Learning incoherent light emission steering from metasurfaces using generative models

Prasad P. Iyer*, Saaketh Desai*, Sadhvikas Addamane, Remi Dingreville and Igal Brener
ppadma@sandia.gov, saadesa@sandia.gov,saddama@sandia.gov,rdingre@sandia.gov ibrener@sandia.gov
Center for Intergrated Nanotechnologies, Sandia National Laboratories
1515 Eubank Blvd SE, Albuquerque, NM, USA 87113
*These authors contributed equally

Abstract

Spatiotemporal control over incoherent light sources is critically important for applications such as displays, remote sensing, clean energy, and illumination. Incoherent light emission made up of randomized wavefronts is incompatible with known beam steering techniques that rely on coherent electromagnetic wave interference. The emerging field of tunable dielectric metasurfaces consisting of subwavelength arrays of optical nanoresonators has recently enabled active re-direction of incoherent light (photoluminescence, PL) emission. This was achieved by illuminating (pumping) the metasurface with a pump laser reflecting off a programmable spatial light modulator (SLM) with sawtooth grating patterns as input. Achieving efficient beam steering requires the generation of optimal pump patterns programmed into the SLM to maximize the PL emitted towards a given direction. Given the innumerable possibilities and the lack of a theoretical physical framework to guide the exploration of pump patterns, we use an active learning algorithm running a closed loop optical experiment with a generative model to explore and optimize novel pump patterns. We achieve up to an order of magnitude enhancement in the steering efficiency by using pump patterns that are generated by a variational auto-encoder, with minimal number of experiments. The results presented in this paper highlight the unique ability of generative models and active learning to dramatically improve steering efficiency by finding novel optical pump patterns that are beyond human intuition. Our combination of advanced machine learning techniques driving closed loop nanophotonic experiments might pave the way to derive the underlying physics of emergent light-matter phenomena.

1. Introduction

The ability to achieve dynamic spatiotemporal control over incoherent light emission has been an elusive target, restricting applications to those using mostly coherent lasers [1, 2]. Typical incoherent emitters like light emitting diodes (LEDs) and light bulbs emit over a wide angular field of view with randomized wavefronts making it impossible to steer using coherent phased-array concepts. Metasurface-based optics with sub-wavelength control over the amplitude, phase [3, 4], and polarization [5] of light raise the possibility of controlling incoherent emission or photoluminescence (PL) [6]. Recently, researchers have demonstrated static control of incoherent emission from light emitting metasurfaces with polarized, directional [7], and focused [8] emission. These results may enable new possibilities for low-cost integrated holographic augmented and virtual reality displays [9, 10], LED-based remote sensing [11] and communication modules for LiFi [12, 13], etc.

Dynamic steering of incoherent emission from within the metasurface was demonstrated by superimposing a refractive index pattern on the metasurface sample. This pattern is created by illuminating (pumping) the metasurface with a spatially patterned wide laser beam. This spatial pattern is created by reflecting the laser beam off a programmable spatial light modulator (SLM) that can be fed arbitrary spatial patterns of varying intensity [14]. The traditional sawtooth grating pattern [15] within the optical pump has been shown to dynamically re-direct the PL from the metasurface over a 70° field of view. In this paper, we develop a machine-learning framework to steer incoherent light emission into a particular direction by structuring the optical pump pattern. We propose a combination of a variational autoencoder (VAE) to create novel optical pump patterns with active learning to optimize the PL steered into a particular direction using minimal experiments.

Metasurfaces limit the emission directions of light to certain fixed angles (or momentum) due to their sub-wavelength arrangement (periodicity) of optical nanoresonators [3]. Reconfigurable metasurfaces [16, 17, 18, 11] with the ability to arbitrarily control the phase (ϕ) between 0 and 2π) can enforce arbitrary momentum ($\vec{k}_M = \frac{\partial \phi}{\partial x}$) on

