Generative Neural Networks for Kerr Combs

Janet Zhong[1], **Eran Lustig**[1], **Shiye Su**[1], **Louise Schul**[1],
Jamison Sloan[1], **Congyue Deng**[2], **Jelena Vuckovic**[1], **Shanhui Fan**[1]
[1]Stanford University [2]MIT

Abstract

Light in integrated micro-resonators can lead to highly nonlinear processes and complex optical spectra including frequency combs, instabilities, solitons, and more. Due to the inherent nonlinearity, bistability, and hysteresis of the system, the mapping of desired optical spectral properties to input physical parameters can be very difficult. Existing approaches for retrieving this mapping use traditional optimization or non-generative neural networks, which struggle with multi-solution landscapes. To solve these issues, we frame this inverse design problem as a generative distribution-learning task for the first time. Using conditional variational autoencoders and flow matching models, we generate input parameters and their spectra for high-bandwidth steady-state solitons trained on Lugiato-Lefever equation simulations. Our approach can be applied to experimental data with little modification. Actual experimental conditions for Kerr combs often deviate from theoretical models, making data-driven machine learning approaches particularly promising for applications in spectroscopy, optical communications, and nonlinear optics research.

1 Introduction

Kerr micro-resonators can exhibit nonlinear optical phenomena ranging from frequency combs, instabilities, dissipative solitons and more [1–8]. These states are used in spectroscopy, optical frequency metrology, and telecommunications, and are also fundamentally interesting for their novel properties such as soliton crystals [9, 10], quantum dynamics [11, 12] and squeezed light [11, 13–15]. However, the bistability and path-dependent dynamics of Kerr combs can create multi-valued inverse problems that challenge theoretical and numerical studies as well as traditional inverse design methods [16–18]. We demonstrate that generative neural networks can offer a solution to the Kerr comb inverse design problem by learning a probabilistic distribution over input parameters.

We train a conditional variational autoencoder and a conditional flow matching model to generate input parameters (single pump power, dispersion, and laser detuning) and the corresponding optical spectra for high-bandwidth steady-state solitons. Broadband solitons are desirable because many applications rely on access to a wide range of modes [19]. Although trained on simulated data, our method is designed to be adaptable to experimental datasets that can deviate from the known theory. While for simplicity we train a model for the single-pump case [20–25], our method is translatable to multi-pump cases [26–29] which have more realistically tunable experimental parameters and richer optical spectral regimes. This generative approach offers promising applications in spectroscopy, communications, and nonlinear optics research [1, 6, 19].

1.1 Kerr combs and Lugiato-Lefever Equation

We use Lugiato–Lefever equation [1–3] (LLE) simulations which can model Kerr combs to create a training dataset. Examples of optical states in Kerr combs modeled by the LLE can be seen in Fig. 1(b), which represents the field intensity in a microring resonator for various optical states. In

Machine Learning and the Physical Sciences Workshop, NeurIPS 2025.

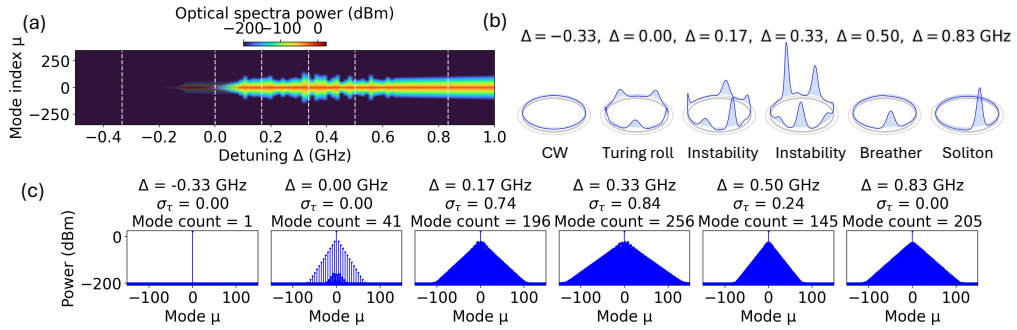

Figure 1: (a) Example of a detuning sweep of an input CW pump (b) The field amplitude in the microring at the detuning values marked by dashed lines in (a). (c) Optical spectra at the same detuning values marked in (a) with corresponding temporal stability measure σ_τ and mode count (a bandwidth measure).

Fig. 1(b), we depict the continuous-wave (CW) states, Turing roll frequency combs, instabilities, breathers and solitons. These optical states are categorized by their characteristic field profiles (Fig. 1(b)) and by their transient properties, such as whether they are steady states.

The LLE is a driven-dissipative nonlinear partial differential equation given by [1–3]:

$$\frac{\partial A}{\partial \tau} + \sum_n (-i)^n \frac{D_n}{n!} \frac{\partial^n A}{\partial \phi^n} - ig|A|^2 A + \left(\frac{\kappa}{2} + i\Delta\right) A = \sqrt{\eta\kappa} s. \tag{1}$$

Here $A(\phi, \tau)$ is the complex electric field inside the resonator, $\phi \in [0, 2\pi)$ is the angular coordinate on the ring, and $\tau = \frac{\kappa}{2} t$ is the normalized time (t is the physical time) where κ is a loss term. The terms D_n describe dispersion coefficients which represent frequency-dependent light propagation in the resonator and g is the nonlinear Kerr coefficient, s is the input field amplitude with $|s|^2 = P/(\hbar\omega_0)$ where P is the pump power, ω_0 is the pump frequency and $\hbar$ is the reduced Planck constant. Δ is the detuning which is the frequency mismatch of the laser pump with the natural frequency of the resonator.

