FACTORIZED INITIAL NOISE OPTIMIZATION WITH DIFFUSION AND FLOW PRIORS FOR MULTI-OBJECTIVE SCIENTIFIC INVERSE DESIGN

Anonymous authors
Paper under double-blind review

ABSTRACT

Scientific inverse design seeks structures in a vast design space that achieve a target property while satisfying domain-specific constraints. Generative models provide expressive priors over such design spaces, but existing inference-time methods steer generation toward either the target property or the constraints. We propose *factorized initial noise optimization* (FINO), which casts constrained inverse design as optimization over the initial noise of a pretrained generative model, and treats the reward and constraints as separate objectives. At each iteration, FINO first verifies whether the current iterate lies within the feasible set: if falls outside, then takes a constraint-restoration step; otherwise take a reward step toward the target property. This factorization confines reward optimization to the feasible region and requires no retraining, no modification of the sampling trajectory, and no constraint weights as trade-off. On topological photonic inverse design and molecule generation, FINO achieves low target-property error without sacrificing constraint satisfaction, and even robustly remains effective for target property values beyond the training distribution that of greatest scientific interest.

1 INTRODUCTION

Scientific inverse design aims to create structures that exhibit desired target properties within high-dimensional design spaces, with applications spanning photonic, molecular, and protein design (Zhang et al., 2026; Lin et al., 2026; Christopher et al., 2026). Despite substantial progress in each of these domains, inverse design problems share two fundamental challenges that remain largely unresolved. First, a candidate design must not only reach the target property, which is always formulated as a reward function, but also *satisfy domain-specific constraints* (Lin et al., 2026). A designed structure that achieves the target property but violates these constraints is of little scientific or practical value. For example, in topological photonic inverse design, a photonic structure needs to meet a target bandgap value (**Reward**), while preserving the C_3 rotational symmetry of each motif and avoiding features smaller than the minimum fabricable size (**Constraints**). Likewise, in molecular design, a molecule exhibiting a desired quantum-chemical property must also remain synthetically accessible. (Li et al., 2026; Nigam et al., 2023; Cardei et al., 2025). Second, the target properties of greatest scientific interest often lie beyond the range of property values observed in the training data. Because scientific data are costly to acquire, available datasets cover only a small fraction of vast design spaces (the space of drug-like small molecules alone is estimated to contain on the order of 10^{60} compounds (Bohacek et al., 1996)) and consequently span a limited range of property values. Designing structures whose properties extrapolate beyond this range, which we refer to as *out-of-distribution* (OOD) design, therefore requires generating samples in regions where the training data provide little or no support. (Klarner et al., 2024; Hong et al., 2026).

Diffusion and flow-based generative models offer a promising framework for scientific inverse design, as they can learn expressive priors over complex, high-dimensional design distributions, which have attracted growing attention across scientific domains (Lin et al., 2026; Ben-Hamu et al., 2024; Yang et al., 2026; Wang et al., 2025b). Within this framework, the target property is typically formulated as a reward that guides the sampling process of diffusion models. Along this line, much work has focused on reward alone, steering generation toward target properties while neglecting necessary

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

physical constraints (Ben-Hamu et al., 2024; Han et al., 2024; Wang et al., 2025b; Amiri & Nasirinia, 2026). A complementary line of work instead focuses on constrained generation, ensuring that generated samples remain scientifically feasible, but not optimize toward a target property (Christopher et al., 2024; Zhang et al., 2025; Cardei et al., 2025; Christopher et al., 2026). More recent work has begun to address both simultaneously, as requested by many real scientific inverse design tasks, yet merely combines them by folding the constraints into the guidance objective as additional weighted terms alongside the reward (Mazé & Ahmed, 2023). Such a formulation treats constraints as soft penalties, so constraint satisfaction is traded off against reward, where the balance becomes more challenging as the number of constraints grows. Therefore, there remains a need for a principled framework to solve **multi-objective problems**, optimizing the target property while systematically satisfying possibly multiple constraints that practical scientific applications demand.

As shown in previous works (Molesky et al., 2018; Zunger, 2018; Sanchez-Lengeling & Aspuru-Guzik, 2018; Ma et al., 2018), addressing the multi-objective problems described above directly in the sample space is challenging, where the data distribution over structured scientific designs has a complex landscape that renders both the reward optimization and the imposition of constraints difficult. This motivates to seek an alternative space to pose this problem. Previous work Song et al. (2025) shows that the initial noise perturbations demonstrates a principled property for noise-generation relationship, which naturally facilitate constrained sampling through the controllability in the surrogate space of initial noise. Because this space is endowed with a Gaussian prior with fully characterized geometry, for the targeted multi-objective inverse design problems, it is appealing to transfer and reformulate both the reward optimization and constraint satisfaction in terms of the initial noise. Furthermore, Jia et al. (2026) observes that initial noise optimization, which manipulates only the initial noise and leaves the sampling trajectory unaltered, exhibits promising potential for generation beyond the training domain, an attractive property for the out-of-distribution regime characteristic desired by scientific inverse design tasks. These findings motivate the central hypothesis of this work: the initial noise space serves as an effective surrogate to address multi-constraint, reward-guided inverse design problems through initial noise optimization (INO).

Motivated by this, we propose *factorized initial noise optimization* (FINO), which treats the target-property reward and the constraints as separate objectives through an interleaved iterative optimization framework. This factorization preserves the asymmetric roles of reward and constraints: the reward is optimized only within the feasible region, so progress toward the target property does not come at the expense of feasibility. Our contributions are summarized as follows:

- **A new formulation for scientific inverse design.** We introduce a new formulation of constrained scientific inverse design as an initial noise optimization problem over diffusion and flow models, transferring both the target property and the constraints in the initial noise space for effective optimization.
- **A principled framework for multi-objective inverse design.** Building on this formulation, we propose FINO, which optimizes the target property while systematically satisfying the multiple constraints that practical scientific applications impose. By factorizing the objective rather than collapsing it into a single weighted reward, FINO preserves the distinct roles of reward and constraints, and requires no retraining, and no modification of the sampling trajectory.
- **Strong empirical performance, including beyond the training distribution.** Across topological photonic inverse design and molecular generation, FINO matches or exceeds the target-property accuracy of five inference-time baselines while satisfying all constraints, which the baselines frequently violate. On topological photonic inverse design, it further produces feasible designs for target values in the tail of, and beyond, the training range, where baseline methods degrade sharply, supporting the robustness of our method.