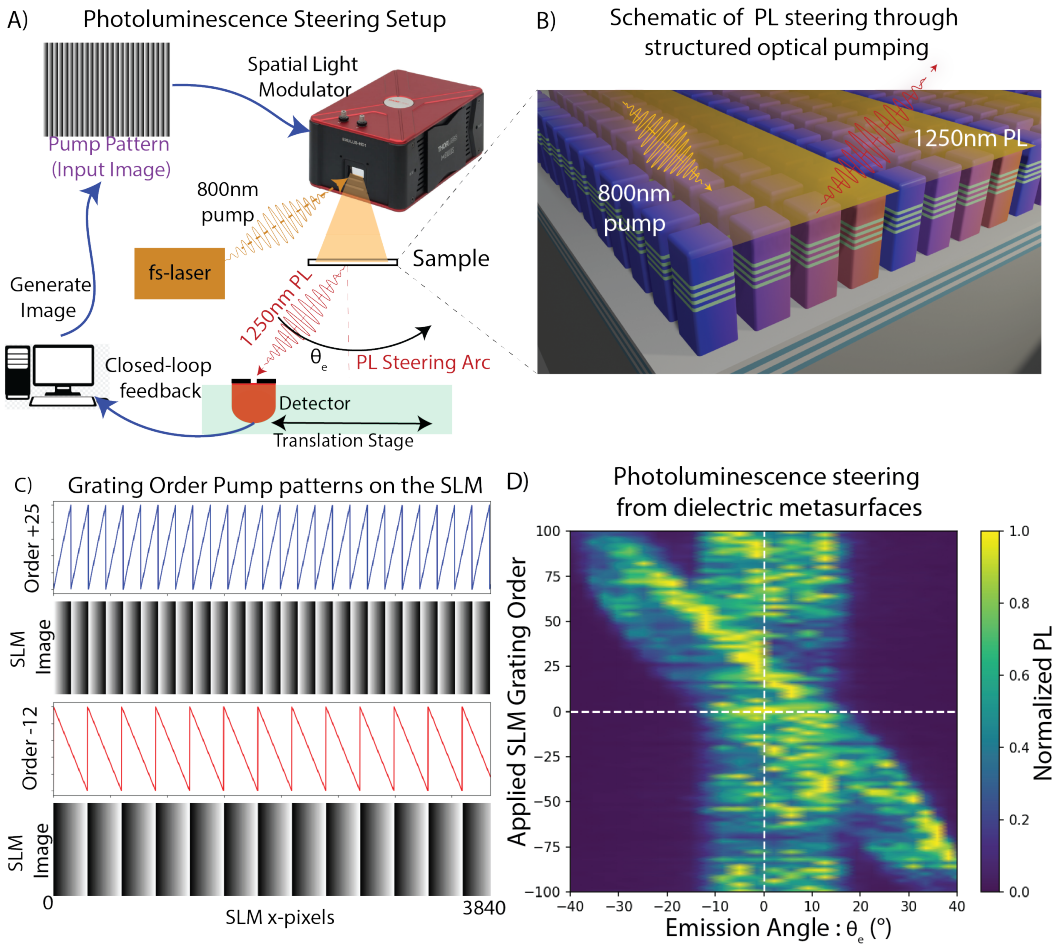

Figure 1. A) Photoluminescence steering setup: The schematic shows how the optical pump pattern on the 800nm pump laser is imaged onto the sample using the SLM. The PL from the sample is collected by a detector (on a translation stage) at different emission angles. The blue arrows indicates interface of the machine learning framework with the experiment where the image generated by the VAE forms the input for the SLM and the active learning leverages the closed-loop feedback of the detector. B) Schematic of the array of GaAs metasurface resonators with embedded InAs quantum dots (QDs) grown epitaxially on a distributed Bragg reflector (DBR) made of $Al_{0.3}Ga_{0.7}As$ and AlAs layers on a GaAs substrate. The yellow semi-transparent sheet on top of the resonators represents the spatial pattern of the optical pump inducing the emission from the InAs QDs and producing a spatially varying refractive index pattern in the resonators. C) Example grating order pump patterns: The two sub-panels show two example pump patterns (line profiles) and the corresponding SLM images used for steering PL. The pump patterns are characterized by the grating order (+25 and -12 for the top and bottom sub-panel respectively), where the grating order relates to the spatial periodicity of the pump pattern. Note that pattern has variations only along the x-axis of 3840 pixels while it is repeated along the Y-axis for 2160 pixels to generate the SLM image. D) Photoluminescence steering from the dielectric metasuface: The angular emission characteristics of the metasurface are normalized per optical pump pattern from the SLM.

the emitted light. This imposed momentum is used to steer the incoherent emission that is generated from within the metasurface. Here, we design and fabricate a reconfigurable semiconductor (GaAs) metasurface with embedded InAs quantum dot (QD) emitters. The refractive index change within the resonator caused by the absorption of the optical pump from the SLM results in an optical phase shift at the emission wavelength of the QDs [16, 17, 19]. The 800nm laser reflected from the SLM, carries the information of the spatial structure of the refractive index profile which gets translated into dynamic spatial momentum distributions on the sample. See Fig. 1A for a schematic of our experimental setup. Hence by spatially structuring the optical pump using a SLM (Fig. 1B) [20], we can impart additional momentum to the light emission. More details about the growth, fabrication and measurement setup used here are available in the supplementary information. We measure the far-field emission pattern from the metasurface to demonstrate continuous PL over an 80° field of view (limited by the setup) as a function of the spatial frequency of a saw tooth wave

(grating order) applied using the SLM (Fig. 1D).