In Fig. 1(a), we show the optical spectra as a function of pump detuning. The optical spectra is calculated with the LLE. The detuning Δ is swept from blue to red. The order of transitions as the detuning is swept, is: CW $\rightarrow$ Turing rolls $\rightarrow$ instabilities $\rightarrow$ solitons $\rightarrow$ CW. In Fig. 1(c), we also show the optical spectra of these characteristic phases.

1.2 Inverse design using generative modeling

Inverse design approaches in photonics: Traditional inverse design or optimization approaches include gradient-based optimization or heuristic algorithms [30–32]. Neural network-assisted methods [32–35] include optimizing on surrogate forward model outputs [36, 37], optimizing on latent spaces of unconditional generative models [38], conditional generative models [39, 40] and more [32–35]. We use conditional generative models as they learn the inverse mapping more directly [41] without requiring an extra optimization step.

Kerr comb inverse design: Previous works on Kerr comb inverse design have used genetic algorithms [16–18] and non-generative neural networks for predicting soliton regimes [42], dispersion engineering [16, 43] or as a surrogate forward model [44, 45]. These methods can struggle with the one-to-many inverse mapping inherent in Kerr comb dynamics, where multiple parameter configurations can produce similar states. Generative models offer a promising alternative as they can naturally represent the multi-valued inverse mapping.

Generative model architectures for inverse design: Variational autoencoders (VAEs) [46] have been a popular choice for inverse design [38, 39, 47–49] but can produce blurry interpolations [50]. Diffusion models [51–53] have been shown to outperform VAEs [54–56] and can surpass training distribution performance [57, 58]. Meanwhile, flow matching models [59–61] have recently gained popularity in various scientific applications [62–64] as they can have more efficient sampling while

achieving comparable or better performance to diffusion models. In this work, we focus on VAEs for their computational efficiency and flow matching models for their potential generation quality.

Variational autoencoders: VAEs [46] are generative models that learn probabilistic latent representations by combining an encoder network $q(\mathbf{z}|\mathbf{x})$ that approximates the posterior distribution of latent variables $\mathbf{z}$ given data $\mathbf{x}$, and a decoder network $p(\mathbf{x}|\mathbf{z})$ that reconstructs the data from the latent representation. Both q and p can be conditioned on additional information y, enabling controlled generation. VAEs are trained by maximizing the evidence lower bound, which balances reconstruction accuracy with regularization of the latent space:

$$\mathcal{L}_{\text{VAE}} = \mathbb{E}_{\mathbf{z}\sim q(\mathbf{z}|\mathbf{x},y)} \left[-\log p(\mathbf{x}|\mathbf{z},y) + D_{\text{KL}}(q(\mathbf{z}|\mathbf{x},y) \| p(\mathbf{z}|y)) \right], \tag{2}$$

where D_{KL} denotes the Kullback-Leibler (KL) divergence between distributions.

Flow matching models: Flow matching models [59, 65] are generative models that learn to transport samples from a prior distribution to a data distribution via a time-dependent vector field $v_\theta(\mathbf{x}, t)$ parameterized by a neural network with parameters θ. The model defines continuous normalizing flows through ordinary differential equations and is trained in a simulation-free manner by regressing v_θ to a target velocity field u:

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{t\sim\mathcal{U}[0,1],\mathbf{x}_1\sim p_{\text{data}},\mathbf{x}_t\sim p_t(\mathbf{x}_t|\mathbf{x}_1),y} \left\| v_\theta(\mathbf{x}_t, t, y) - u(\mathbf{x}_t|\mathbf{x}_1, y) \right\|^2, \tag{3}$$

where $t \in [0, 1]$ is the flow time, $\mathbf{x}_1$ denotes data samples, $\mathbf{x}_t$ represents intermediate states along the flow path, and y is the conditioning information. In this work, we use a Gaussian prior and the conditional optimal transport path $p_t(\mathbf{x}_t|\mathbf{x}_1) = \mathcal{N}(t\mathbf{x}_1, (1-t)^2\mathbf{I})$ [59], where $\mathbf{I}$ is the identity matrix. We condition the flow model on y to sample from conditional distributions with particular desirable properties.

2 Dataset and Methods

Dataset construction: We generate a dataset using the open-source LLE solver PyCORe [66]. Other alternatives include PyLLE [67] and PyGLLE [68]. We create detuning sweeps from –500 MHz to 1 GHz in 150 steps, with randomized pump powers between 0.1 and 0.2 W. We vary the dispersion coefficients within the ranges: $D_2 \in$ [20, 40] MHz, $D_3 \in$ [-20, 20] kHz, $D_4 \in$ [-2, 2] kHz, and $D_5 \in$ [-500, 500] Hz. We calculate a temporal stability measure σ_τ at each detuning point, from the normalized temporal variance of the transient dynamics found in the LLE (see Appendix for more details), where $\sigma_\tau \in [0, 1]$ with $\sigma_\tau \approx 0$ indicating steady states and larger values indicating increasing instability. A similar measure could be found experimentally from the RF spectra. We calculate bandwidth via mode count, defined as the number of modes in the optical spectra exceeding the noise floor (10th percentile) by 10 dB. We vary the dispersion which influences the nonlinear states and can be readily engineered [20–25]. We include dispersion terms up to quintic order, which is a realistic cutoff with respect to current scientific studies [25].

Objective: Our objective is to find the input parameters $P, D_2, D_3, D_4, D_5, \Delta$ that maximize the steady-state soliton bandwidth. High bandwidth is often desirable in applications such as spectroscopy, communication, and ultrafast light generation [69]. We look for regimes with both low σ_τ (steady-states) and high mode count (broad bandwidth). Both conditions are necessary for broadband solitons as high mode count alone would include modulation instabilities (which have high σ_τ) while low σ_τ alone would include Turing rolls (which have lower mode count than solitons). For example, in Fig. 1(c), the optical spectra in the 3rd and 4th subpanels from the left show instabilities, and the 6th subpanel shows a soliton. While the instabilities and the soliton spectra are similarly broadband, they differ in their temporal stability measure σ_τ.

