TRAIN SEPARATELY, COMPOSE AT SAMPLING: MULTI-PROPERTY CRYSTAL GENERATION WITH ORTHOGONAL FLOW GUIDANCE

Junjie Guo
The Chinese University of Hong Kong
1155246214@link.cuhk.edu.hk

ABSTRACT

Multi-property crystal generation is often bottlenecked by data: crystals satisfying multiple constraints jointly are rare, making it impractical to curate training sets or labels for $p(x \mid c_1, \ldots, c_m)$. We avoid this joint-data bottleneck by learning *reusable, per-property* guidance modules from *single-property* signals only. Starting from a pretrained unconditional base vector field, we fine-tune one flow model per property via an online, RL-style extension of Energy-Weighted Flow Matching (EWFM)(Zhang et al., 2025b). Specifically, we sample trajectories from the current flow, score the resulting terminal crystals with fast property evaluators, and convert these scores into threshold-shaped importance weights to reweight the flow-matching objective. This yields property-specific correction fields without differentiable property gradients and any jointly labeled multi-property dataset. In inference, we satisfy new constraint sets by composing these fields (no retraining), and we reduce cross-objective interference by projecting each residual field to be orthogonal to the base flow before aggregation.

On MP-20, composing stability and band-gap modules improves both objectives simultaneously: the formation-energy success rate increases from 0.754 to 0.924 (mean $-3.42 \rightarrow -4.35$ eV/atom), while the fraction with band gap > 3.0 eV rises from 0.042 to 0.157 (mean $0.58 \rightarrow 1.19$ eV), with only a modest drop in diversity/coverage, validating that modular per-property fields can be composed to achieve multi-constraint generation without joint data or retraining.

1 INTRODUCTION

Generative modeling for crystalline materials has emerged as a promising paradigm for inverse materials design, enabling direct sampling of crystal structures from learned data distributions rather than explicit search or optimization procedures (Xie et al., 2021; Jiao et al., 2023; Zeni et al., 2025; Miller et al., 2024; Antunes et al., 2024). In practice, this is often formulated as a conditional generation problem, where the generative model is guided toward target regions in design space (e.g., MOFs(Fu et al., 2023; Park et al., 2025), Quantum Materials(Okabe et al., 2025), and 2D materials(Xu et al., 2025)) or the property space (e.g., band gap(Zeni et al., 2025; Cao & Wang, 2025; Chen et al., 2025), density(Zhang et al., 2025a; Govindarajan et al., 2025; Chen et al., 2025), and HHI(Zeni et al., 2025; Chen et al., 2025)).

Despite the progress in single-condition generation, handling multiple conditions simultaneously is still less explored. Existing approaches to multi-conditional generation can be broadly categorized into two paradigms. The first paradigm incorporates multiple conditions directly into a single generative model(Zhang et al., 2025a; Zeni et al., 2025), which can be written as $P(x \mid A, B, \ldots)$ and typically relies on training signals that couple these conditions (e.g., data labeled under the joint set of constraints). This approach usually needs training data that already satisfies all constraints at the same time. In some cases, such examples are rare or even nonexistent, making it difficult or impossible to collect enough data to train the model. The second paradigm aggregates multiple property-specific rewards into a single scalar objective $R = \sum_i w_i R_i$ and then maximizes this objective through reinforcement learning(Chen et al., 2025; Cao & Wang, 2025), effectively performing joint multi-objective optimization. In practice, combining multiple rewards into a single scalar may cause

conflicting guidance signals, which will result in unstable optimization and sample-inefficient RL exploration. As a result, for multi-conditional material generation, both data coupling and joint RL fine-tuning become increasingly hard to scale as constraints grow more stringent.

Here, we argue that such repeated data collection and retraining is not strictly necessary. Instead of training a single model on jointly-labeled data or collapsing multiple objectives into one reward, we adopt a modular view: we learn property-specific guidance fields from single-property supervision, and compose them at sampling time to enforce multiple constraints. Under flow matching, multi-property conditioning can be modeled by composing multiple property-specific vector fields. Formally, let $v_0(t, x)$ denote the base (unconditional) field and $\{v_{\theta_k}(t, x)\}_{k=1}^m$ denote the learned property guidance fields $\{c_k\}_{k=1}^m$. We can obtain samples that satisfy multiple constraints by aggregating the corresponding vector fields together:

$$\frac{dx_t}{dt} = v_0(t, x_t) + \sum_{k=1}^{m} \lambda_k \, v_{\theta_k}(t, x_t)$$

Specifically, we adopt a decoupled vector-field view for multi-property control. First, we fine-tune separate flow-matching models for each target property using an online, RL-style variant of Energy-Weighted Flow Matching(Zhang et al., 2025b), and interpret their deviations from the base generator as property-specific correction fields. At inference time, we compose these corrections—via orthogonal projection onto the complement of the base vector field—to steer sampling toward crystals that jointly satisfy multiple constraints. This design mitigates gradient interference during training, reduces the need for jointly labeled multi-property data, and enables flexible trade-offs between objectives by tuning composition strengths.