The dynamic steering of PL from the GaAs metasurfaces based on the applied grating order pump pattern forms the forward physical model to estimate the far-field emission based on single periodic pattern in the SLM pump image. Theoretically, a grating saw-tooth should provide the correct pattern to cause angular steering. However, given the physics of light generation and several other physical processes that are interconnected in such complex light emitting metasurfaces, we have observed that this pattern is far from optimal. Using traditional optimization techniques leads to an intractable problem given the innumerable possibilities of pump patterns and the lack of a theoretical physical framework to guide this optimization. Even most advanced electromagnetic wave simulation tools tackle incoherent sources one at a time rendering them unsuitable for an array of incoherent source emitting from the metasurface. Thus a closed loop active learning process in combination with a image generator forms an ideal combination to enable us to rapidly find the optimal SLM image to maximize the PL steering.

2. Machine Learning Framework

The PL steering problem can be characterized within a machine learning framework as follows: A pump pattern (input image) is passed to the optical system through the SLM, which generates a PL (output) intensity along multiple emission angles. Beam steering is achieved by maximizing the output intensity at one specific emission angle, while minimizing the intensity everywhere else. The state-of-the art has explored a set of grating order (spatial frequency) based pump profiles (Fig. 1D) to identify that beam steering of PL is a possibility. However, there has been no attempt made to optimize this property or explore the response of the emitter to a vast combinatorial search space of $> 2^{2^{12000}}$ images of input pump patterns. The large search space prohibits brute force searches for an optimal input profile. Additionally, current human intuition and subject matter expertise are unable to reduce the dimensionality of the search space. This implies that we do not know a set of basis functions that can be used to generate new input pump pattern image. Hence, human intuition or subject-matter expertise cannot be used to explore more input profiles since the underlying relationship between input profile and output beam steering is unknown. Our problem is thus twofold: (1) We wish to find novel input pump patterns that maximize beam steering in a specific direction and (2) we wish to find these profiles in an efficient manner since a brute force search across pump-pattern space is unfeasible. To this end, we develop a machine learning workflow shown in Fig. 2 A.

To solve problem (1) we use a generative model (VAE [21] in this work) to (a) characterize input optical profiles

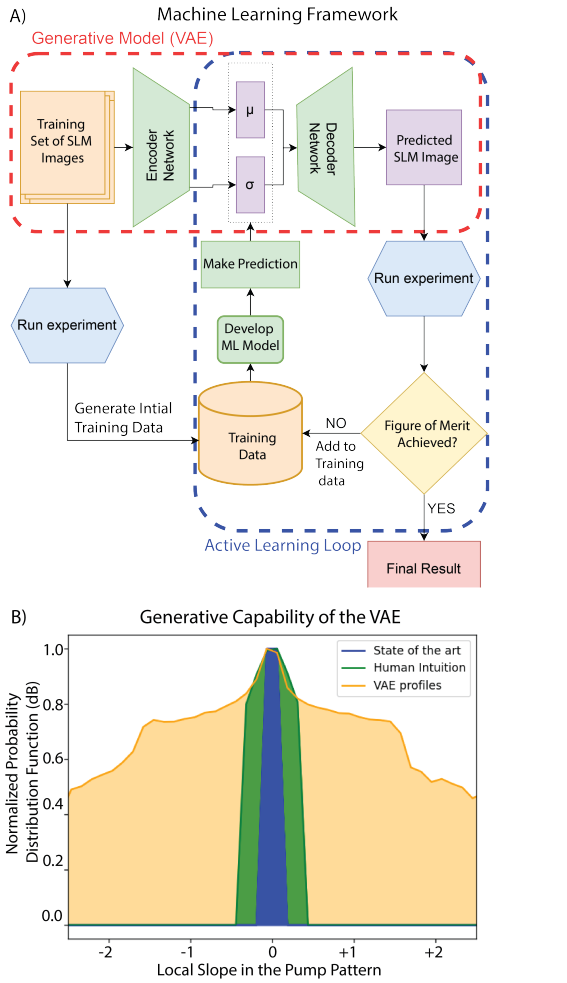

Figure 2. A) Machine learning flowchart highlighting the combination of the generative model (red dashed box) for image generation (VAE) used as the optical pump pattern in combination with a closed-loop active learning experiment (blue dashed box). B) The generative capability of the VAE illustrated by showing the distribution of patterns explored by the VAE (orange) compared to the distribution of state-of-the-art saw-tooth patterns (blue), and the distribution of human-intuition based patterns (green) used to hand-craft the training set for the VAE.

that can be expressed in a low dimensional, compact representation (latent dimension) unknown to subject matter experts and (b) sample the learnt latent dimension distribution to generate optical profiles beyond a limited, hand-crafted training set. While the VAE allows us to generate an infinite number of profiles by sampling the latent dimension, we need an efficient manner of exploring these profiles such that beam steering is optimized. To solve problem (2) above, we use an active learning scheme that is coupled to the latent dimension of the VAE. The active learning scheme finds the latent dimension values such that beam

steering is optimized, following which we use the trained VAE to decode the latent dimension and visualize the input pump pattern.