Conditional generative modeling task: Our inverse design objective of steady-state soliton bandwidths can be posed as a conditional generative modeling problem:

$$(\sigma_\tau, \text{mode count}) \rightarrow (P, D_2, D_3, D_4, D_5, \Delta, \text{optical spectra}) \tag{4}$$

Here, σ_τ and mode count (the desired spectral properties) are the conditioning variables. The model then generates both the physical system parameters $P, D_2, D_3, D_4, D_5, \Delta$ and the optical spectra at those values (which has shape 700×1). We pick the ranges $\sigma_\tau < 0.1$ (at or close to steady-state) and mode count $\in [300, 400]$ (upper range of soliton bandwidths in our dataset). These ranges were selected based on the distribution in our dataset to target high-performing regimes without extrapolating too far beyond training data distribution which may yield unphysical results.

Model training: Our models are trained on a 9:1 train-test split of 100 detuning sweeps (150 detuning steps each), yielding 15,000 optical spectra with 700 modes each. For training, physical parameters (P, D_2, D_3, D_4, D_5, Δ) and conditioning variables (σ_τ, mode count) are globally standardized to zero mean and unit variance, while optical spectra are normalized per detuning sweep. The latter is chosen because we are more interested in the optical phase transitions within a detuning sweep as opposed to how the absolute intensity varies across the dataset. For the VAE, we use a 4-layer encoder-decoder architecture. The encoder takes 700-mode spectra concatenated with 6 parameters (P, D_2, D_3, D_4, D_5, Δ) and maps to a latent space of dimension 32, while the decoder reconstructs both spectrum and parameters. The VAE loss combines spectrum reconstruction (weight 1), parameter reconstruction (weight 1), and KL divergence (weight β). Both encoder and decoder are conditioned on mode count and temporal stability σ_τ. For flow matching, we use a 2D U-Net architecture, which empirically outperformed 1D architectures. The 700 spectral modes are zero-padded and concatenated with 6 parameters (768 total), then reshaped into 24×32 tensors for U-Net processing. We condition on mode count and σ_τ by concatenating them with the time embeddings. For the parameter and spectral loss in both VAE and flow models, we tested standard mean-squared error (MSE) loss, L1 loss and a peak-preserving L1 loss (PLL1). More details on model architectures are in the Appendix.

Generation evaluation methods: To our knowledge, no prior work addresses generative models for Kerr combs and previous studies on inverse design in Kerr combs focus on slightly different problem formulations [16–18]. We propose several evaluation metrics focusing on distribution matching, conditioning quality, and LLE consistency. We evaluate on a fixed test set subsample of 1000 steady-state spectra detuning slices ($\sigma_\tau < 0.1$) rather than the entire test set as steady-state regimes have less variability and are more relevant for applications. For each of the test subsample entries, we generate one sample conditioned on its mode count and σ_τ value. We use Wasserstein distance [55, 70] to quantify distributional differences between generated and test sets, computed on normalized spectra (flattened across samples and 700 modes) and normalized parameters (each parameter flattened across samples, then averaged across the 6 parameters: P, D_2, D_3, D_4, D_5, Δ). The Wasserstein distance is the 'cost' to transform one distribution into another. For conditioning quality, we calculate the actual mode count of unnormalized generated spectra and compare it to the conditioned target mode count and take the L1 error ('Mode Count Error' in Table. 1). For LLE consistency, we would ideally reverify the spectra using the LLE of the generated parameters and compare the error. However, this is computationally costly and we instead employ a nearest-neighbor approach. For each generated parameter set, we find its nearest neighbor (by L2 distance in normalized parameter space) from the full test set and compare their spectra using MSE, Pearson correlation, and structural similarity index (SSIM).

3 Results

Model	W. Spectra	W. Params Avg	Mode Count Error	LLE MSE	LLE Pearson	LLE SSIM
VAE-MSE-β0.5	0.019	0.082	9.608	0.130	0.933	0.506
VAE-L1-β0.5	0.020	0.167	7.950	0.078	0.941	0.887
VAE-PLL1-β0.5	0.022	0.111	14.391	0.147	0.890	0.466
VAE-MSE-β5.0	0.025	0.232	8.240	0.176	0.897	0.739
VAE-L1-β5.0	0.026	0.287	**6.320**	0.071	**0.952**	**0.917**
VAE-PLL1-β5.0	0.025	0.312	15.928	0.110	0.879	0.646
Flow-MSE	0.028	0.071	15.180	0.071	0.935	0.251
Flow-L1	**0.012**	0.057	15.680	**0.064**	0.944	0.391
Flow-PLL1	0.025	**0.051**	16.210	0.067	0.938	0.282

Table 1: Comparison of evaluation metrics for various models. Lower is better (↓) except Pearson and SSIM where higher is better (↑).