2 Related Works

Diffusion and Flow Models Diffusion and flow models are generative models that transport samples from a noise distribution to the data distribution by training a neural network to reverse a prescribed corruption process. We denote a clean sample by $\boldsymbol{x}_0$ and pure noise by $\boldsymbol{x}_T \sim \mathcal{N}(0, \boldsymbol{I})$. For flow models (Lipman et al., 2023), we set $T = 1$ for data distribution to unify the notation. Denoising diffusion probabilistic models (DDPMs) (Ho et al., 2020) corrupt data with the forward

process $\boldsymbol{x}_t = \sqrt{\bar{\alpha}_t}\boldsymbol{x}_0 + \sqrt{1-\bar{\alpha}_t}\boldsymbol{\epsilon}$, $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \boldsymbol{I})$, and train a neural network $\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t)$ to predict the noise. Although the learned reverse chain $p_\theta(\boldsymbol{x}_{t-1} \mid \boldsymbol{x}_t)$ is stochastic, the corresponding diffusion SDE admits a deterministic probability-flow ODE (Song et al., 2021b),

$$d\boldsymbol{x}_t = \Big[f(t)\boldsymbol{x}_t - \frac{1}{2}g(t)^2 \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t)\Big]dt, \tag{1}$$

and DDIM (Song et al., 2021a) is its discretization,

$$\boldsymbol{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}}\hat{\boldsymbol{x}}_0(\boldsymbol{x}_t) + \sqrt{1-\bar{\alpha}_{t-1}}\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t), \quad \hat{\boldsymbol{x}}_0(\boldsymbol{x}_t) = \frac{\boldsymbol{x}_t - \sqrt{1-\bar{\alpha}_t}\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t)}{\sqrt{\bar{\alpha}_t}}, \tag{2}$$

which injects no additional noise and permits skipping reverse timesteps. A parallel line of work, flow matching (Lipman et al., 2023), instead directly learns a velocity field $v_\theta(\boldsymbol{x}_t, t)$. The model is typically trained on the linear interpolant

$$\boldsymbol{x}_t = (1-t)\boldsymbol{x}_0 + t\boldsymbol{\epsilon}, \qquad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \boldsymbol{I}), \tag{3}$$

where $\boldsymbol{x}_0$ corresponds to data and $\boldsymbol{x}_1$ to noise. At inference, samples are generated by integrating the ODE

$$d\boldsymbol{x}_t = v_\theta(\boldsymbol{x}_t, t)dt \tag{4}$$

backward in time from $\boldsymbol{x}_1 \sim \mathcal{N}(0, \boldsymbol{I})$ to $\boldsymbol{x}_0$.

Inference-time reward alignment Inference-time alignment steers the pretrained diffusion and flow model towards high-reward samples without retraining. There are three representative methods for inference-time alignment: classifier guidance, particle-based sampling, and initial noise optimization. Classifier guidance (Dhariwal & Nichol, 2021) adds reward gradient for every reverse step. Particle-based sampling generates multiple sampling trajectories and selects or resamples particles according to their reward. For example, TDS-SMC (Wu et al., 2023) resamples particles using intermediate reward estimates, allowing computation to focus more on promising trajectories. Initial noise optimization (Eyring et al., 2024; Tang et al., 2025) considers the generative model as a black box and directly optimizes on the initial noise space to maximize the reward of the final sample without modifying the reverse trajectory. Optimizing the initial Gaussian noise has been extensively explored in text-to-image (T2I) generation, typically targeting a single reward such as text-image semantic alignment, aesthetic quality, or spatial consistency (Wallace et al., 2023; Eyring et al., 2024; Tang et al., 2025; Guo et al., 2024), and several works have extended it to scientific domains with a single scientific reward, including molecule generation (Ben-Hamu et al., 2024; Wang et al., 2025b), protein design (Wang et al., 2025b), and aerodynamic inverse design (Yang et al., 2026). None of these works treats a constraint separately from the reward.

Multi-Objective and Constraint-Aware Generation Scientific design problems often involve multiple objectives or explicit constraints. One representative example is molecule generation, where it is inherently a multi-objective problem (Fromer & Coley, 2023; Wang et al., 2025a). Training-free approaches combine multiple property guidance as a weighted sum during the reverse sampling (Han et al., 2024), whereas Amiri & Nasirinia (2026) fine-tune the model with Pareto-based multi-objective reinforcement learning. A separate line of work enforces constraints by incorporating constraint-enforcement or projection steps into the reverse process on molecule generation (Cardei et al., 2025), path planning and control (Christopher et al., 2024; Zhang et al., 2025) and protein design (Christopher et al., 2026). FINO differs from both lines as it optimizes a target property under a constraint in the initial noise space.

3 Problem Definition

In many scientific applications, the objective is to achieve a desired property while satisfying constraints. Let $F_p : \mathbb{R}^d \to \mathbb{R}$ be a differentiable target property predictor with target value $y^\star$, and let $F_c : \mathbb{R}^d \to \mathbb{R}$ be a differentiable constraint predictor. Here we want to solve

$$\boldsymbol{x}^\star = \arg\min_{\boldsymbol{x} \in \mathcal{C}} \ell(F_p(\boldsymbol{x}), y^\star), \quad \mathcal{C} = \{\boldsymbol{x} : F_c(\boldsymbol{x}) \geq \tau_c\} \tag{5}$$

where $\mathcal{C} = \{\boldsymbol{x} : F_c(\boldsymbol{x}) \geq \tau_c\}$ is a feasible set defined by a threshold τ_c, and $\ell(\cdot, \cdot)$ is an operator that calculates the distance. In molecule generation, F_p predicts the quantum chemical property and

F_c is the synthetic accessibility score; in topological photonic design, F_p is a bandgap value and F_c is the C_3 symmetry and fabrication constraint. Here we define the **reward** and the **constraint** as

$$R(\boldsymbol{x}) := -\ell(F_p(\boldsymbol{x}), y^\star), \qquad C(\boldsymbol{x}) := F_c(\boldsymbol{x}) - \tau_c, \tag{6}$$

so that $\boldsymbol{x} \in \mathcal{C} \Leftrightarrow C(\boldsymbol{x}) \geq 0$. Let $G_\theta(\cdot\,; M) : \mathbb{R}^d \to \mathbb{R}^d$ be a deterministic sampler of a pretrained diffusion or flow model, which maps initial noise $\boldsymbol{z} \sim \mathcal{N}(0, \boldsymbol{I})$ to a clean sample with M ODE steps. As our method is in an initial noise optimization framework, we express the reward and the constraint as functions of the initial noise,

$$r(\boldsymbol{z}) := R(G_\theta(\boldsymbol{z}; M)), \quad c(\boldsymbol{z}) := C(G_\theta(\boldsymbol{z}; M)), \tag{7}$$

and search over the initial noise $\boldsymbol{z}$ instead of the sample $\boldsymbol{x}$, so that Eq. 5 becomes

$$\boldsymbol{z}^\star = \arg\max_{\boldsymbol{z}} \; r(\boldsymbol{z}) \quad \text{s.t.} \quad c(\boldsymbol{z}) \geq 0. \tag{8}$$

We also write $c_N(\boldsymbol{z}) := C(G_\theta(\boldsymbol{z}; N))$ for the margin evaluated with a coarser N-step rollout, $N < M$.