VAEs consist of an encoder that reduces an input pump pattern image x to a latent dimension z by learning a probability distribution $Q(z|x)$, and a decoder that expands the latent dimension z into an output pump pattern image x by learning the probability distribution $P(x|z)$. The latent dimension z is learnt as a Gaussian distribution, and sampling this distribution allows the VAE to generate new pump-pattern samples similar to those in the training set. Training a VAE involves maximizing the likelihood of generating pump pattern images similar to the training set X, formulated as minimizing the negative log-likelihood of $P(X)$ over the network parameters θ

$$P(X) = \prod_{x=X_i} \int P(x|z;\theta)P(z)dz. \quad (1)$$

This minimization can be formulated as minimizing the Evidence Lower Bound (ELBO) loss:

$$\text{ELBO} = E_{z\sim Q(z|X)} \log P(X|z) - D\left[Q(z|X)||P(z)\right], \quad (2)$$

where the first term is a reconstruction error that measures the ability of the decoder to generate images close to the input image, and the second term is the Kullback-Liebler (KL) divergence term that measures the similarity of latent space learnt by the VAE to a zero mean unit covariance Gaussian prior:

$$D[Q(z|X)||P(z)] = E_{z\sim Q}[\log Q(z|X) - \log P(z)]. \quad (3)$$

The reconstruction term here is formulated as a mean squared error between the input and the reconstructed image, assuming $P(X|z) = N(f(X|z;\theta), I)$. We leverage the innate properties of the VAE [21] to encode the input pump patterns into a low dimensional continuous latent space – dramatically reducing the search space by orders of magnitude. Additionally, the Gaussian distribution of the latent dimension allows us to generate novel pump pattern profiles beyond the set of training profiles without imposing constraints on the functional forms used. This generative property of VAEs derived from the continuous distribution of its latent dimensions are specifically useful for generating images in the physical sciences [22, 23, 24]. VAEs additionally allow for interpretability of learned representations [25, 26, 27], which is critical to developing an understanding of the underlying physics.

Here we optimize a one dimensional (1D) input pump pattern by using a VAE that generates an input pattern of 3840 pixels. The 2D SLM pump pattern input is generated by repeating the 1D-VAE generated pattern along the Y axis of the image with 2160 pixels. We limit ourselves to this repetition of the 1D curve along the Y-axis since our experimental setup (Fig. 1A) measures the PL variations only along the X-axis. Our VAE uses an encoder network with three standard feedforward hidden layers consisting of 1000, 1000, and a 100 units respectively. The decoder mirrors this architecture. The training set for the VAE consists of a variety of periodic 1D profiles generated by sampling the previously measured grating order profiles (Fig. 1D) and extending these to include polynomial Bezier curves defining the intensity profile within each period of the periodic pattern. We used 50,000 profiles as a part of our training set, normalizing profiles between 0 and 1. This VAE was implemented using Pytorch [28], using a mean squared error loss to quantify the reconstruction error, and using the Adam optimizer [29] with a learning rate of 0.001. For further details on the training, please refer to the Supplementary material. To illustrate the generative capability of the VAE in our context, Fig. 2 B shows the range of patterns explored by VAE. Each input pattern can be characterized by a set of local slopes (calculated after unwrapping the signal to account for periodicity). A distribution of these slopes across all the patterns in a dataset gives us a measure of the range of patterns explored. The blue region Fig. 2 B shows that the state-of-the art saw-tooth patterns only form a minor region of the design space. The green region demonstrates human intuition in generating the training set for the VAE. While this range is greater than the state of the art (blue), the VAE is clearly able to learn key features of patterns from the training set and use this to explore a significantly wider range of patterns (orange) compared to state-of-the art or human intuition.

We minimize the number of experiments required to find the optimal optical pattern using a closed loop active learning process (Fig. 2) [30, 31, 32, 33] We choose active learning in this work because the use of a Gaussian-process based active learning scheme allows us to handle experimental noise. The active learning protocol consists of a Gaussian process surrogate model that predicts the output beam steering, and an acquisition function that uses the Gaussian process surrogate model to identify which experiment to conduct next in order to achieve a target figure of merit, such as maximizing beam steering. This protocol begins with a small initial training set, following which subsequent experiments are identified using acquisition functions. Acquisition functions can be of various types depending on the need to balance exploration and exploitation, some examples are Maximum Expected Improvement (MEI) and Maximum Uncertainty (MU):

$$\text{MEI} : x^* = \text{argmax}\, E\left(\max(F(x) - F(x^*)), 0\right) \quad (4)$$