Compared with existing methods, our approach requires only per-property training signals (either supervision or reward evaluation). New property combinations can be obtained by composing learned components, without retraining a new joint model. As a result, our framework is particularly attractive in settings where multi-property labels are scarce, expensive, or difficult to obtain.

2 Related work

2.1 Deep Generative Model for Materials

Deep generative models have become a central tool for data-driven materials discovery. A common way to categorize crystal generative models is by their representation. Structure-based models treat a crystal as atom types together with periodic lattices and atomic coordinates, and incorporate symmetry constraints such as permutation invariance, (E/SE)(3) equivariance, and periodic boundary conditions (Xie et al., 2021; Jiao et al., 2023; Zeni et al., 2025; Miller et al., 2024; Luo et al., 2025; Yang et al., 2023). CDVAE(Xie et al., 2021) is an early influential approach that combines a VAE with a noise-conditional score network decoder (Song & Ermon, 2019), generating crystals via Langevin dynamics while using periodic SE(3)-equivariant GNNs to enforce the required symmetries (Xie et al., 2021). Subsequent works further improve fidelity by learning a joint generative distribution over atom types, coordinates, and lattices (Jiao et al., 2023; Zeni et al., 2025; Luo et al., 2025; Miller et al., 2024; Yang et al., 2023). For example, MatterGen performs joint diffusion over these variables and supports broad conditional generation via adapter fine-tuning and classifier-free guidance (Zeni et al., 2025; Ho & Salimans, 2022a). Beyond diffusion, flow-based approaches such as FlowMM and CrystalFlow leverage (Riemannian/conditional) flow matching with symmetry-aware parameterizations to model the same joint space more efficiently at inference time, while enabling flexible conditional generation (Miller et al., 2024; Luo et al., 2025).

A different research line uses sequence representations, treating crystal structures as strings and performing autoregressive token prediction, e.g., CIF-based models (CrystalLLM) (Antunes et al., 2024) and alternative crystal string languages such as SLICES (Xiao et al., 2023); symmetry-aware sequences based on space groups and Wyckoff sites have also been explored (Wyckoff Transformer, CrystalFormer) (Kazeev et al., 2025; Cao et al., 2025a). LLM-based text generation for materials design has also been studied in related works (Choudhary, 2024; Gruver et al., 2024; Chen et al., 2024; Sriram et al., 2024).

2.2 CONDITIONAL GENERATION IN DIFFUSION AND FLOW MODELS

The conditional generation for diffusion models can be classified into three categories (Lai et al., 2025). The first category is classifier guidance (CG) (Ho & Salimans, 2022b), which aims to learn a classifier to guide the model toward the target condition. This can be denoted as $\nabla_x \log p_t(x_t \mid c) = \nabla_x \log p_t(x_t) + \nabla_x \log p_t(c \mid x_t)$. Then, classifier-free guidance (CFG) (Ho & Salimans, 2022a) further simplifies the CG process by treating the unconditional case as a special token $\varnothing$, and training both conditional and unconditional predictions within the same model. During sampling, CFG linearly combines the conditional and unconditional predictions to obtain the guided result. Currently, CFG is the most widely used guidance method for diffusion-based generative models (Nichol et al., 2021; Saharia et al., 2022; Rombach et al., 2022; Blattmann et al., 2023; Ramesh et al., 2022). Many following works aim to improve the efficiency of this method (Chung et al., 2024; Fan et al., 2025; Lin et al., 2024; Hong et al., 2023). Finally, the last category of methods is called training-free guidance (Chung et al., 2022; Yu et al., 2023; He et al., 2023; Ye et al., 2024; Bansal et al., 2023). The core idea for this line of research is to use an additional gradient from a differentiable scoring function to guide the sampling dynamics, which can be denoted as $\nabla_x \log p_t(x_t \mid c) = \nabla_x \log p_t(x_t) + \nabla_x f_{\text{score}}(x)$. However, this type of method requires a differentiable scoring function. Another issue is that, in diffusion models, CFG may involve a trade-off between efficiency and guidance strength compared with classifier-based guidance methods (Lai et al., 2025).

2.3 CONDITIONAL AND MULTI-PROPERTY MATERIAL GENERATION

Currently, multi-property material generation can be broadly classified into two categories. The first category formulates material generation as a conditional generation problem, e.g., conditioning on desired property values (Luo et al., 2025; Xiao et al., 2023; Zhang et al., 2025a; Zeni et al., 2025) or natural-language descriptions (Gruver et al., 2024; Sriram et al., 2024). Typically, these methods collect or construct a dataset that satisfies the target conditions (e.g., via property labels or evaluator-based scores), and then adapt an existing material generation model to support such design. A key practical requirement is to automatically obtain these supervision signals; for novel material control, this process can be particularly challenging. In multi-conditional settings, many works use two conditional inputs jointly to guide generation.

Another mainstream line of research leverages reinforcement learning for inverse material design, either by fine-tuning pre-trained generative models or by directly learning a design policy in the material space. Such RL-based approaches can be classified into three categories: (1) autoregressive token policies, which treat crystal generation as next-token prediction in a sequence model (Cao & Wang, 2025; Karpovich et al., 2024; Cao et al., 2025b); (2) diffusion/flow policies, which regard the reverse diffusion or flow process as the policy and adjust its trajectory via RL (Chen et al., 2025; Park & Walsh, 2025); and (3) direct RL on design actions, which trains an agent to perform editing actions rather than directly outputting structures (Govindarajan et al., 2025; Karpovich et al., 2024; Zamaraeva et al., 2023; Govindarajan et al., 2024).