4 Method

4.1 Factorized Reward Optimization

Eq. 8 contains two objectives with different roles. The reward r is to be *maximized*: the larger the better, while the constraint c is to be *satisfied*: once $c(\boldsymbol{z}) \geq 0$, a larger margin brings no advantage. This asymmetry suggests that the two should not be optimized at the same time. Instead of maximizing a single weighted combination of r and c, FINO treats Eq. 8 as two separate objectives. At every iteration, FINO takes a gradient step on either the reward or the constraint, chosen according to how many of the perturbed candidates remain feasible after a short rollout. We call this scheme *factorized* because the objective is split into a constraint-restoration part and a reward part, each handled by its own update and never combined. Algorithm 1 summarizes the procedure and Fig. 1 illustrates it.

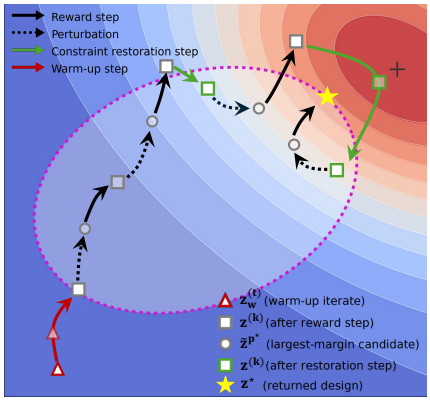

Figure 1: **Illustration of the optimization trajectory in the initial noise space for our proposed FINO algorithm.** To address multi-objective inverse design, FINO steps on either the reward optimization or constraint restoration, depending on the constraint feasibility estimated from a few-step rollout.

Warm up A randomly drawn sample $\boldsymbol{z}_w \sim \mathcal{N}(0, \boldsymbol{I})$ is generally infeasible: for constraints like C_3 symmetry or synthetic accessibility, a sample from the prior satisfies them only by chance. We therefore first run gradient ascent with respect to the constraint

$$\boldsymbol{z}_w^{(t+1)} = \boldsymbol{z}_w^{(t)} + \eta_c \, \nabla(c(\boldsymbol{z}_w^{(t)}) + \lambda \Omega(\boldsymbol{z}_w^{(t)})), \qquad \Omega(\boldsymbol{z}) = (d-1)\log\|\boldsymbol{z}\|_2 - \tfrac{1}{2}\|\boldsymbol{z}\|_2^2, \tag{9}$$

until $c(\boldsymbol{z}_w^{(t)}) \geq 0$ to ensure the starting point of our algorithm lies on the feasible set. Here, Ω denotes the log-density of the norm of a standard Gaussian vector. Maximizing the regularizer keeps the noise on the typical shell $||\boldsymbol{z}||_2 \approx \sqrt{d}$, preventing the optimized initial noise from deviating from the Gaussian distribution (Eyring et al., 2024; Samuel et al., 2023). Not every draw, however, is an equally good starting point for reaching the feasible set. For some draws, the local constraint landscape may be flat or poorly conditioned, causing gradient ascent to make little progress. To avoid spending excessive optimization budget, we cap each warm-up attempt at T_w steps. If feasibility is not reached within this budget, we resample $\boldsymbol{z}_w \sim \mathcal{N}(0, \boldsymbol{I})$ and restart the warm-up procedure.

Perturb and Filter Before performing any gradient update we first verify the current iteration is in the feasible set. Rather than trusting a pointwise check, FINO first examines a small neighborhood around it. We form $P+1$ candidates: the current noise itself, $\tilde{\boldsymbol{z}}^0 = \boldsymbol{z}^{(k)}$, and P variance-preserving perturbations

$$\tilde{\boldsymbol{z}}^p = \sqrt{1-\rho}\,\boldsymbol{z}^{(k)} + \sqrt{\rho}\,\boldsymbol{\epsilon}^p, \quad \boldsymbol{\epsilon}^p \sim \mathcal{N}(0, \boldsymbol{I}), \quad p = 1, \ldots, P, \tag{10}$$

with strength $\rho \in (0, 1)$. The variance-preserving step keeps every candidate on the typical shell of the prior, so that each is a valid input to the sampler. Each candidate is then given as an input to the N-step ODE sampler and tested against the constraint

$$\mathcal{F}^{(k)} = \{p \in \{0, \dots, P\} : c_N(\tilde{\boldsymbol{z}}^p) \geq 0\}. \tag{11}$$

Constraint Restoration Step If no candidate is feasible, $\mathcal{F}^{(k)} = \emptyset$, we assume that the previous gradient update has moved the current sample outside the feasible set. We therefore take a gradient step

$$\boldsymbol{z}^{(k+1)} = \boldsymbol{z}^{(k)} + \eta_c \nabla (c(\boldsymbol{z}^{(k)}) + \lambda \Omega(\boldsymbol{z}^{(k)})). \tag{12}$$

toward the constraint satisfaction region to guide the sample back into the feasible set.

Reward Step Otherwise, we select the largest margin $p^\star = \arg\max_p c(\tilde{\boldsymbol{z}}^p)$ and take a gradient step on the reward from that candidate to approach the target property,

$$\boldsymbol{z}^{(k+1)} = \tilde{\boldsymbol{z}}^{p^\star} + \eta_r \nabla (r(\tilde{\boldsymbol{z}}^{p^\star}) + \lambda \Omega(\tilde{\boldsymbol{z}}^{p^\star})). \tag{13}$$

Selecting the largest margin rather than the best current $r(\cdot)$ gives the target property step the most budget allowing more freedom to the reward update. Even if the reward step leaves the feasible set after the gradient step, the next constraint step will handle this issue. To ensure our final sample stays on the feasible set, at every iteration we save the highest reward sample that satisfies the feasible set.

4.2 Theoretical Analysis

In this section, we show that the constraint restoration step brings the current sample into the feasible set after a finite number of iterations. For clarity, we set $\lambda = 0$ in this section.