$$\text{MU} : x^* = \text{argmax}\, \sigma(F(x)), \quad (5)$$

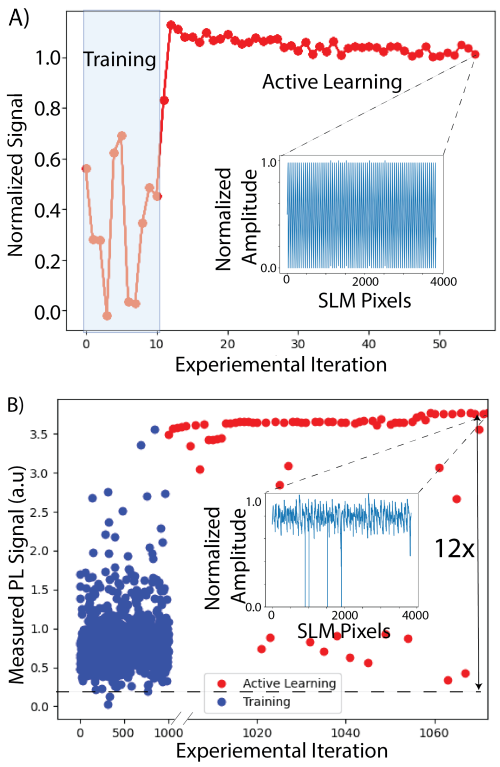

Figure 3. A) Normalized PL signal measured with the active learning optimizing over grating orders. Active learning is rediscovering a known result with 10% the number of experimental iterations required for a brute force search at an emission angle of -40°. The inset shows the normalized pump amplitude at the end of the active learning process, corresponding to a grating order of 80. B) Measured PL signal at the same emission angle with the active learning with pump patterns generated using the VAE. The inset here again shows the optimized pump pattern at the end of the active learning process with a 12× improvement in the signal w.r.t grating order based optimization in panel A represented by the black dashed line. The blue dots in this figure represent the training data while the final 100 experiments in red demonstrates how quickly the active learning optimizes the pump pattern.

where x^* is the next selected experiment, and x represents the set of possible experiments, and $F(x)$ is the collection of predicted outputs for all possible experiments x, as predicted by a surrogate model.

In this work we perform active learning on the latent dimension of the trained VAE, implementing the active learning using the Ax library [34]. We use the MEI acquisition function, where the initial training set is generated using Sobol sequences [35]. We use 1000 initial training points and query the active learning acquisition function for 100 subsequent experiments. The code for our work is available at `https://github.com/saakethdesai/ml-optics-WACV2023`.

3. Results and Discussion

We initially prove the capability of the active learning process to optimize for beam steering at a fixed emission angle (Fig. 3A). This is done to demonstrate that we can rediscover a previously known result (Fig. 1D), where a specific input pattern periodicity (grating order) optimizes beam steering at this angle. We find that the active learning scheme can identify this specific grating order using only 10% of the experiments required to obtain the result in Fig. 1D. Even though this is a trivial one-dimensional optimization, the result shows that the active learning is working as intended. Subsequently, we maximize the signal steered along the same angle using patterns generated from the VAE. We find that the wide range of patterns generated by the VAE enables us to find a pattern that shows a significant 12× improvement in the PL signal strength (Fig. 3B) compared to the saw tooth patterns measured previously in Ȧ closer look at the pump pattern that leads to this dramatic increase (inset, Fig. 3B) shows that the discovered pattern has a higher average value than the saw tooth patterns explored before. This implies that the active learning scheme learnt that a higher average pump power results in more PL emission from the metasurface, as would be expected from physics. Interestingly, this also means that our ML framework found a way to 'cheat' and obtain high output intensities at a specific angle. In this case, our framework obtained a high intensity at a specific angle not by steering light, but by increasing the overall photoluminescence across all emission angles. As shown in Fig. 3A, the active learning scheme was able to identify this pattern within the first 10 experiments, a feat that would not be possible with the current understanding of the underlying physics. The ability of the VAE and active learning to maximize beam steering at a specific angle effectively represents the ability of our framework to learn one of the basic energy conservation laws of physics.

Encouraged by this success, we now update the active learning figure of merit to optimize for the directivity of emission, defined as:

$$\text{Directivity} = \frac{\mathrm{f}(\theta_\mathrm{e})}{\sum_{\theta_\mathrm{e}} \mathrm{f}(\theta_\mathrm{e})}. \tag{6}$$

where $f(\theta_e)$ represents output signal at a specific steering angle θ_e

Using Eq. (6) as our target for the active learning results in a clear maximization of directivity, as shown in Fig. 4 A, bottom panel. The pump pattern that resulted in this maximum directivity shows interesting saw-tooth-like characteristics, but with perturbations to the pattern inside each period, see Fig. 4 A, top panel. This pattern shows a steering efficiency that is 3.4 times greater than the state-of-the-art saw-tooth patterns. A similar result is obtained across