Evaluation metrics for a selection of the best-performing conditional VAE and conditional flow matching models are given in Table 1. Both approaches generated physically reasonable spectra, as illustrated in Fig. 2(a-b). Within the VAEs, smaller β has better distribution learning properties (Wasserstein metrics), but larger β had better nearest-neighbour LLE metrics. We note that spectra generated by flow models appeared slightly noisier than VAEs and the small noise did not disappear despite hyperparameter tuning. VAEs also have the tendency to produce smoother spectra, which may or may not be physically accurate. Metrics like SSIM and Mode Count Error strongly penalize

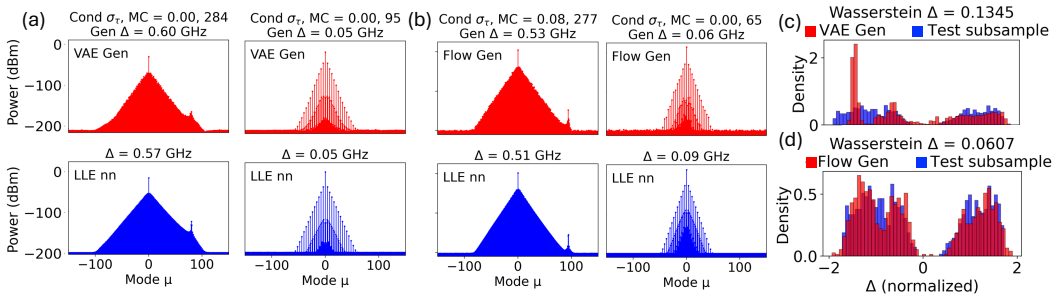

Figure 2: (a-b) Generated examples for VAE and flow models (red) with LLE nearest-neighbour comparison (blue). (c-d) Detuning distribution for VAE and flow generations (red) compared to test subsample of steady-states (blue). 'Cond σ_τ, MC' are the conditional variables σ_τ and mode count, 'Gen Δ' is the unnormalized generated detuning value. Other parameters are omitted from the title for brevity. The models are VAE-MSEβ0.5 and Flow-PLL1.

noise, and the flow models perform worse in these. We note that the relative performance of the VAE and flow models in these metrics may be altered through post-processing (e.g., smoothing flow-generated spectra) without retraining the models. Flow models have slightly better distribution learning properties in the parameter reconstruction than VAEs. In Fig 2(c-d), we show parameter density histograms for detuning for generated evaluation spectra (red) and the test set subsample (blue) for both a VAE and a flow model. For steady-states, the detuning distribution should show a greater density at the lower and higher end of the detuning range (for CW, Turing rolls and solitons), and both VAE and flow model show this but the flow model generations match the test set distribution more closely. The VAE models are much more computationally efficient than flow models for both training and sampling. Table 1 highlights the importance of employing diverse evaluation metrics, as different properties (spectral smoothness, distribution learning, computational efficiency) can have mixed trade offs, and one may place more importance on specific metrics depending on the application at hand.

4 Conclusion

In summary, we demonstrate that generative neural networks can effectively be used for multi-objective Kerr comb inverse design. While we use simulation data as a proof of concept, our approach is designed for direct application to experimental datasets and can extend beyond the single-pump regime to multi-pump parameter optimization [26–29]. Application to experimental data is particularly compelling because the LLE often deviates from fabricated devices. Furthermore, experimental data acquisition occurs on faster timescales than LLE simulations [3, 7, 8], potentially alleviating data scarcity concerns that often limit machine learning approaches. Future extensions may include training on full detuning sweeps rather than individual slices (which may incorporate more information on path history [71]), incorporating transient information such as RF spectra [72], and more systematic comparisons with traditional optimization methods or sensitivity analysis. The interpretable latent space of some generative models [73–76] and simulation-based inference or active learning approaches [77] may also be of interest. Our work establishes both a methodological framework and evaluation metrics for this new generative modeling task. Given recent advances in generative models for scientific applications [41, 57, 58, 62, 63], we believe that generative neural networks are an exciting area of potential for Kerr comb inverse design and nonlinear photonics research and applications.

Acknowledgements

We acknowledge Stanford Research Computing Center Sherlock and Marlowe cluster for computing resources. JZ acknowledges Matt Beutel, Olivia Long and Chenkai Mao for helpful discussions.

Appendix

4.1 Lugiato-Lefever equation details

Initialize:
$A^{\text{seed}} := \text{seed}(P, \Delta_0) + \text{noise}$
for $i = 0$ **to** $N_d - 1$ **do**
  $\Delta_i := \texttt{detuning}[i]$
  $A^{(i)}(0) := \begin{cases} A^{\text{seed}}, & i = 0 \\ A^{(i-1)}(N_\tau), & i > 0 \end{cases}$
  for $k = 0$ **to** $N_\tau - 1$ **do**
    $A^{(i)}(k+1) := \texttt{split\_step\_Fourier}(A^{(i)}(k), \Delta_i, D_2, D_3, D_4, D_5, g, \kappa)$
    $\texttt{time\_evolution\_step}\,[i, k, :] := A^{(i)}(k+1)$
  end
  $\texttt{detuning\_snapshot}[i, :] := A^{(i)}(N_\tau)$
end

Algorithm 1: Lugiato-Lefever equation (LLE) detuning (Δ) sweep pseudocode.

We show the LLE pseudocode in Algorithm 1. Here, `detuning` (shape N_d) is the list of detuning values to sweep through, $A^{(i)}(k)$ (shape N_μ) is the electric field at time step k and detuning step i, `time_evolution_step` (shape N_d, N_τ, N_μ) stores the field evolution across all detuning steps and time, and `detuning_snapshot` (shape N_d, N_μ) saves only the state at time step N_τ for each detuning step. The detuning sweep in Fig. 1(a) shows an example `detuning_snapshot` array. The initial seed is the steady-state CW solution at Δ_0 with added noise representing quantum vacuum fluctuations which seeds modulation instability and enables soliton formation. The algorithm evolves the optical field through time steps at each detuning using the split-step Fourier method and uses the final state for some specified time span as the initial condition for the next detuning value. The time evolution at selected detuning values is shown in Fig. 3, which illustrates the temporal characteristics of the Turing rolls, instabilities, breather, and soliton states seen in Fig. 1(b, c). In Fig. 3, σ_τ is the normalized variance across time steps averaged over all spatial positions. The time steps are down-sampled for plotting but must be sufficiently fine in simulations to capture accurate physics, which is a computational bottleneck. Each detuning sweep in our dataset took 4-5 minutes to generate.