Proposition 1 (Finite-step Constraint Satisfaction) *Assuming (i) $c(\boldsymbol{z})$ is L_c-smooth, and (ii) exists $\tau_{\min} > 0$ such that $\|\nabla c(\boldsymbol{z})\|_2 \geq \tau_{\min}$. Then, for any $0 \leq \eta_c \leq 2/L_c$, the constraint restoration step satisfies*

$$c(\boldsymbol{z} + \eta_c \nabla c(\boldsymbol{z})) \geq c(\boldsymbol{z}) + \eta_c (1 - L_c \eta_c / 2) ||\nabla c(\boldsymbol{z}))||_2^2 \geq c(\boldsymbol{z}) \tag{14}$$

and starting from any $\boldsymbol{z}_0$ with $c(\boldsymbol{z}) < 0$, the iterate enters the feasible set after at most

$$T_c = \left\lceil \frac{-c(\boldsymbol{z}_0)}{\eta_c \left(1 - \frac{L_c \eta_c}{2}\right) \tau_{\min}^2} \right\rceil \tag{15}$$

steps

Remarks. Proposition 1 states that constraint restoration steps make monotone progress. Assumption *(ii)* is a local condition that may fail in the vanishing gradient region of the constraint landscape. This is the reason why we need a warm-up budget T_w and the resampling rule is designed.

5 Experiments

In this section, we evaluate FINO on two scientific inverse design tasks: topological photonic inverse design (Section 5.1) and constrained molecule generation (Section 5.2). We compare our method against five inference-time guided or constrained diffusion sampling methods: (1) classifier guidance (CG) (Dhariwal & Nichol, 2021); (2) augmented Lagrangian method(ALM) (Cardei et al., 2025) which handles the constraint through a Lagrangian term; (3) best-of-N (BoN); (4) twisted sequential Monte-Carlo (SMC) (Wu et al., 2023); and (5) initial noise optimization method (InitNO) (Ben-Hamu et al., 2024; Eyring et al., 2024) which optimize against the weighted combination of the reward and constraint. For CG and ALM, we implement the methods following their respective papers, build SMC upon on the code provided by Wu et al. (2023) and implement InitNO using the code on Jia et al. (2026). For a fair comparison, all baseline methods use the same weighted combination of the reward and constraint, with the weighting selected based on the InitNO setting. The specific details for the hyperparameters are defined in the Appendix B.

5.1 TOPOLOGICAL PHOTONIC INVERSE DESIGN

Background Topological photonics is a field that applies concepts of band-structure topology to engineer optical transport Lu et al.; Ozawa et al.; Wang et al.. Specifically, valley photonic crystals (VPCs) break the inversion symmetry and open a bandgap at the K and K' valleys of the Brillouin zone, resulting in robustness to structural perturbations Dong et al. (2017); Xiao et al.. Albeit to rapid progress, optimal topological design is not guaranteed from such symmetry-restricted, human-intuitive templates. This motivates the shift to inverse design, where desired properties are specified a priori and the underlying photonic structure is predicted by searching over physically realizable configurations. Here the objective is to generate topological photonic structures that achieve target bandgap values while simultaneously satisfying C_3 rotational symmetry and the minimum feature-size constraints. The C_3 constraint is inherent to the dataset construction, in which the designs are defined on a honeycomb lattice with three-fold rotationally symmetric motifs. The minimum feature-size constraint arises from fabrication requirements, ensuring that each motif satisfies a minimum feature size of 16 pixels to remain physically manufacturable. Prior work (Zhang et al., 2026) incorporates classifier guidance to sample a specific bandgap value topological structure, but limits its performance on out-of-distribution bandgap value samples.

Problem Setting The original dataset consists of 7689 VPC designs that are constructed by random three-fold rotationally symmetric Bezier motifs at the two inequivalent sites of a honeycomb unit cell. Illustrated in Fig. 3, each design is represented as a binary mask $M \in \{0,1\}^{128\times64}$, where (1) denotes SiN and (0) denotes air. To obtain a training distribution with well-defined tails, we subsample the negative bandgap designs so that their density matches that of the positive bandgap designs. The modified data distributions are shown in Fig 2.

We evaluate the methods under three target regimes. We define the in-distribution (InD) targets as the 25th and 75th percentiles of the training bandgap distribution, and the tail-distribution (Tail) targets as the 0.1th and 99.9th percentiles. Out-of-distribution (OOD) targets are selected outside the range of bandgap values that are not observed during diffusion-model training. We train the diffusion model following the architecture of (Zhang et al., 2026). In our experiment, we set the bandgap predictor as a reward and C_3 rotational symmetry reward and fabrication constraint reward as constraint predictors. During inference time, the bandgap values are calculated by the differentiable surrogate model proposed by Zhang et al. (2026), which we retrain on the complete original dataset to provide accurate bandgap estimates over a broader range of structures. For final evaluation, we perform full-wave electromagnetic simulations using Lumerical FDTD on the generated samples. To evaluate the constraints, we extract patches centered at each motif. The C_3 symmetry constraint is evaluated by comparing each motif with its rotated counterparts, while the fabrication constraint is computed from the number of SiN. The specific details for calculating the bandgap value and the constraints are explained in the Appendix D E.

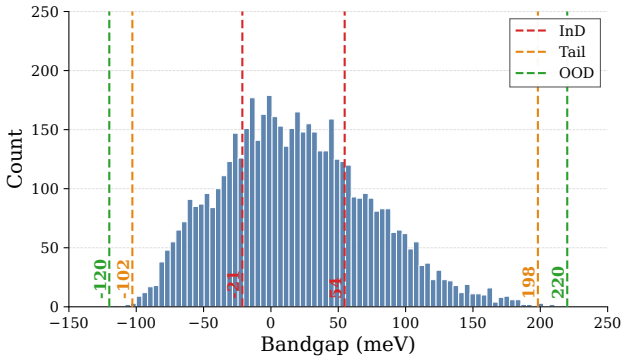

Figure 2: **Distribution of the target property: bandgap value for topological photonic inverse design.** Red line indicates the bandgap value of in-distribution samples $(-21\text{meV}, 54\text{meV})$, orange line indicates the bandgap value of in-distribution samples $(-102\text{meV}, 198\text{meV})$, and green line indicates the bandgap value of in-distribution samples $(-120\text{meV}, 220\text{meV})$.

