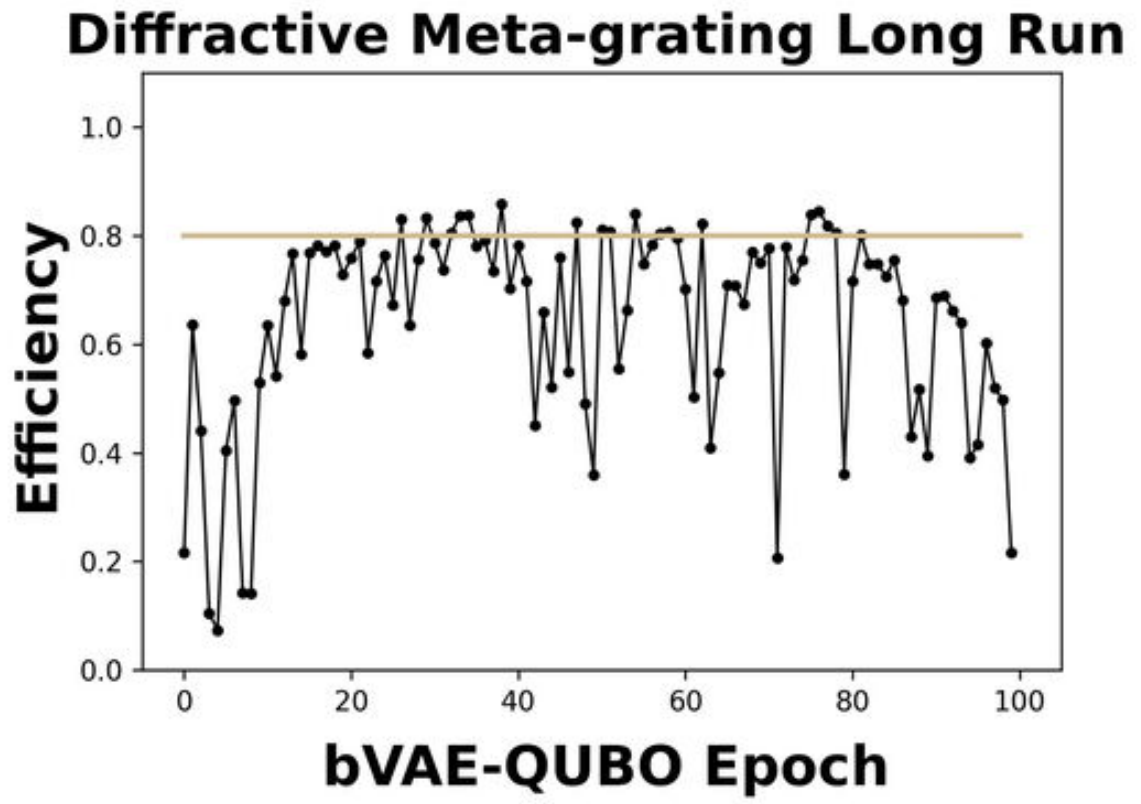
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0060481



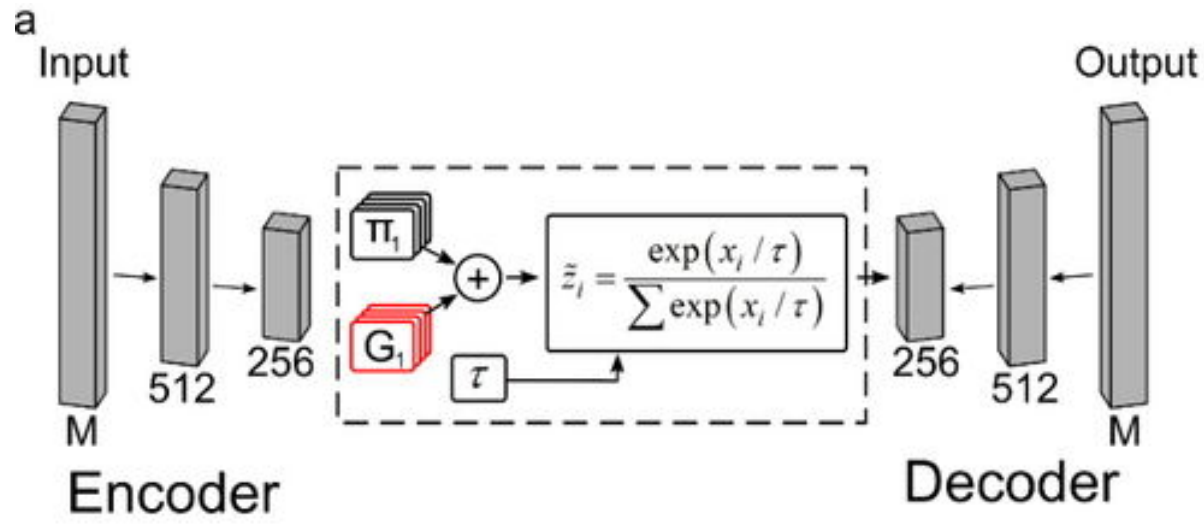
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0060481



This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0060481



This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0060481