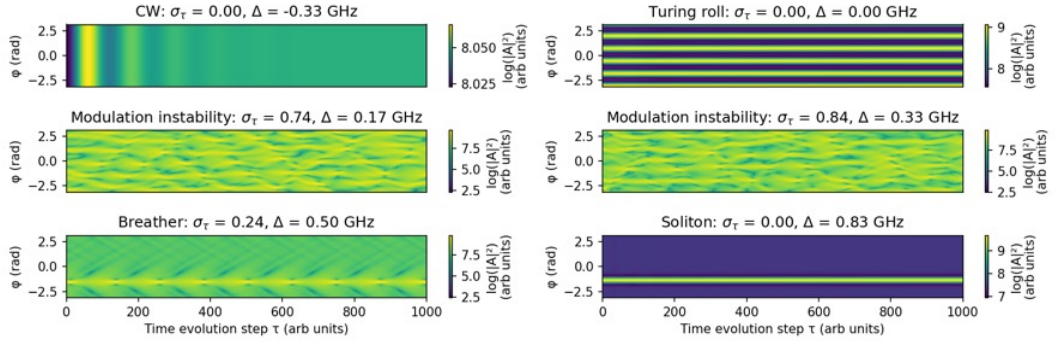

Figure 3: Time-evolution of the detuning snapshots in Fig. 1, showing how σ_τ (the normalized temporal variance, averaged across spatial domain ϕ) can capture the instabilities of the time-evolution.

Other simulation parameters: For the resonator, we used intrinsic loss rate $\kappa_0/2\pi = 70$MHz, external coupling rate $\kappa_{\text{ex}}/2\pi = 70$ MHz (critically coupled), free spectral range FSR $= 200$ GHz, pump frequency $\omega_0/2\pi = 192$ THz, linear refractive index $n_0 = 1.9$, Kerr nonlinearity coefficient $n_2 = 2.0 \times 10^{-19}$ m^2/W, waveguide cross-section of 1.5 μm width and 0.85 μm height. In Eq. (1), $\kappa = \kappa_0 + \kappa_{\text{ex}}$ and $\eta = \kappa_{\text{ex}}/\kappa$. In the data set, we used a scan time of 1μs across the whole detuning sweep with a normalized time step $d\tau = 0.00015$ for the split-step Fourier method.

4.2 Model hyperparameters

VAE models: All VAE models use a 4-layer encoder/decoder architecture with hidden dimensions [1024, 512, 256, 128], latent dimension of 32, Adam optimizer with learning rate of 0.001, and are trained for 4000 epochs using full-batch updates.

Flow models: The flow models use a U-Net architecture with 64 base channels, channel multipliers [1, 2, 3], 2 residual blocks per resolution, single-resolution attention at 4×4, trained with batch size 32 for 2000 epochs using AdamW optimizer with learning rate 0.0001, and evaluated with classifier-free guidance scale 0.2 using Heun's solver with 50 steps. Our implementation is based on Ref. [65].

Loss functions: We weight spectrum and parameter loss equally (by 1) in all models. We use MSE loss, L1 loss and a peak-preserving loss PLL1 in Table. 1. The PLL1 loss is L1 loss for spectra and parameters, but the spectral modes with larger intensity than 10% of the maximum spectral intensity have an additional $30\times$ weight.

References

[1] L. A. Lugiato, F. Prati, M. L. Gorodetsky, and T. J. Kippenberg. From the Lugiato-Lefever equation to microresonator-based soliton Kerr frequency combs. *Philosophical Transactions of the Royal Society of London Series A*, 376(2135):20180113, December 2018.

[2] L. A. Lugiato and R. Lefever. Spatial dissipative structures in passive optical systems. *Phys. Rev. Lett.*, 58:2209–2211, May 1987.

[3] Tobias J. Kippenberg, Alexander L. Gaeta, Michal Lipson, and Michael L. Gorodetsky. Dissipative kerr solitons in optical microresonators. *Science*, 361(6402):eaan8083, 2018.

[4] T. Herr, K. Hartinger, J. Riemensberger, C. Y. Wang, E. Gavartin, R. Holzwarth, M. L. Gorodetsky, and T. J. Kippenberg. Universal formation dynamics and noise of Kerr-frequency combs in microresonators. *Nature Photonics*, 6(7):480–487, July 2012.

[5] Wieslaw Krolikowski, Ole Bang, Jens Juul Rasmussen, and John Wyller. Modulational instability in nonlocal nonlinear kerr media. *Phys. Rev. E*, 64:016612, Jun 2001.

[6] Xucheng Zhang, Chunxue Wang, Zhibo Cheng, Congyu Hu, Xingchen Ji, and Yikai Su. Advances in resonator-based kerr frequency combs with high conversion efficiencies. *npj Nanophotonics*, 1(1), July 2024.

[7] Alessia Pasquazi, Marco Peccianti, Luca Razzari, David J. Moss, Stéphane Coen, Miro Erkintalo, Yanne K. Chembo, Tobias Hansson, Stefan Wabnitz, Pascal Del'Haye, Xiaoxiao Xue, Andrew M. Weiner, and Roberto Morandotti. Micro-combs: A novel generation of optical sources. *Physics Reports*, 729:1–81, January 2018.