Experiment Results We report quantitative results on InD, Tail, and OOD bandgap targets in Tab. 4 and qualitative results in Fig. 3. For each method and target, we generate 50 designs from a fixed set of random seeds and report the mean and standard deviation of the FDTD-simulated bandgap error together with the fraction of designs that violate either constraint. FINO is the only

Table 1: **Photonic inverse-design results for InD, Tail, and OOD target property**. MAE is computed from FDTD simulations and reported as mean ± standard deviation over 50 designs; Viol. (%) is the fraction of designs that violate either the C_3 symmetry or the fabrication constraint. Best and second-best MAE are shown in bold and underlined, respectively. †One ALM run diverged and produced no valid design; it is excluded from the statistics.

	InD				Tail				OOD			
	−21meV		54meV		−102meV		198meV		−120meV		220meV	
Method	MAE ↓	Viol. (%) ↓	MAE ↓	Viol. (%) ↓	MAE ↓	Viol. (%) ↓	MAE ↓	Viol. (%) ↓	MAE ↓	Viol. (%) ↓	MAE ↓	Viol. (%) ↓
CG	2.87 ± 7.01	40	2.06 ± 1.77	64	2.66 ± 2.75	6	12.66 ± 21.42	80	$\underline{6.26 \pm 7.00}$	2	19.33 ± 28.81	84
ALM	$2.92 \pm 7.07^{\dagger}$	40	2.02 ± 1.77	62	$\underline{2.63 \pm 2.71}$	6	12.67 ± 21.41	80	8.10 ± 15.34	2	19.34 ± 28.81	84
BoN	$\underline{1.80 \pm 1.44}$	34	$\mathbf{1.60 \pm 1.27}$	46	6.59 ± 9.45	16	10.81 ± 8.38	56	14.83 ± 16.67	16	27.09 ± 13.35	60
SMC	1.85 ± 1.43	46	1.84 ± 1.28	48	3.25 ± 2.11	12	$\underline{3.73 \pm 2.01}$	54	6.69 ± 3.73	0	$\mathbf{4.20 \pm 4.02}$	78
InitNO	24.56 ± 31.60	58	18.24 ± 28.77	70	26.81 ± 32.31	64	41.84 ± 54.70	80	56.52 ± 43.71	90	60.06 ± 71.20	84
FINO	$\mathbf{1.62 \pm 1.20}$	**0**	$\underline{1.82 \pm 1.69}$	**0**	$\mathbf{2.06 \pm 2.12}$	**0**	$\mathbf{3.41 \pm 2.47}$	**0**	$\mathbf{3.58 \pm 3.53}$	**0**	$\underline{5.48 \pm 4.24}$	**0**

method with zero constraint violation at every target, and it attains the lowest MAE at four of the six targets and the second-lowest at the remaining two. As the target moves into the upper tail and beyond the training range (198 and 220 meV), every baseline degrades in both accuracy and violation rate, whereas FINO keeps its error below 5.5 meV with no violations. InitNO, which also optimizes the initial noise but through a weighted sum of reward and constraint, performs worst throughout (MAE 18-60 meV, violation 58-90%): the constraint is not enforced strongly enough to preserve the C_3 symmetry of the motifs, and the resulting asymmetric structures yield bandgaps far from the target in FDTD simulation. This indicates that the noise space alone does not resolve the conflict between reward and constraint; yet the factorized update does. Fig. 3 visualize one randomly selected design from the 50 samples generated by each method at each target value. FINO produces a feasible design at every target, including 198 and 220 meV where every baseline violates a constraint, while keeping the simulated bandgap within 5 meV of the target. The baselines reach high bandgaps only with asymmetric or fragmented motifs, and InitNO drifts far from the target.

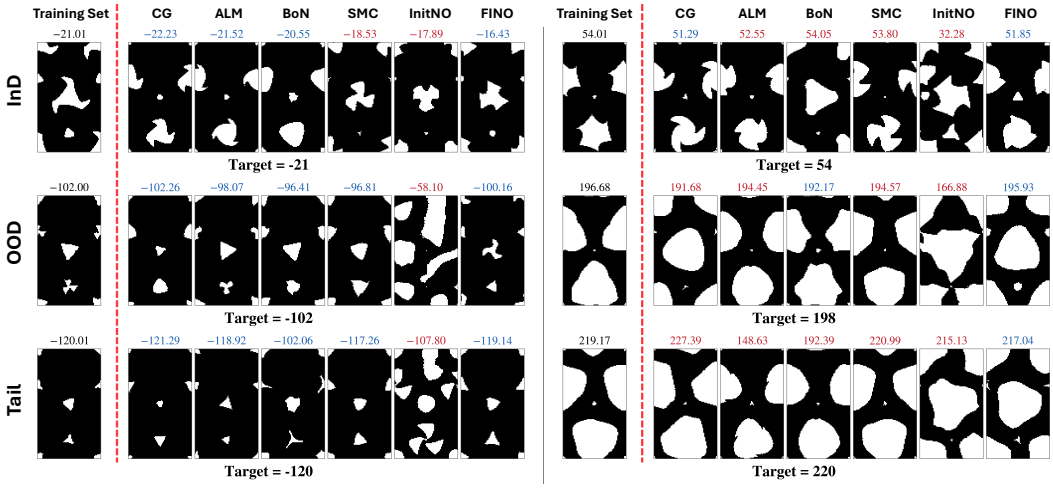

Figure 3: **Qualitative results on topological photonic inverse design.** Columns are grouped by target bandgap; rows correspond to InD, Tail, and OOD targets. The title of each design is its bandgap in meV obtained from full-wave FDTD simulation; blue indicates that the design satisfies both the C_3 symmetry and the minimum feature-size constraints, and red that it violates at least one of them.

5.2 Molecule Generation

Problem Setting We experiment on molecule generation using the QM9 dataset (Ruddigkeit et al., 2012; Ramakrishnan et al., 2014), which is a well-known dataset consisting small molecules with at most 9 heavy atoms from C, O, N, F. We set the reward as the quantum chemical property values, which include polarizability α, orbital energies $\varepsilon_{\text{HOMO}}$, $\varepsilon_{\text{LUMO}}$ and their gap $\Delta\varepsilon$, dipole moment μ, and heat capacity C_v. For the constraint, we use the synthetic accessibility (SA) score (Ertl &

Table 2: **Molecular property optimization under the synthetic accessibility (SA) constraint.** MAE measures target-property error; Viol. Surr. and Viol. Real denote the percentage of molecules whose surrogate-predicted and true RDKit (Ertl & Schuffenhauer, 2009) SA scores exceed 3.5, respectively. Best and second-best are in bold and underline.