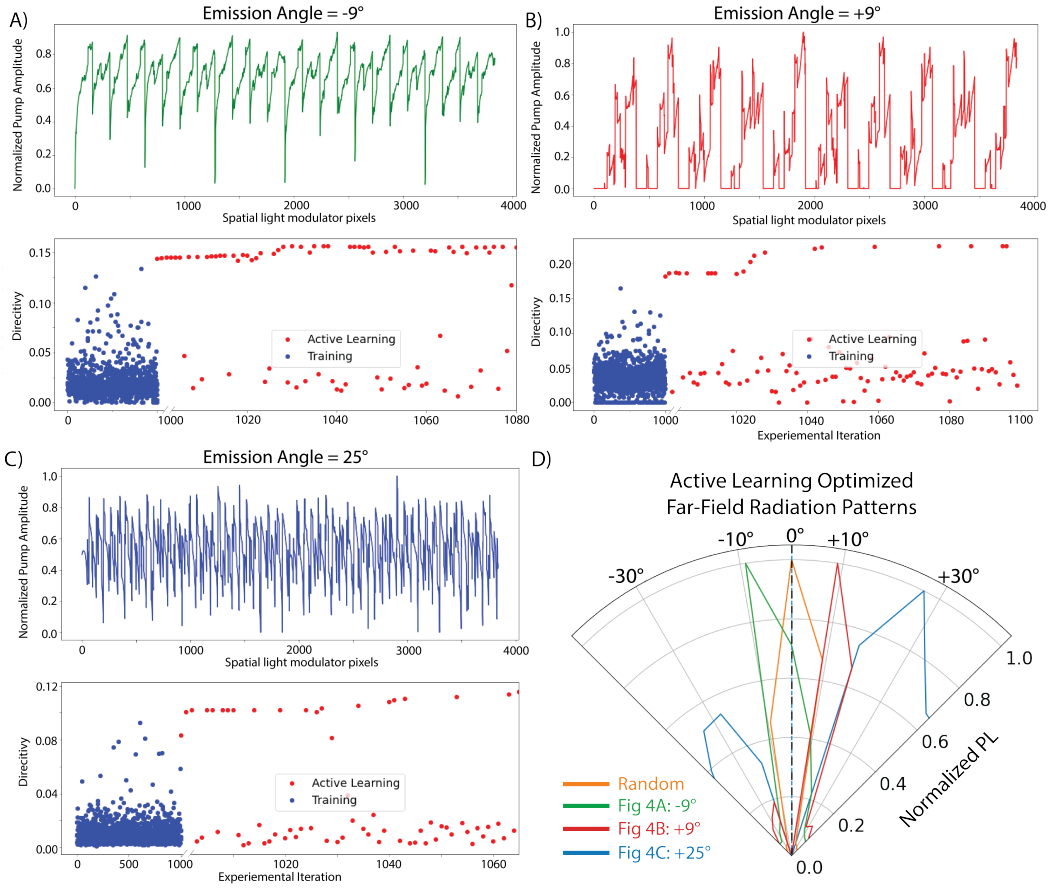

Figure 4. **Active Learning PL directivity**: A),B) & C) represent the active learning process to maximize the directivity of the far-field emission from the metasurface at -9°, +9° & +25° respectively. The top panel in each of these figures represent the optimized optical pump pattern generated by the VAE whose far-field emission patterns are plotted in panel D as indicated by the colored legends. As a control experiment, the orange curve in panel D shows the far-field emission of an unoptimized pump pattern generated by the VAE

multiple emission angles, as illustrated in Fig. 4 B and C. Our framework improves the steering efficiency by 5.6 and 8.7 times respectively w.r.t the grating order based steering demonstrated in Fig. 1D. Interestingly, the patterns that achieve this steering are different from the pattern in Fig. 4 A, with the pattern in 4 B showing a mixture of pulse and saw-tooth characteristics, while the pattern in 4 C is a very high-frequency pattern with sharply varying spatial intensity. These results are even more encouraging when we consider that the figure of merit (directivity) involves a summation of output intensities across multiple emission angles. This summation implies that the inherent experimental noise associated with capturing intensity at each emission angle is now amplified, and we find it encouraging that the active learning framework is able to optimize for specific patterns despite this noise. The final optimized far-field emission profiles are plotted in Fig. 4D, showcasing how the various pump patterns optimize for steering at various emission angles, highlighting the capability of this machine learning framework.