[8] Bill Corcoran, Arnan Mitchell, Roberto Morandotti, Leif K. Oxenløwe, and David J. Moss. Optical microcombs for ultrahigh-bandwidth communications. *Nature Photonics*, 19(5):451–462, May 2025.

[9] Daniel C. Cole, Erin S. Lamb, Pascal Del'Haye, Scott A. Diddams, and Scott B. Papp. Soliton crystals in Kerr resonators. *Nature Photonics*, 11(10):671–676, October 2017.

[10] Maxim Karpov, Martin H. P. Pfeiffer, Hairun Guo, Wenle Weng, Junqiu Liu, and Tobias J. Kippenberg. Dynamics of soliton crystals in optical microresonators. *Nature Physics*, 15(10):1071–1077, October 2019.

[11] Yanne K. Chembo. Quantum dynamics of kerr optical frequency combs below and above threshold: Spontaneous four-wave mixing, entanglement, and squeezed states of light. *Phys. Rev. A*, 93:033820, Mar 2016.

[12] Melissa A. Guidry, Daniil M. Lukin, Ki Youl Yang, Rahul Trivedi, and Jelena Vučković. Quantum optics of soliton microcombs. *Nature Photonics*, 16(1):52–58, January 2022.

[13] Melissa A. Guidry, Daniil M. Lukin, Ki Youl Yang, and Jelena Vučković. Multimode squeezing in soliton crystal microcombs. *Optica*, 10(6):694, June 2023.

[14] Eran Lustig, Melissa A. Guidry, Daniil M. Lukin, Shanhui Fan, and Jelena Vučković. Quadrature-dependent lattice dynamics of dissipative microcombs. *Nature Photonics*, 19(11):1247–1254, November 2025.

[15] A. Bensemhoun, S. Cassina, C. Gonzalez-Arciniegas, M. F. Melalkia, G. Patera, J. Faugier-Tovar, Q. Wilmart, S. Olivier, A. Zavatta, A. Martin, J. Etesse, L. Labonté, O. Pfister, V. D'Auria, and S. Tanzilli. Multipartite quantum correlated bright frequency combs. *Phys. Rev. Res.*, 7:033173, Aug 2025.

[16] Cheng Zhang, Guoguo Kang, Jin Wang, Yijie Pan, and Jifeng Qu. Inverse design of soliton microcomb based on genetic algorithm and deep learning. *Optics Express*, 30(25):44395, December 2022.

[17] Erwan Lucas, Su-Peng Yu, Travis C. Briles, David R. Carlson, and Scott B. Papp. Tailoring microcombs with inverse-designed, meta-dispersion microresonators. *Nature Photonics*, 17(11):943–950, November 2023.

[18] Celine Mazoukh, Luigi Di Lauro, Imtiaz Alamgir, Bennet Fischer, Nicolas Perron, A. Aadhi, Armaghan Eshaghi, Brent E. Little, Sai T. Chu, David J. Moss, and Roberto Morandotti. Genetic algorithm-enhanced microcomb state generation. *Communications Physics*, 7(1):81, December 2024.

[19] Zhaoyang Sun, Yang Li, Benfeng Bai, Zhendong Zhu, and Hongbo Sun. Silicon nitride-based kerr frequency combs and applications in metrology. *Advanced Photonics*, 4(06), November 2022.

[20] Shaofei Wang, Hairun Guo, Xuekun Bai, and Xianglong Zeng. Broadband Kerr frequency combs and intracavity soliton dynamics influenced by high-order cavity dispersion. *Optics Letters*, 39(10):2880, May 2014.

[21] Mulong Liu, Leiran Wang, Qibing Sun, Siqi Li, Zhiqiang Ge, Zhizhou Lu, Chao Zeng, Guoxi Wang, Wenfu Zhang, Xiaohong Hu, and Wei Zhao. Influences of high-order dispersion on temporal and spectral properties of microcavity solitons. *Optics Express*, 26(13):16477, June 2018.

[22] A. V. Cherenkov, V. E. Lobanov, and M. L. Gorodetsky. Dissipative kerr solitons and cherenkov radiation in optical microresonators with third-order dispersion. *Phys. Rev. A*, 95:033810, Mar 2017.

[23] Qingjie Liu, Quan Lei, Pu Zhang, and Yingquan Ao. Soliton stabilization in microresonators with high order dispersion via pump phase modulation. *Results in Physics*, 46:106289, March 2023.

[24] Pedro Parra-Rivas, Damià Gomila, François Leo, Stéphane Coen, and Lendert Gelens. Third-order chromatic dispersion stabilizes Kerr frequency combs. *Optics Letters*, 39(10):2971, May 2014.

[25] Shuangyou Zhang, Toby Bi, and Pascal Del'Haye. Quintic Dispersion Soliton Frequency Combs in a Microresonator. *Laser & Photonics Reviews*, 17(10):2300075, October 2023.

[26] Ewelina Obrzud, Steve Lecomte, and Tobias Herr. Temporal solitons in microresonators driven by optical pulses. *Nature Photonics*, 11(9):600–607, September 2017.

[27] Hossein Taheri, Andrey B. Matsko, and Lute Maleki. Optical lattice trap for kerr solitons. *The European Physical Journal D*, 71(6), June 2017.

[28] Gregory Moille, Edgar F. Perez, Jordan R. Stone, Ashutosh Rao, Xiyuan Lu, Tahmid Sami Rahman, Yanne K. Chembo, and Kartik Srinivasan. Ultra-broadband kerr microcomb through soliton spectral translation. *Nature Communications*, 12(1), December 2021.

[29] Wenle Weng, Romain Bouchand, and Tobias J. Kippenberg. Formation and collision of multistability-enabled composite dissipative kerr solitons. *Phys. Rev. X*, 10:021017, Apr 2020.