3 METHOD

3.1 OVERVIEW AND PROBLEM SETUP

Our goal is to sample crystal structure from the multi-property conditional distribution $p(x \mid c_1, \ldots, c_m)$, where where each constraint c_k specifies a desired regime of a material property (e.g., formation energy, band gap). We utlize the Crystlflow(Luo et al., 2025) as the base model, which is a a pre-trained flow-matching model $v_0(x, t)$ for unconditional crystal generation, which defines an ODE $\frac{dx_t}{dt} = v_0(x_t, t),\ t \in [0, 1]$ mapping a simple prior $p_{\text{prior}}(x_0)$ to the data distribution $p_{\text{data}}(x_1)$ at $t = 1$. Instead of training a single large model directly on multi-property objectives, we learn a set of property-specific correction fields $\{g_k(x, t)\}_{k=1}^{K}$ on top of v_0. At inference time, we combine these fields to steer sampling towards crystals that satisfy multiple constraints simultaneously.

3.2 LEARNING SINGLE-PROPERTY GUIDANCE FIELDS

First, recall that flow matching models generation as an ordinary differential equation (ODE) over $t \in [0, 1]$:

$$\frac{\mathrm{d}x_t}{\mathrm{d}t} = v_\theta(t, x_t), \qquad t \in [0, 1] \tag{0}$$

Here, x_t denotes the state (e.g., lattice, coordinates, and atom types) at time t, and v_θ is a learnable time-dependent vector field. Formally, conditional (property-guided) generation can be understood as modifying the base vector field by adding a guidance term:

$$\frac{\mathrm{d}x_t}{\mathrm{d}t} = v_\theta(t, x_t) + \gamma \nabla_x R(x_t) \tag{1}$$

where $R(x)$ is the target reward/property and $\gamma > 0$ controls the guidance strength. However, directly using $\nabla_x R(x)$ is often impractical because (i) $R(x)$ may be non-differentiable, and (ii) even when differentiable, computing $\nabla_x R(x)$ can be prohibitively expensive.

To address this, we adopt Energy-Weighted Flow Matching (EFM)(Zhang et al., 2025b) and extend it from the offline setting to an online setting for learning the guided vector field. Specifically, we define a tilted target distribution over terminal structures x_1:

$$q(x_1) \propto p(x_1)\, \exp(-\beta\, \mathcal{E}(x_1)) \tag{2}$$

where $p(\cdot)$ denote the original data distribution, $\mathcal{E}(\cdot)$ is an energy surrogate, and $\beta > 0$ controls the strength of reweighting (often $\mathcal{E}(x) = -R(x)$).

Let $p_t(x \mid x_1)$ denote the forward interpolation path used in flow matching, and let $u_t(x \mid x_1)$ be the corresponding conditional (oracle) vector field. The (unweighted) marginal vector field is the posterior average

$$u_t(x_t) = \int p_t(x_1 \mid x_t)\, u_t(x_t \mid x_1)\, \mathrm{d}x_1 \tag{3}$$

According to formula (3), under the tilted target q_0, the corresponding energy-weighted marginal vector field becomes

$$\hat{u}_t(x) = \int q_t(x_1 \mid x_t)\, u_t(x_t \mid x_1)\, \mathrm{d}x_1 = \frac{\int p_t(x_1 \mid x_t)\, \exp(-\beta \mathcal{E}(x_1))\; u_t(x_t \mid x_1)\, \mathrm{d}x_1}{\int p_t(x_1 \mid x_t)\, \exp(-\beta \mathcal{E}(x_1))\; \mathrm{d}x_1} \tag{4}$$

Recall that the standard flow matching objective under p_{data} is

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{t\sim\text{Uniform}(0,1),\; x_1\sim p_{\text{data}},\; x_t\sim p_t(\cdot|x_1)} \left[\| v_\theta(t, x_t) - u_t(x_t \mid x_1) \|_2^2 \right] \tag{5}$$

Since

$$q(x_1) = \frac{p(x_1)\, \exp(-\beta\, \mathcal{E}(x_1))}{Z}, \qquad Z = \mathbb{E}_{x_1\sim p_{\text{data}}}[\exp(-\beta\, \mathcal{E}(x_1))] \tag{6}$$

we can write $\mathcal{L}_q$ as an importance-weighted expectation over samples from p_{data}:

$$\mathcal{L}_q(\theta) = \mathbb{E}_{t\sim\text{Uniform}(0,1),\; x_1\sim p_{\text{data}},\; x_t\sim q_t(\cdot|x_1)} \left[\frac{\exp(-\beta\, \mathcal{E}(x_1))}{Z} \, \| v_\theta(t, x_t) - u_t(x_t \mid x_1) \|_2^2 \right] \tag{7}$$

Because Z does not depend on θ, it only rescales the objective and does not affect the optimizer. Thus, up to a constant factor, minimizing $\mathcal{L}_q$ is equivalent to minimizing the energy-weighted flow matching objective

$$\mathcal{L}_{\text{EWFM}}(\theta) = \mathbb{E}_{t\sim p(t),\; x_0\sim p_{\text{data}},\; x\sim p_{t0}(\cdot|x_0)} \left[w(x_1)\; \| v_\theta(t, x_t) - u_t(x_t \mid x_1) \|_2^2 \right] \tag{8}$$

with importance weights $w(\cdot) \propto \exp(-\beta\, \mathcal{E}(\cdot))$.