	α (Bohr^3)			$\Delta\varepsilon$ eV			$\varepsilon_{\text{HOMO}}$ eV		
Method	MAE ↓	Viol. Surr. (%) ↓	Viol. Real (%) ↓	MAE ↓	Viol. Surr. (%) ↓	Viol. Real (%) ↓	MAE ↓	Viol. Surr. (%) ↓	Viol. Real (%) ↓
CG	5.580	78.0	82.8	1.171	80.0	83.7	0.596	81.0	84.8
ALM	5.334	77.0	55.6	1.293	76.0	55.1	0.569	81.0	56.0
BoN	**1.202**	83.0	88.3	**0.166**	77.0	80.0	**0.089**	88.0	87.9
SMC	4.733	75.0	76.4	0.691	69.0	73.9	0.329	75.0	79.1
InitNO	3.891	85.0	85.3	0.625	87.0	95.3	0.211	88.0	89.5
FINO	1.697	**0.0**	**12.8**	0.343	**0.0**	**17.0**	0.177	**0.0**	**20.2**

	$\varepsilon_{\text{LUMO}}$ eV			μ (D)			C_v $\frac{\text{cal}}{\text{K·mol}}$		
Method	MAE ↓	Viol. Surr. (%) ↓	Viol. Real (%) ↓	MAE ↓	Viol. Surr. (%) ↓	Viol. Real (%) ↓	MAE ↓	Viol. Surr. (%) ↓	Viol. Real (%) ↓
CG	1.028	81.0	83.2	1.360	80.0	80.9	3.136	81.0	82.6
ALM	1.090	81.0	65.3	1.562	80.0	52.3	2.555	77.0	62.0
BoN	**0.145**	78.0	79.6	**0.182**	78.0	83.0	**0.310**	85.0	86.7
SMC	0.566	62.0	74.5	0.527	56.0	58.8	1.438	62.0	71.4
InitNO	0.512	89.0	90.6	0.314	95.0	92.8	1.613	88.0	93.5
FINO	0.375	**0.0**	**26.6**	0.317	**0.0**	**14.9**	0.596	**0.0**	**21.3**

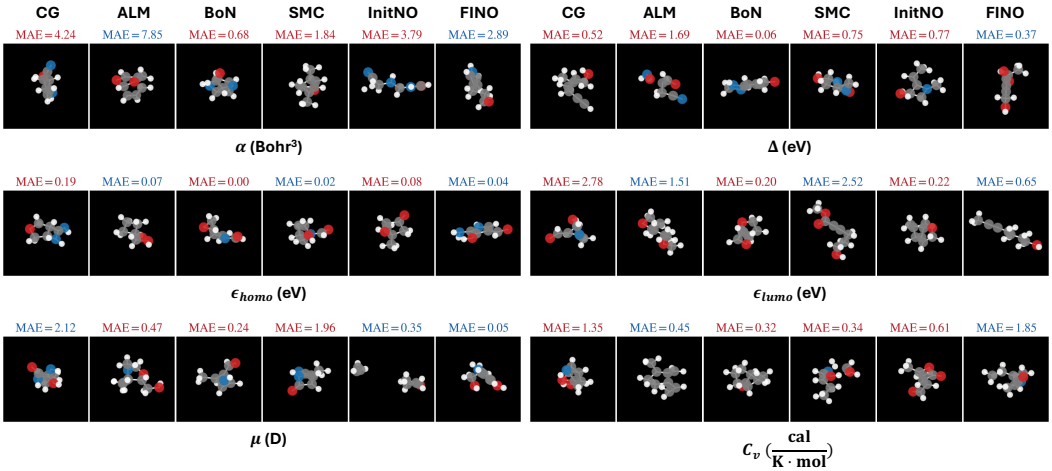

Figure 4: **Qualitative results on molecule generation.** One generated molecule per method for each of the six target properties (α, $\Delta\varepsilon$, $\varepsilon_{\text{HOMO}}$, $\varepsilon_{\text{LUMO}}$, μ, C_v). The title of each molecule is the absolute error between its predicted property and the target. Blue indicates a molecule that satisfies the constraint(RDKit SA score ≤ 3.5) and red indicates a violation.

Schuffenhauer, 2009), which ranges from 1 (easy to synthesize) to 10 (difficult to synthesize). We set the threshold of this constraint as 3.5. For calculating the property values during inference time, we use the surrogate model architecture proposed by Hoogeboom et al. (2022). For the generative model for molecular generation, we used a flow-based generative model from EquiFM (Song et al., 2023).

Experiment Result We report the quantitative results on Tab. 2 and qualitative results in Fig. 4. Under the surrogate SA model that all methods optimize, FINO satisfies the constraint for every generated molecule across all six target properties, whereas every baseline violates it on 56-95% of its samples. At the same time, FINO attains the second-lowest MAE on five of the six properties. The only method with lower MAE, BoN, reaches it by disregarding the constraint, violating it on 77-88% of its molecules; the remaining baselines are worse on both metrics. The gap exists under the true RDKit SA score. Since the constraint is optimized through a differentiable surrogate rather than the RDKit evaluator itself, a molecule that satisfies the surrogate model is not guaranteed to satisfy the true evaluator. This gap is due to the surrogate model error rather than the optimizer. Compared with

the surrogate model result, FINO solves the problem exactly, with zero violations. Fig. 4 presents one randomly chosen molecule for each method and target property. FINO consistently generates molecules that satisfy the RDKit SA threshold while remaining close to the desired property value, whereas the baseline methods more frequently produce molecules that violate the constraint.

5.3 Generalization to OOD Design for Scientific Discovery

In this section, we evaluate FINO on the topological photonic inverse-design task with targets that lie outside the training distributions of both the diffusion model and the CNN surrogate. Whereas the OOD targets in Section 5.1 (-120 and 220 meV) are unseen by the diffusion model but still within the range of the original dataset on which the surrogate is trained (-130 to 220 meV), here we target -150 and 250 meV, beyond the range of any training data.

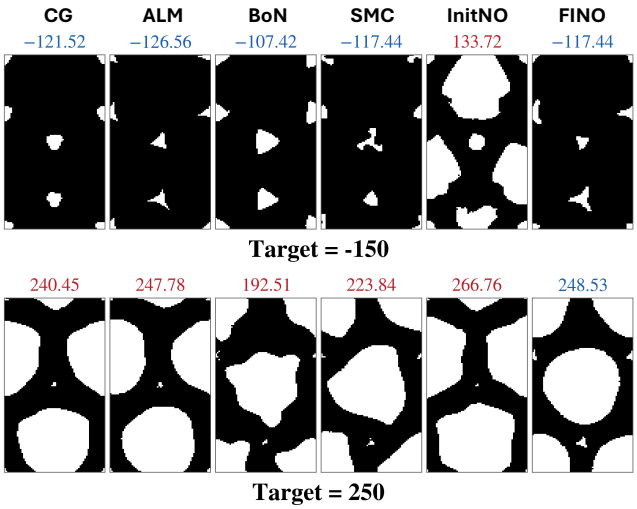

Figure 5: **Qualitative results on topological photonic inverse design.** Columns are grouped by target bandgap. The title of each design is its bandgap in meV obtained from full-wave FDTD simulation; blue indicates that the design satisfies both the C_3 symmetry and the minimum feature-size constraints, and red that it violates at least one of them.