In conclusion, we have demonstrated nearly an order of magnitude improvement in the PL steering efficiency from metasurfaces using an active learning process combined with a variational auto-encoder to generate novel pump patterns that go beyond subject matter expert intuition. This is one of the first implementations of a machine-learning framework to actively improve a macroscopic observable such as PL steering without a guiding theoretical framework. We have demonstrated for the first time that non-trivial structuring of the the optical pump pattern can be lead to emergent light-matter interactions at the nano-scale. This demonstration proved that the active learning process working off a generative model can indeed learn and guide us towards new physical possibilities with minimal experimental iterations. Additionally, this architecture lays down a foundational template to explore novel physical processes driven by structured illumination while accounting for large experimental noise. This improvement in PL steering effi-

ciency opens up the possibility of using these low-form factor metasurface emitters in real-world applications including displays for augmented and virtual reality holographic displays. Moving forward, we are planning on leveraging the continuous distribution of the latent dimensional space of the VAE to develop a structure-property relationship between the PL steering efficiency and the pump pattern features. Ultimately, the results presented in this paper highlight the exceptional ability of computer vision based generative models to accelerate and advance our understanding of new physical processes.

Acknowledgements

This work was supported by the U.S. Department of Energy, Office of Basic Energy Sciences, Division of Materials Sciences and Engineering and performed, in part, at the Center for Integrated Nanotechnologies, an Office of Science User Facility operated for the U.S. Department of Energy (DOE) Office of Science. This work was also supported by the Beyond Fingerprinting Sandia Grand Challenge Laboratory Directed Research and Development (GC LDRD) program. Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

References

[1] Arun KS Iyengar, Hiroatsu Sugimoto, David B Smith, and Michael S Sacks. Dynamic in vitro quantification of bioprosthetic heart valve leaflet motion using structured light projection. *Annals of Biomedical Engineering*, 29(11):963–973, 2001.

[2] Maria Emília de Abreu Chaves, Angélica Rodrigues de Araújo, André Costa Cruz Piancastelli, and Marcos Pinotti. Effects of low-power light therapy on wound healing: Laser x led. *Anais Brasileiros de Dermatologia*, 89:616–623, 2014.

[3] Amir Arbabi, Yu Horie, Mahmood Bagheri, and Andrei Faraon. Dielectric metasurfaces for complete control of phase and polarization with subwavelength spatial resolution and high transmission. *Nature Nanotechnology*, 10(11):937–943, 2015.

[4] Tong Wu, Xueqian Zhang, Quan Xu, Eric Plum, Kaiji Chen, Yuehong Xu, Yongchang Lu, Huifang Zhang, Ziying Zhang, Xieyu Chen, et al. Dielectric metasurfaces for complete control of phase, amplitude, and polarization. *Advanced Optical Materials*, 10(1):2101223, 2022.

[5] Carl Pfeiffer and Anthony Grbic. Cascaded metasurfaces for complete phase and polarization control. *Applied Physics Letters*, 102(23):231116, 2013.

[6] Sheng Liu, Aleksandr Vaskin, Sadhvikas Addamane, Benjamin Leung, Miao-Chan Tsai, Yuanmu Yang, Polina P Vabishchevich, Gordon A Keeler, George Wang, Xiaowei He, et al. Light-emitting metasurfaces: simultaneous control of spontaneous emission and far-field radiation. *Nano Letters*, 18(11):6906–6914, 2018.

[7] Prasad P Iyer, Ryan A DeCrescent, Yahya Mohtashami, Guillaume Lheureux, Nikita A Butakov, Abdullah Alhassan, Claude Weisbuch, Shuji Nakamura, Steven P DenBaars, Jon Schuller, et al. Unidirectional luminescence from InGaN/GaN quantum-well metasurfaces. *Nature Photonics*, 14(9):543–548, 2020.

[8] Yahya Mohtashami, Ryan A DeCrescent, Larry K Heki, Prasad P Iyer, Nikita A Butakov, Matthew S Wong, Abdullah Alhassan, William J Mitchell, Shuji Nakamura, Steven P DenBaars, et al. Light-emitting metalenses and meta-axicons for focusing and beaming of spontaneous emission. *Nature Communications*, 12(1):1–7, 2021.

[9] Gyeongtae Kim, Seokwoo Kim, Hongyoon Kim, Jihae Lee, Trevon Badloe, and Junsuk Rho. Metasurface-empowered spectral and spatial light modulation for disruptive holographic displays. *Nanoscale*, 2022.

[10] Shoufeng Lan, Xueyue Zhang, Mohammad Taghinejad, Sean Rodrigues, Kyu-Tae Lee, Zhaocheng Liu, and Wenshan Cai. Metasurfaces for near-eye augmented reality. *ACS Photonics*, 6(4):864–870, 2019.

[11] Junghyun Park, Byung Gil Jeong, Sun Il Kim, Duhyun Lee, Jungwoo Kim, Changgyun Shin, Chang Bum Lee, Tatsuhiro Otsuka, Jisoo Kyoung, Sangwook Kim, et al. All-solid-state spatial light modulator with independent phase and amplitude control for three-dimensional lidar applications. *Nature Nanotechnology*, 16(1):69–76, 2021.