In practice, given a mini-batch $\{x_1^{(i)}\}_{i=1}^B$ sampled from p_{data}, we use self-normalized weights

$$\tilde{w}_i = \frac{\exp\Big(-\beta\, \mathcal{E}(x_1^{(i)})\Big)}{\sum_{j=1}^B \exp\Big(-\beta\, \mathcal{E}(x_1^{(j)})\Big)} \tag{9}$$

yielding the empirical loss

$$\widehat{\mathcal{L}}_{\mathrm{EWFM}}(\theta) = \frac{1}{B}\sum_{i=1}^{B} \tilde{w}_i \left\| v_\theta(t_i, x_t^{(i)}) - u_{t_i}(x_t^{(i)} \mid x_0^{(i)}) \right\|_2^2, \quad t_i \sim \text{Uniform(0,1)}, \;\; x_t^{(i)} \sim p_{t_i}(\cdot \mid x_1^{(i)}) \tag{10}$$

Therefore, samples with lower $\mathcal{E}(x_0)$ (or higher reward since we define $\mathcal{E} = -R$) contribute more strongly to the training signal, leading to a learned vector field that implicitly steers trajectories toward the desired high-reward region without requiring explicit gradients $\nabla_x R(x)$.

For each individual property, we define different reward format which will be clearly defined within Appendix B.

3.3 Multi-Property Vector-Field Composition

At inference time, we combine the base vector field $v_0(t, x)$ with multiple property-specific guidance fields $\{d_i(t, x)\}$. A naive classifier-free-guidance (CFG)-style(Ho & Salimans, 2022a) addition can be written as

$$\frac{dx}{dt} = v_0(t, x) + \sum_i \lambda_i \, d_i(t, x) \tag{11}$$

but empirically, this performs poorly in the multi-conditional setting because different guidance fields often point in conflicting directions. We quantify this conflict by measuring the instantaneous angle between two fields d_A and d_B:

$$\cos\phi(t, x) = \frac{\langle d_A(t, x),\, d_B(t, x)\rangle}{\|d_A(t, x)\|_2 \, \|d_B(t, x)\|_2}, \qquad \phi(t, x) = \arccos(\cos\phi(t, x)) \tag{12}$$

and observe that $\phi(t, x)$ is typically large and unstable, indicating strong disagreement.

In order to mitigate such interference, we build an orthogonalization strategy inspired by Chae et al. (2025): for each guidance field d_i, we remove its component parallel to the base field v_0 and keep only the orthogonal correction. Concretely, define the projection

$$\mathrm{proj}_{v_0}(d_i) = \frac{\langle d_i, v_0\rangle}{\langle v_0, v_0\rangle} \, v_0, \qquad d_i^{\perp} = d_i - \mathrm{proj}_{v_0}(d_i) \tag{13}$$

and perform inference using the mixed field

$$\frac{dx}{dt} = v_0(t, x) + \sum_i \lambda_i \, d_i^{\perp}(t, x) \tag{14}$$

This preserves the main transport direction of the base generator while allowing each property model to contribute non-redundant corrections, which empirically improves multi-conditional controllability and stability.

4 Experiments

4.1 Experimental Setup

As mentioned previously, we adopt CrystalFlow (Luo et al., 2025) as our base unconditional generator $v_0(t, x)$ for crystalline materials. All experiments are conducted on the MP-20(Xie et al., 2021) dataset, following the original train/validation/test split used in Luo et al. (2025). For each crystal x we consider three scalar target properties: (i) electronic band gap (in eV), (ii) formation energy per atom (in eV/atom), and (iii) Herfindahl–Hirschman Index (HHI) over elemental production. We focus on steering the generator toward the following desirable regimes:

$$\text{band gap} > 3.0 \text{ eV}, \qquad E_{\mathrm{form}} < -2.5 \text{ eV/atom}, \qquad \mathrm{HHI} < 1500 \tag{15}$$

Based on the three targets, we design specific rewards for each of the target, the detail will be put in the Appendix B.

	Band Gap (BG)		Formation Energy (FE)		HHI	
Model	mean (eV)↑	rate↑	mean (eV/atom)↓	rate↑	mean↓	rate↑
BG only	**2.96**	**47.79**	-	-	-	-
FE only	-	-	**-4.83**	**91.60**	-	-
HHI only	-	-	-	-	**1196.39**	**91.64**
Original	0.56	4.98	-3.36	74.22	3167.97	9.25

Table 1: Single-property EWFM results on MP-20. We report the mean value and threshold success rate (Eq. (15)) for each target property. Original is the unconditional base model.