This setting reflects a central challenge in scientific inverse design: discovering designs with target properties entirely beyond the training range, which lack data support yet are often of greatest scientific interest. Tab. 5 reports the quantitative results and Fig. 5 the qualitative results. At 250 meV, FINO generates feasible designs within 12 meV of the target, while every baseline violates the constraints on 60-100% of its designs. At -150 meV, no method reaches the target accurately, but FINO still attains the lowest MAE and remains the only method with zero constraint violations. These results indicate that FINO continues to produce feasible designs and to move in the direction of the target even where neither the prior nor the surrogate has support.

6 Conclusion, Limitation & Future Work

We presented FINO, a framework for multi-objective scientific inverse design that optimizes the initial noise of a pretrained diffusion or flow model while treating the target-property reward and the physical constraints as separate objectives. The factorization confines the reward optimization to the feasible region, allowing it to generate low-property-error samples while satisfying the domain design constraints. On topological photonic inverse design and molecule generation, FINO is the only method that achieves low target-property error while satisfying the constraints at every target, and it remains robustly effective for target values beyond the training distribution, where the baselines degrade sharply. A limitation of FINO is its reliance on differentiable surrogate predictors to approximate the reward or constraints. Constraint satisfaction under a surrogate does not always transfer to the true evaluator, as observed in the molecular generation task. Thus, a natural direction for future work is to extend our proposed multi-objective framework to non-differentiable or simulation-based rewards and constraints, which is also relevant yet complementary with ongoing research on reinforcement learning for diffusion and flow-based models

AI use statement

In this work, we used generative AI tools to assist with several aspects of the research and preparation of the manuscript. Specifically, (1) We used LLM for editing and polishing the writing for clarity

and grammer, (2) We used AI assistance for formatting and designing figures, (3) We used AI for searching related works but not used it to start it as our motivation, (4) We used AI assistance for debugging, and editing research and visualization code. We have not used generative AI tools for other things than listed. Generative AI was not used to generate experimental results or data and we take full responsibility for the methodology, experiments, results, claims, citations, and final content of the paper.

REFERENCES

Masoud Amiri and Zahra Nasirinia. Hybridmolgen: a unified framework for goal-directed molecular generation via multi-objective reinforcement learning. *Bioinformatics*, 42, 2026. URL https://api.semanticscholar.org/CorpusID:287193791.

Heli Ben-Hamu, Omri Puny, Itai Gat, Brian Karrer, Uriel Singer, and Yaron Lipman. D-flow: Differentiating through flows for controlled generation. In *International Conference on Machine Learning*, 2024. URL https://api.semanticscholar.org/CorpusID:267770268.

Regine S. Bohacek, Colin McMartin, and Wayne C. Guida. The art and practice of structure-based drug design: A molecular modeling perspective. *Medicinal Research Reviews*, 16, 1996. URL https://api.semanticscholar.org/CorpusID:44271689.

Michael Cardei, Jacob K Christopher, Bhavya Kailkhura, Thomas Hartvigsen, and Ferdinando Fioretto. Constrained discrete diffusion. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=Es4s9dtCjR.

Jacob K Christopher, Stephen Baek, and Ferdinando Fioretto. Constrained synthesis with projected diffusion models, 2024. URL https://arxiv.org/abs/2402.03559.

Jacob K Christopher, Austin Seamann, Jingyi Cui, Sagar D Khare, and Ferdinando Fioretto. Constrained diffusion for protein design with hard structural constraints. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=kkvqVRu2Zy.

Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021. URL https://openreview.net/forum?id=AAWuCvzaVt.

Jian-Wen Dong, Xiao-Dong Chen, Hanyu Zhu, Yuan Wang, and Xiang Zhang. Valley photonic crystals for control of spin and topology. *Nat. Mater.*, 16(3):298–302, March 2017. ISSN 1476-4660. doi: 10.1038/nmat4807.

Peter Ertl and Ansgar Schuffenhauer. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of Cheminformatics*, 1:8 – 8, 2009. URL https://api.semanticscholar.org/CorpusID:7423230.

Luca Eyring, Shyamgopal Karthik, Karsten Roth, Alexey Dosovitskiy, and Zeynep Akata. Reno: Enhancing one-step text-to-image models through reward-based noise optimization. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 125487–125519. Curran Associates, Inc., 2024. doi: 10.52202/079017-3987. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/e31bdea0a93741c2157eea705dd219eb-Paper-Conference.pdf.

Jenna C. Fromer and Connor W. Coley. Computer-aided multi-objective optimization in small molecule discovery. *Patterns*, 4(2):100678, 2023. ISSN 2666-3899. doi: https://doi.org/10.1016/j.patter.2023.100678. URL https://www.sciencedirect.com/science/article/pii/S2666389923000016.

Xiefan Guo, Jinlin Liu, Miaomiao Cui, Jiankai Li, Hongyu Yang, and Di Huang. Initno: Boosting text-to-image diffusion models via initial noise optimization. In *CVPR*, 2024.

Xu Han, Caihua Shan, Yifei Shen, Can Xu, Han Yang, Xiang Li, and Dongsheng Li. Training-free multi-objective diffusion model for 3d molecule generation. In *ICLR 2024*, April 2024. URL https://www.microsoft.com/en-us/research/publication/training-free-multi-objective-diffusion-model-for-3d-molecule-generation/.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html.

Haokai Hong, Wanyu Lin, Ming Yang, and Kay Chen Tan. Distributional priors guided diffusion for generating 3d molecules in low data regimes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 21743–21751, 2026.

Emiel Hoogeboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d. *ArXiv*, abs/2203.17003, 2022. URL https://api.semanticscholar.org/CorpusID:247839510.

Jing Jia, Wei Yuan, Sifan Liu, Liyue Shen, and Guanyang Wang. Weak diffusion priors can still achieve strong inverse-problem performance. *International Conference on Machine Learning*, 2026.