[30] Stefanie Kroker, Stéphane Lanteri, Owen Miller, Jens Niegemann, and Lora Ramunno. Inverse design in photonics: introduction. *Journal of the Optical Society of America B*, 41(2):IDP1, January 2024.

[31] Sean Molesky, Zin Lin, Alexander Y. Piggott, Weiliang Jin, Jelena Vucković, and Alejandro W. Rodriguez. Inverse design in nanophotonics. *Nature Photonics*, 12(11):659–670, October 2018.

[32] Junhyeong Kim, Jae-Yong Kim, Jungmin Kim, Yun Hyeong, Berkay Neseli, Jong-Bum You, Joonsup Shim, Jonghwa Shin, Hyo-Hoon Park, and Hamza Kurt. Inverse design of nanophotonic devices enabled by optimization algorithms and deep learning: recent achievements and future prospects. *Nanophotonics*, 14(2):121–151, January 2025.

[33] Yuheng Chen, Alexander Montes McNeil, Taehyuk Park, Blake A. Wilson, Vaishnavi Iyer, Michael Bezick, Jae-Ik Choi, Rohan Ojha, Pravin Mahendran, Daksh Kumar Singh, Geetika Chitturi, Peigang Chen, Trang Do, Alexander V. Kildishev, Vladimir M. Shalaev, Michael Moebius, Wenshan Cai, Yongmin Liu, and Alexandra Boltasseva. Machine-learning-assisted photonic device development: a multiscale approach from theory to characterization. *Nanophotonics*, July 2025.

[34] Peter R. Wiecha, Arnaud Arbouet, Christian Girard, and Otto L. Muskens. Deep learning in nano-photonics: inverse design and beyond. *Photonics Research*, 9(5):B182, April 2021.

[35] Jiaqi Jiang, Mingkun Chen, and Jonathan A. Fan. Deep neural networks for the evaluation and design of photonic devices. *Nature Reviews Materials*, 6(8):679–700, December 2020.

[36] Ravi S. Hegde. Photonics inverse design: Pairing deep neural networks with evolutionary algorithms. *IEEE Journal of Selected Topics in Quantum Electronics*, 26(1):1–8, 2020.

[37] Yannick Augenstein, Taavi Repän, and Carsten Rockstuhl. Neural operator-based surrogate solver for free-form electromagnetic inverse design. *ACS Photonics*, 10(5):1547–1557, 2023.

[38] Parinaz Naseri and Sean V. Hum. A Generative Machine Learning-Based Approach for Inverse Design of Multilayer Metasurfaces. *IEEE Transactions on Antennas and Propagation*, 69(9):5725–5739, September 2021.

[39] Wei Ma, Feng Cheng, Yihao Xu, Qinlong Wen, and Yongmin Liu. Probabilistic Representation and Inverse Design of Metamaterials Based on a Deep Generative Model with Semi-Supervised Learning Strategy. *Advanced Materials*, 31(35):1901111, August 2019.

[40] Myeonghun Lee and Kyoungmin Min. Mgcvae: Multi-objective inverse design via molecular graph conditional variational autoencoder. *Journal of Chemical Information and Modeling*, 62(12):2943–2950, 2022. PMID: 35666276.

[41] Hyunsoo Park, Zhenzhu Li, and Aron Walsh. Has generative artificial intelligence solved inverse materials design? *Matter*, 7(7):2355–2367, 2024.

[42] Teng Tan, Cheng Peng, Zhongye Yuan, Xu Xie, Hao Liu, Zhenda Xie, Shu-Wei Huang, Yunjiang Rao, and Baicheng Yao. Predicting Kerr Soliton Combs in Microresonators via Deep Neural Networks. *Journal of Lightwave Technology*, 38(23):6591–6599, December 2020.

[43] Yixuan Xiang, Biyan Zhan, Haoxuan Zhang, and Xianwen Liu. Machine learning-assisted dispersion engineering for predicting ultra-flat soliton microcombs. *Optics Communications*, 581:131622, May 2025.

[44] Tianye Huang, Lin Chen, Mingkong Lu, Jianxing Pan, Chaoyu Xu, Pei Wang, and Perry Ping Shum. Rapid prediction of complex nonlinear dynamics in kerr resonators using the recurrent neural network. *Frontiers of Optoelectronics*, 18(1), September 2025.

[45] Ruibo Lan, Hongbin Hu, Yubin Zang, and Zuxing Zhang. Data-driven prediction of dissipative soliton generation in kerr fiber cavities. *Physica Scripta*, 100(8):086007, Aug 2025.

[46] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.

[47] Jong-Hoon Kim and Ic-Pyo Hong. Inverse design of electromagnetic metasurfaces utilizing infinite and separate latent space yielded by a machine learning-based generative model. *Journal of Electromagnetic Engineering and Science*, 24(2):178–190, 2024.

[48] Shenghui Zhao, Ruizhi Zhang, Meng Guo, Xuewei Zhang, Lifeng Nie, and Xiaoyu Pang. Deep learning enabled bandpass electromagnetic metasurface inverse design. In *2024 Cross Strait Radio Science and Wireless Technology Conference (CSRSWTC)*, pages 1–4. IEEE, 2024.

[49] Huan Huan Zhang, He Ming Yao, Lijun Jiang, and Michael Ng. Fast full-wave electromagnetic forward solver based on deep conditional convolutional autoencoders. *IEEE Antennas and Wireless Propagation Letters*, 22(4):779–783, 2022.

[50] Diederik P. Kingma and Max Welling. An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning*, 12(4):307–392, 2019.