[12] Harald Haas. LiFi is a paradigm-shifting 5G technology. *Reviews in Physics*, 3:26–31, 2018.

[13] Shuji Nakamura. Current status of GaN-based solid-state lighting. *MRS Bulletin*, 34(2):101–107, 2009.

[14] Prasad P Iyer, Nicholas Karl, Sadhvikas Addamane, Sylvain D Gennaro, Michael B Sinclair, and Igal Brener. Sub picosecond steering of ultrafast incoherent emission from semiconductor metasurfaces. *arXiv preprint arXiv:2206.04645*, 2022.

[15] Christopher Palmer and Erwin G Loewen. Diffraction grating handbook. 2005.

[16] Prasad P Iyer, Nikita A Butakov, and Jon A Schuller. Reconfigurable semiconductor phased-array metasurfaces. *ACS Photonics*, 2(8):1077–1084, 2015.

[17] Maxim R Shcherbakov, Sheng Liu, Varvara V Zubyuk, Aleksandr Vaskin, Polina P Vabishchevich, Gordon Keeler, Thomas Pertsch, Tatyana V Dolgova, Isabelle Staude, Igal Brener, et al. Ultrafast all-optical tuning of direct-gap semiconductor metasurfaces. *Nature Communications*, 8(1):1–6, 2017.

[18] Shi-Qiang Li, Xuewu Xu, Rasna Maruthiyodan Veetil, Vytautas Valuckas, Ramón Paniagua-Domínguez, and Arseniy I Kuznetsov. Phase-only transmissive spatial light modulator based on tunable dielectric metasurface. *Science*, 364(6445):1087–1090, 2019.

[19] Nicholas Karl, Polina P Vabishchevich, Maxim R Shcherbakov, Sheng Liu, Michael B Sinclair, Gennady Shvets, and Igal Brener. Frequency conversion in a time-variant dielectric metasurface. *Nano Letters*, 20(10):7052–7058, 2020.

[20] Dmitriy Mikhaylov, Baifan Zhou, Thomas Kiedrowski, Ralf Mikut, and Andrés-Fabián Lasagni. High accuracy beam splitting using spatial light modulator combined with machine learning algorithms. *Optics and Lasers in Engineering*, 121:227–235, 2019.

[21] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[22] Pankaj Rajak, Aravind Krishnamoorthy, Aiichiro Nakano, Priya Vashishta, and Rajiv Kalia. Structural phase transitions in a MoWSe2 monolayer: Molecular dynamics simulations and variational autoencoder analysis. *Physical Review B*, 100(1):014108, 2019.

[23] Sergei V Kalinin, Ondrej Dyck, Stephen Jesse, and Maxim Ziatdinov. Exploring order parameters and dynamic processes in disordered systems via variational autoencoders. *Science Advances*, 7(17):eabd5084, 2021.

[24] Sebastian J Wetzel. Unsupervised learning of phase transitions: From principal component analysis to variational autoencoders. *Physical Review E*, 96(2):022140, 2017.

[25] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. 2016.

[26] Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. *arXiv preprint arXiv:1509.00519*, 2015.

[27] Gregory P Way and Casey S Greene. Extracting a biologically relevant latent space from cancer transcriptomes with variational autoencoders. In *Pacific Symposium on Biocomputing 2018: Proceedings of the Pacific Symposium*, pages 80–91. World Scientific, 2018.

[28] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32, 2019.

[29] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[30] Julia Ling, Maxwell Hutchinson, Erin Antono, Sean Paradiso, and Bryce Meredig. High-dimensional materials and process optimization using data-driven experimental design with well-calibrated uncertainty estimates. *Integrating Materials and Manufacturing Innovation*, 6(3):207–217, 2017.

[31] Turab Lookman, Prasanna V Balachandran, Dezhen Xue, and Ruihao Yuan. Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. *npj Computational Materials*, 5(1):1–17, 2019.

[32] Yongcheng Ding, José D Martín-Guerrero, Yujing Song, Rafael Magdalena-Benedicto, and Xi Chen. Active learning for the optimal design of multinomial classification in physics. *Physical Review Research*, 4(1):013213, 2022.

[33] David A Cohn, Zoubin Ghahramani, and Michael I Jordan. Active learning with statistical models. *Journal of Artificial Intelligence Research*, 4:129–145, 1996.

[34] Eytan Bakshy, Lili Dworkin, Brian Karrer, Konstantin Kashin, Benjamin Letham, Ashwin Murthy, and Shaun Singh. AE: A domain-agnostic platform for adaptive experimentation. In *Conference on Neural Information Processing Systems*, pages 1–8, 2018.

[35] Il'ya Meerovich Sobol'. On the distribution of points in a cube and the approximate evaluation of integrals. *Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki*, 7(4):784–802, 1967.