4.2 Evaluation metrics

We evaluate our models on MP-20(Xie et al., 2021) using the same family of metrics as CrystalFlow(Luo et al., 2025), and group them into two categories: target-property metrics and generation-quality metrics.

Target-property metrics: For each experiment setup, we report the mean formation energy per atom $\mathrm{FE}(x_0)$, the mean band gap $\mathrm{BG}(x_0)$, and the mean criticality score $\mathrm{HHI}(x_0)$ of generated crystals. In addition, we measure the fraction of samples that satisfy the desired constraints in equation 15: (i) FE rate, defined as the proportion of structures with $\mathrm{FE}(x_0) < E_{\mathrm{cut}}$, (ii) BG rate, defined as the proportion of structures with $\mathrm{BG}(x_0) > 3.0$ eV, and (iii) HHI rate, defined as the proportion of structures with $\mathrm{HHI}(x_0) < \mathrm{HHI}_{\mathrm{cut}}$ (e.g., $\mathrm{HHI}_{\mathrm{cut}} = 1500$).

Generation-quality metrics To assess physical plausibility, we report composition validity (fraction of chemically valid compositions within the MP-20(Xie et al., 2021) element set) and structural validity (fraction of structures that pass basic geometric sanity checks). To quantify distribution shift, we compute 1D Wasserstein distances between the generated and test sets for two scalar descriptors: density and the number of distinct elements per structure. Finally, we report coverage and diversity via coverage recall/precision and average minimum structural/compositional distances (AMSD/AMCD) between generated and test structures, together with simple generation statistics.

4.3 Single-Property EWFM Fine-Tuning

We first validate whether EWFM fine-tuning can reliably steer CrystalFlow toward a desired property regime when optimizing a single constraint at a time. As shown in Table 4.3, each single-property model achieves a substantial shift in the intended direction and improves the corresponding threshold success rate, indicating effective controllability rather than just a change in mean. Concretely, BG guidance raises the mean band gap from 0.56 to 2.96 eV and increases the fraction of samples with $\mathrm{BG}(x_0) > 3.0$ eV from 4.98% to 47.79%. FE guidance lowers the mean formation energy from -3.36 to -4.83 eV/atom and boosts the stability success rate from 74.22% to 91.60%. HHI guidance reduces the mean HHI from 3167.97 to 1196.39 while increasing the fraction below $\mathrm{HHI} < 1500$ from 9.25% to 91.64%. Overall, these results support our premise that energy-weighted reweighting with threshold-shaped rewards can learn property-specific correction fields that meaningfully reshape the terminal distribution toward the target regime.

Stronger controllability inevitably comes with a distribution shift away from the original MP-20(Xie et al., 2021) test distribution. Table 4.3 shows that the unconditional model remains the closest to the test-set statistics, achieving the smallest Wasserstein distances (e.g., 0.224 on density and 0.270 on element count) and the highest coverage recall (97.17%). In contrast, guided samplers increase Wasserstein distances and reduce coverage recall, consistent with steering toward a smaller, constraint-favorable subset of the data manifold. Importantly, structural validity stays high across all runs ($\approx$ 97.5–99.8%), suggesting that EWFM does not collapse generation into geometrically implausible structures; instead, it primarily redistributes probability mass toward the target region. We therefore interpret the reduced coverage as an expected trade-off when enforcing stringent property thresholds.

	Validity (%)↑		Coverage (%)↑		Property↓	
Model	Structural	Compositional	Recall	Precision	wdist(ρ)	wdist(N_{el})
BG only	99.77	70.53	74.66	99.63	3.315	0.762
FE only	**99.80**	88.11	69.90	98.46	1.632	1.271
HHI only	97.52	**91.72**	48.47	**99.92**	1.022	1.128
Original	99.63	81.25	**97.17**	**99.92**	**0.224**	**0.270**

Table 2: Generation-quality metrics for single-property models on MP-20: validity, coverage (recall/precision), and Wasserstein distance to the test set for density ρ and element count N_{el}.

	Band Gap (BG)		Formation Energy (FE)		HHI	
Model	mean (eV)↑	rate↑	mean (eV/atom)↓	rate↑	mean↓	rate↑
BG + FE	1.06	12.24	**-4.44**	**95.28**	-	-
FE + HHI	-	-	-3.80	87.57	2327.89	33.49
HHI + BG	**1.11**	**12.40**	-	-	**2128.78**	**35.62**
Original	0.56	4.98	-3.36	74.22	3167.97	9.25

Table 3: Two-property composition results on MP-20. We report mean values and threshold success rates (Eq. (15)) for the involved properties.

4.4 Multi-Property Vector-Field Composition

We next test whether composing two learned guidance fields can simultaneously improve multiple constraints. As shown in Table 4.4, all property pairs improve over the original model on the involved objectives, but the success rates are markedly lower than the strongest single-property results. This gap is expected: the feasible region defined by the *intersection* of two threshold constraints is substantially smaller than either single constraint alone, and the corresponding guidance fields may induce competing corrections along the sampling trajectory. For BG+FE, we obtain strong stability control (mean FE -4.44 eV/atom with a 95.28% success rate) while increasing the band-gap success rate from 4.98% to 12.24%. For HHI+BG, we simultaneously reduce criticality (HHI success rate increases from 9.25% to 35.62%) and improve band gap (BG success rate to 12.40%). For FE+HHI, we observe consistent but weaker improvements (FE success rate 87.57%, HHI success rate 33.49%), suggesting that this pair is particularly challenging under our current thresholds. These results highlight the central difficulty of multi-property crystal generation: effective single-property guidance is necessary but not sufficient, because multi-objective control is constrained by feasibility shrinkage and guidance interference.