Leo Klarner, Tim G. J. Rudner, Garrett M. Morris, Charlotte M. Deane, and Yee Whye Teh. Context-guided diffusion for out-of-distribution molecular and protein design. *ArXiv*, abs/2407.11942, 2024. URL https://api.semanticscholar.org/CorpusID:271218010.

Yanheng Li, Haojia Dong, Xiaohan Lin, Yize Hao, Yue Xue, Jun Zhang, Yundong Wu, Jinming Zhou, and Yi Qin Gao. Molsculptor: An adaptive diffusion–evolution framework enabling generative drug design for multitarget affinity and selectivity. *Journal of Chemical Theory and Computation*, 22(10):5311–5325, 2026.

Weihao Lin, Liuchao Jin, Zhifei Jiang, Mingjing Cai, Zhihui Lai, Daniil Yurchenko, Bo Yan, Shengxi Zhou, Kostya S. Novoselov, Wei-Hsin Liao, and Shitong Fang. Machine learning-based inverse design for functional materials: Methods, challenges, and engineering applications. *Advanced Functional Materials*, 36(40):e75070, 2026. doi: https://doi.org/10.1002/adfm.75070. URL https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/adfm.75070.

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. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Ling Lu, John D Joannopoulos, and Marin Soljačić. Topological photonics. *Nature photonics*, 8 (11):821–829.

Wei Ma, Feng Cheng, and Yongmin Liu. Deep-learning-enabled on-demand design of chiral metamaterials. *ACS nano*, 12(6):6326–6334, 2018.

François Mazé and Faez Ahmed. Diffusion models beat gans on topology optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023.

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, 2018.

AkshatKumar Nigam, Robert Pollice, Gary Tom, Kjell Jorner, John Willes, Luca Thiede, Anshul Kundaje, and Alan Aspuru-Guzik. Tartarus: A benchmarking platform for realistic and practical inverse molecular design. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL https://openreview.net/forum?id=hjFQoqPIgL.

Tomoki Ozawa, Hannah M Price, Alberto Amo, Nathan Goldman, Mohammad Hafezi, Ling Lu, Mikael C Rechtsman, Dirk Schuster, Jonathan Simon, Oded Zilberberg, and Iacopo Carusotto. Topological photonics. *Reviews of Modern Physics*, 91(1):015006.

Raghunathan Ramakrishnan, Pavlo O. Dral, Matthias Rupp, and O. Anatole von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1, 2014. URL https://api.semanticscholar.org/CorpusID:15367821.

Lars Ruddigkeit, Ruud van Deursen, Lorenz C. Blum, and Jean-Louis Reymond. Enumeration of 166 billion organic small molecules in the chemical universe database gdb-17. *Journal of Chemical Information and Modeling*, 52(11):2864–2875, 10 2012. ISSN 1549-9596. doi: 10.1021/ci300415d. URL https://doi.org/10.1021/ci300415d.

Dvir Samuel, Rami Ben-Ari, Nir Darshan, Haggai Maron, and Gal Chechik. Norm-guided latent space exploration for text-to-image generation, 2023. URL https://arxiv.org/abs/2306.08687.

Benjamin Sanchez-Lengeling and Alán Aspuru-Guzik. Inverse molecular design using machine learning: Generative models for matter engineering. *Science*, 361(6400):360–365, 2018.

Bowen Song, Zecheng Zhang, Zhaoxu Luo, Jason Hu, Wei Yuan, Jing Jia, Zhengxu Tang, Guanyang Wang, and Liyue Shen. Ccs: Controllable and constrained sampling with diffusion models via initial noise perturbation. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 70516–70552. Curran Associates, Inc., 2025. doi: 10.52202/085713-2371. URL https://proceedings.neurips.cc/paper_files/paper/2025/file/663ebbb266aec936d726fba950ca693a-Paper-Conference.pdf.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021a.

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*, 2021b. URL https://openreview.net/forum?id=PxTIG12RRHS.

Yuxuan Song, Jingjing Gong, Minkai Xu, Ziyao Cao, Yanyan Lan, Stefano Ermon, Hao Zhou, and Wei-Ying Ma. Equivariant flow matching with hybrid probability transport for 3d molecule generation. *Advances in Neural Information Processing Systems*, 36:549–568, 2023.

Zhiwei Tang, Jiangweizhi Peng, Jiasheng Tang, Mingyi Hong, Fan Wang, and Tsung-Hui Chang. Inference-time alignment of diffusion models with direct noise optimization. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=JpbqiD7n9r.

Bram Wallace, Akash Gokul, Stefano Ermon, and Nikhil Naik. End-to-end diffusion latent optimization improves classifier guidance. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 7246–7256. IEEE, 2023.

Liang Wang, Chao Song, Zhiyuan Liu, Yu Rong, Qiang Liu, Shu Wu, and Liang Wang. Diffusion models for molecules: A survey of methods and tasks, 2025a. URL https://arxiv.org/abs/2502.09511.

Luran Wang, Chaoran Cheng, Yizhen Liao, Yanru Qu, and Ge Liu. Training free guided flow-matching with optimal control. In *The Thirteenth International Conference on Learning Representations*, 2025b. URL https://openreview.net/forum?id=61ss5RA1MM.

Zheng Wang, Yidong Chong, John D Joannopoulos, and Marin Soljačić. Observation of unidirectional backscattering-immune topological electromagnetic states. *Nature*, 461(7265):772–775.

Luhuan Wu, Brian L. Trippe, Christian A Naesseth, John Patrick Cunningham, and David Blei. Practical and asymptotically exact conditional sampling in diffusion models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=eWKqr1zcRv.

Di Xiao, Wang Yao, and Qian Niu. Valley-contrasting physics in graphene: magnetic moment and topological transport. *Physical Review Letters*, 99(23):236809.

Aobo Yang, Zhen Wei, Rhea P. Liem, and Pascal Fua. Physics-guided generative design with gradient-based source-space optimization on flow matching. *Structural and Multidisciplinary Optimization*, 69:136, 2026. doi: 10.1007/s00158-026-04318-6. URL https://doi.org/10.1007/s00158-026-04318-6.

Jichen Zhang, Liqun Zhao, Antonis Papachristodoulou, and Jack Umenberger. Constrained diffusers for safe planning and control, 2025. URL https://arxiv.org/abs/2506.12544.

Yuhan Zhang, Xin Xie, Bowen Song, Lixuan Chen, Liyue Shen, and Hui Deng. Label-efficient generative inverse design and design-rule discovery in freeform topological photonics, 2026. URL https://arxiv.org/abs/2609.28401.

Alex Zunger. Inverse design in search of materials with target functionalities. *Nature Reviews Chemistry*, 2(4):0121, 2018.