[51] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020.

[52] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR.

[53] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.

[54] Jiawen Li, Jiang Guo, Yuanzhe Li, Zetian Mao, Jiaxing Shen, Tashi Xu, Diptesh Das, Jinming He, Run Hu, Yaerim Lee, Koji Tsuda, and Junichiro Shiomi. Inverse design of metamaterials with manufacturing-guiding spectrum-to-structure conditional diffusion model, 2025.

[55] Mohammad Abu-Mualla, Ellis Crabtree, Fredrick Michael, Yayue Pan, and Jida Huang. Inverse design of alloys via generative algorithms: Optimization and diffusion within learned latent space. *Advanced Intelligent Discovery*, n/a(n/a):202500069.

[56] Lu Zhu, Wei Hua, Cong Lv, and Yuanyuan Liu. Intelligent on-demand inverse design of high-degree-of-freedom meta-atoms: multi-constraint-latent diffusion model and contrastive learning. *Opt. Express*, 33(10):20239–20257, May 2025.

[57] Jan-Hendrik Bastek and Dennis M Kochmann. Inverse design of nonlinear mechanical metamaterials via video denoising diffusion models. *Nature Machine Intelligence*, 5(12):1466–1475, 2023.

[58] Li Zheng, Siddhant Kumar, and Dennis M. Kochmann. Diffumeta: Algebraic language models for inverse design of metamaterials via diffusion transformers, 2025.

[59] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023.

[60] Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *The Eleventh International Conference on Learning Representations*, 2023.

[61] Xingchao Liu, Chengyue Gong, and qiang liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations*, 2023.

[62] Xiaoshan Luo, Zhenyu Wang, Qingchang Wang, Xuechen Shao, Jian Lv, Lei Wang, Yanchao Wang, and Yanming Ma. Crystalflow: a flow-based generative model for crystalline materials. *Nature Communications*, 16(1), October 2025.

[63] Joonyoung F. Joung, Mun Hong Fong, Nicholas Casetti, Jordan P. Liles, Ne S. Dassanayake, and Connor W. Coley. Electron flow matching for generative reaction mechanism prediction. *Nature*, 645(8079):115–123, August 2025.

[64] Benjamin Kurt Miller, Ricky T. Q. Chen, Anuroop Sriram, and Brandon M Wood. FlowMM: Generating materials with riemannian flow matching. In *Forty-first International Conference on Machine Learning*, 2024.

[65] Yaron Lipman, Marton Havasi, Peter Holderrieth, Neta Shaul, Matt Le, Brian Karrer, Ricky T. Q. Chen, David Lopez-Paz, Heli Ben-Hamu, and Itai Gat. Flow matching guide and code. *CoRR*, abs/2412.06264, 2024.

[66] Aleksandr Tusnin, Alexey Tikan, Kenichi Komagata, and Tobias J. Kippenberg. Nonlinear dynamics and Kerr frequency comb formation in lattices of coupled microresonators. *Communications Physics*, 6(1):317, December 2023.

[67] Gregory Moille, Qing Li, Lu Xiyuan, and Kartik Srinivasan. pylle: A fast and user friendly lugiato-lefever equation solver. *Journal of Research of the National Institute of Standards and Technology*, 124, May 2019.

[68] Oliver Melchert and Ayhan Demircan. pyGLLE: A Python toolkit for solving the generalized Lugiato-Lefever equation. *SoftwareX*, 15:100741, July 2021.

[69] Xiaoxiao Xue, Pei-Hsun Wang, Yi Xuan, Minghao Qi, and Andrew M. Weiner. Microresonator Kerr frequency combs with high conversion efficiency. *Laser & Photonics Reviews*, 11(1):1600276, January 2017.

[70] Alexander Shmakov, Kevin Greif, Michael Fenton, Aishik Ghosh, Pierre Baldi, and Daniel Whiteson. End-to-end latent variational diffusion models for inverse problems in high energy physics. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 65102–65127. Curran Associates, Inc., 2023.

[71] H. Guo, M. Karpov, E. Lucas, A. Kordts, M. H. P. Pfeiffer, V. Brasch, G. Lihachev, V. E. Lobanov, M. L. Gorodetsky, and T. J. Kippenberg. Universal dynamics and deterministic switching of dissipative Kerr solitons in optical microresonators. *Nature Physics*, 13(1):94–102, January 2017.

[72] Mengjie Yu, Jae K. Jang, Yoshitomo Okawachi, Austin G. Griffith, Kevin Luke, Steven A. Miller, Xingchen Ji, Michal Lipson, and Alexander L. Gaeta. Breather soliton dynamics in microresonators. *Nature Communications*, 8:14569, February 2017.

[73] Mohammadreza Zandehshahvar, Yashar Kiarashinejad, Muliang Zhu, Hossein Maleki, Tyler Brown, and Ali Adibi. Manifold Learning for Knowledge Discovery and Intelligent Inverse Design of Photonic Nanostructures: Breaking the Geometric Complexity. *ACS Photonics*, 9(2):714–721, February 2022.

[74] Eran Lustig, Or Yair, Ronen Talmon, and Mordechai Segev. Identifying topological phase transitions in experiments using manifold learning. *Phys. Rev. Lett.*, 125:127401, Sep 2020.

[75] Peter Y. Lu, Rumen Dangovski, and Marin Soljačić. Discovering conservation laws using optimal transport and manifold learning. *Nature Communications*, 14(1), August 2023.

[76] Nicolas Romeo, Chris Chi, Aaron R. Dinner, and Elizabeth R. Jerison. Characterizing nonlinear dynamics by contrastive cartography. *PRX Life*, 3(3), September 2025.

[77] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, May 2020.