4.5 Ablation Studies

We ablate the *instruction strength* (mixture weights) used to compose multiple guidance fields at inference time, and study the resulting trade-off between (i) target-property satisfaction and (ii) distribution shift / generation quality. Concretely, we vary the coefficients $\{\lambda_k\}$ in our composed ODE (Eq. 14) and report both property outcomes (Table 4.5) and generation-quality metrics (Table 4.5).

Varying BG-FE composition reveals a strong controllability trade-off. Across the BG/FE sweeps in Table 4.5, increasing the FE weight gradually improves stability: the mean formation energy decreases from -3.84 (BG$=0.8$, FE$=0.2$) to -5.48 (BG$=0.2$, FE$=0.8$), while the FE success rate rises from 91.98% to 99.19%. However, this stability gain comes at a steep cost in band-gap control: the BG mean drops from 2.32 to 0.40 eV and the BG>3.0 eV success rate collapses from 36.27% to 0.03% as the FE weight increases. Intermediate settings (e.g., BG$=0.6$, FE$=0.4$) yield a more balanced outcome (BG mean 1.54 eV; FE mean -4.15 eV/atom), indicating that the composition weights provide a practical way to balance the multi-objective trade-off, but also highlighting that strong control of both objectives remains challenging.

	Validity (%)↑		Coverage (%)↑		Property↓	
Model	Structural	Compositional	Recall	Precision	wdist(ρ)	wdist(N_{el})
HHI + BG	99.94	75.90	92.27	99.77	2.359	1.335
BG + FE	99.92	79.49	93.39	99.22	1.178	1.380
HHI + FE	**99.96**	**84.88**	84.83	99.61	0.753	1.466
Original	99.63	81.25	**97.17**	**99.92**	**0.177**	**0.279**

Table 4: Generation-quality metrics for two-property compositions on MP-20 (same metrics as Table 4.3).

	Band Gap (BG)		Formation Energy (FE)		HHI	
Model	mean (eV)↑	rate↑	mean (eV/atom)↓	rate↑	mean↓	rate↑
CFG (BG=0.5, FE=0.5)	1.05	10.31	-3.38	86.50	-	-
BG(0.33)+FE(0.33)+HHI(0.34)	0.82	0.07	-3.85	89.68	2708.99	19.16
BG(0.8)+FE(0.2)	2.32	36.27	-3.84	91.98	-	-
BG(0.6)+FE(0.4)	1.54	20.40	-4.15	93.62	-	-
BG(0.4)+FE(0.6)	0.72	0.07	-4.82	97.15	-	-
BG(0.2)+FE(0.8)	0.40	0.03	-5.48	99.19	-	-

Table 5: Ablation of composition weights for band gap-formation energy (and one three-way setting). We report mean values and threshold success rates (Eq. (15)). CFG means classifier-free guidance.

CFG baseline vs orthogonalized vector-field composition The CFG-style baseline (BG$=0.5$, FE$=0.5$) provides only modest multi-objective improvement (BG rate 10.31%, FE rate 86.50% in Table 4.5) and exhibits lower coverage recall (80.09% in Table 4.5) compared to most of our decomposed compositions. This supports our motivation that CFG-style inference can be suboptimal in multi-property settings, whereas separating and composing learned guidance components enables more controllable trade-offs.

Three-way composition remains difficult. Finally, composing BG, FE, and HHI together (BG$=0.33$, FE$=0.33$, HHI$=0.34$) yields simultaneous improvements on FE and partial improvement on HHI (Table 4.5), with high generation validity and strong coverage (Table 4.5). However, the overall success rates remain far from the corresponding single-property optima, underscoring the central challenge of multi-property crystal generation: as constraints accumulate, feasibility shrinkage and guidance interference can dominate, motivating more advanced composition and scheduling strategies beyond fixed global weights.

5 Discussion and Limitations

Our method may be useful in application scenarios where high-quality training data are extremely limited, such as materials discovery problems with rare or difficult-to-obtain labels. By learning property-specific guidance separately, the framework does not require large amounts of jointly labeled multi-property data.

However, in this work we evaluate generated structures mainly using machine-learning-based property predictors. These predictors provide a practical way to compare different methods at scale, but they cannot fully guarantee physical validity.

A necessary next step is to perform large-scale density functional theory (DFT) calculations to verify the stability and properties of the generated crystals, and to assess whether the observed improvements persist under first-principles evaluation.

We also observe that multi-property generation generally leads to weaker performance on individual properties compared to single-property optimization. This degradation reflects the trade-offs intro-

	Validity (%)↑		Coverage (%)↑		Property↓	
Model	Structural	Compositional	Recall	Precision	wdist(ρ)	wdist(N_{el})
CFG (BG=0.5, FE=0.5)	99.77	77.56	80.09	99.04	1.170	1.498
BG(0.33)+FE(0.33)+HHI(0.34)	99.88	80.16	92.94	99.61	1.362	1.504
BG(0.2)+FE(0.8)	99.92	**86.88**	88.60	99.16	1.106	1.357
BG(0.4)+FE(0.6)	**99.96**	82.21	**93.82**	99.08	**0.459**	1.360
BG(0.6)+FE(0.4)	99.94	77.77	93.79	99.39	1.889	1.224
BG(0.8)+FE(0.2)	99.94	75.59	92.41	**99.69**	2.702	**0.974**

Table 6: Ablation of composition weights: validity, coverage, and Wasserstein distances on MP-20 (same metrics as Table 4.3); includes CFG baseline.

duced when enforcing multiple constraints simultaneously. How to better preserve strong single-property performance while satisfying multiple objectives remains an open problem and will be the focus of future work.

6 Conclusion

We presented a decoupled vector-field framework for multi-property crystal generation with flow matching. Our main contributions are: (1) an online, RL-style extension of Energy-Weighted Flow Matching to fine-tune separate property-specific flow models using threshold-shaped rewards, producing reusable correction fields on top of an unconditional base generator; and (2) an inference-time composition rule that orthogonalizes each correction against the base vector field, reducing redundant updates and mitigating interference in multi-conditional sampling.

On the MP-20 benchmark, the learned single-property fields achieve substantial shifts toward the desired regimes (band gap, formation energy, and HHI), and composing two fields improves both involved objectives relative to the original model. The results suggest that modular per-property learning plus principled composition can reduce reliance on jointly labeled multi-property data and avoid unstable optimization from collapsing multiple objectives into a single scalar reward. Remaining challenges include stronger three-way (or higher-order) control under strict thresholds and rigorous validation beyond ML predictors; addressing these will likely require more advanced conflict-aware composition and first-principles evaluation.

Code availability: Code and trained checkpoints will be released publicly after completing the ongoing work.

REFERENCES

Luis M Antunes, Keith T Butler, and Ricardo Grau-Crespo. Crystal structure generation with autoregressive large language modeling. *Nature Communications*, 15(1):10570, 2024.

Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 843–852, 2023.

Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.

Zhendong Cao and Lei Wang. Crystalformer-rl: Reinforcement fine-tuning for materials design. *arXiv preprint arXiv:2504.02367*, 2025.

Zhendong Cao, Xiaoshan Luo, Jian Lv, and Lei Wang. Space group informed transformer for crystalline materials generation. *Science Bulletin*, 2025a.

Zhendong Cao, Shigang Ou, and Lei Wang. Crystalformer-csp: Thinking fast and slow for crystal structure prediction. *arXiv preprint arXiv:2512.18251*, 2025b.

JungWoo Chae, Jiyoon Kim, and Sangheum Hwang. Parallel rescaling: Rebalancing consistency guidance for personalized diffusion models. *arXiv preprint arXiv:2506.00607*, 2025.

Chi Chen, Weike Ye, Yunxing Zuo, Chen Zheng, and Shyue Ping Ong. Graph networks as a universal machine learning framework for molecules and crystals. *Chemistry of Materials*, 31(9):3564–3572, 2019.

Junwu Chen, Jeff Guo, Edvin Fako, and Philippe Schwaller. Accelerating inverse materials design using generative diffusion models with reinforcement learning. *arXiv preprint arXiv:2511.03112*, 2025.

Yan Chen, Xueru Wang, Xiaobin Deng, Yilun Liu, Xi Chen, Yunwei Zhang, Lei Wang, and Hang Xiao. Mattergpt: A generative transformer for multi-property inverse design of solid-state materials. *arXiv preprint arXiv:2408.07608*, 2024.

Kamal Choudhary. Atomgpt: Atomistic generative pretrained transformer for forward and inverse materials design. *The Journal of Physical Chemistry Letters*, 15(27):6909–6917, 2024.

Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022.

Hyungjin Chung, Jeongsol Kim, Geon Yeong Park, Hyelin Nam, and Jong Chul Ye. Cfg++: Manifold-constrained classifier free guidance for diffusion models. *arXiv preprint arXiv:2406.08070*, 2024.

Weichen Fan, Amber Yijia Zheng, Raymond A Yeh, and Ziwei Liu. Cfg-zero*: Improved classifier-free guidance for flow matching models. *arXiv preprint arXiv:2503.18886*, 2025.

Xiang Fu, Tian Xie, Andrew S Rosen, Tommi Jaakkola, and Jake Smith. Mofdiff: Coarse-grained diffusion for metal-organic framework design. *arXiv preprint arXiv:2310.10732*, 2023.

Prashant Govindarajan, Santiago Miret, Jarrid Rector-Brooks, Mariano Phielipp, Janarthanan Rajendran, and Sarath Chandar. Learning conditional policies for crystal design using offline reinforcement learning. *Digital Discovery*, 3(4):769–785, 2024.

Prashant Govindarajan, Mathieu Reymond, Antoine Clavaud, Mariano Phielipp, Santiago Miret, and Sarath Chandar. Crystalgym: A new benchmark for materials discovery using reinforcement learning. *arXiv preprint arXiv:2509.23156*, 2025.

Nate Gruver, Anuroop Sriram, Andrea Madotto, Andrew Gordon Wilson, C Lawrence Zitnick, and Zachary Ulissi. Fine-tuned language models generate stable inorganic materials as text. *arXiv preprint arXiv:2402.04379*, 2024.

Yutong He, Naoki Murata, Chieh-Hsin Lai, Yuhta Takida, Toshimitsu Uesaka, Dongjun Kim, Wei-Hsiang Liao, Yuki Mitsufuji, J Zico Kolter, Ruslan Salakhutdinov, et al. Manifold preserving guided diffusion. *arXiv preprint arXiv:2311.16424*, 2023.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022a.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022b.

Susung Hong, Gyuseong Lee, Wooseok Jang, and Seungryong Kim. Improving sample quality of diffusion models using self-attention guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7462–7471, 2023.

Rui Jiao, Wenbing Huang, Peijia Lin, Jiaqi Han, Pin Chen, Yutong Lu, and Yang Liu. Crystal structure prediction by joint equivariant diffusion. *Advances in Neural Information Processing Systems*, 36:17464–17497, 2023.

Christopher Karpovich, Elton Pan, and Elsa A Olivetti. Deep reinforcement learning for inverse inorganic materials design. *npj Computational Materials*, 10(1):287, 2024.

Nikita Kazeev, Wei Nong, Ignat Romanov, Ruiming Zhu, Andrey Ustyuzhanin, Shuya Yamazaki, and Kedar Hippalgaonkar. Wyckoff transformer: Generation of symmetric crystals. *arXiv preprint arXiv:2503.02407*, 2025.

Chieh-Hsin Lai, Yang Song, Dongjun Kim, Yuki Mitsufuji, and Stefano Ermon. The principles of diffusion models. *arXiv preprint arXiv:2510.21890*, 2025.

Shanchuan Lin, Bingchen Liu, Jiashi Li, and Xiao Yang. Common diffusion noise schedules and sample steps are flawed. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 5404–5411, 2024.

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):9267, 2025.

Benjamin Kurt Miller, Ricky TQ Chen, Anuroop Sriram, and Brandon M Wood. Flowmm: Generating materials with riemannian flow matching. *arXiv preprint arXiv:2406.04713*, 2024.

Mark Neumann, James Gin, Benjamin Rhodes, Steven Bennett, Zhiyi Li, Hitarth Choubisa, Arthur Hussey, and Jonathan Godwin. Orb: A fast, scalable neural network potential. *arXiv preprint arXiv:2410.22570*, 2024.

Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.

Ryotaro Okabe, Mouyang Cheng, Abhijatmedhi Chotrattanapituk, Manasi Mandal, Kiran Mak, Denisse Córdova Carrizales, Nguyen Tuan Hung, Xiang Fu, Bowen Han, Yao Wang, et al. Structural constraint integration in a generative model for the discovery of quantum materials. *Nature Materials*, pp. 1–8, 2025.

Hyunsoo Park and Aron Walsh. Guiding generative models to uncover diverse and novel crystals via reinforcement learning. *arXiv preprint arXiv:2511.07158*, 2025.

Junkil Park, Youhan Lee, and Jihan Kim. Multi-modal conditional diffusion model using signed distance functions for metal-organic frameworks generation. *Nature Communications*, 16(1):34, 2025.

Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.

Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022.

Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.

Anuroop Sriram, Benjamin K Miller, Ricky T Chen, and Brandon M Wood. Flowllm: Flow matching for material generation with large language models as base distributions. *Advances in Neural Information Processing Systems*, 37:46025–46046, 2024.

Hang Xiao, Rong Li, Xiaoyang Shi, Yan Chen, Liangliang Zhu, Xi Chen, and Lei Wang. An invertible, invariant crystal representation for inverse design of solid-state materials using generative deep learning. *Nature Communications*, 14(1):7027, 2023.

Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi Jaakkola. Crystal diffusion variational autoencoder for periodic material generation. *arXiv preprint arXiv:2110.06197*, 2021.

Shihang Xu, Shibing Chu, Rami Mrad, Zhejun Zhang, Zhelin Li, Runxian Jiao, and Yuanping Chen. Discovery of 2d materials via symmetry-constrained diffusion model. *The Journal of Physical Chemistry C*, 129(14):6794–6802, 2025.

Sherry Yang, KwangHwan Cho, Amil Merchant, Pieter Abbeel, Dale Schuurmans, Igor Mordatch, and Ekin Dogus Cubuk. Scalable diffusion for materials generation. *arXiv preprint arXiv:2311.09235*, 2023.

Haotian Ye, Haowei Lin, Jiaqi Han, Minkai Xu, Sheng Liu, Yitao Liang, Jianzhu Ma, James Y Zou, and Stefano Ermon. Tfg: Unified training-free guidance for diffusion models. *Advances in Neural Information Processing Systems*, 37:22370–22417, 2024.

Jiwen Yu, Yinhuai Wang, Chen Zhao, Bernard Ghanem, and Jian Zhang. Freedom: Training-free energy-guided conditional diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 23174–23184, 2023.

Elena Zamaraeva, Christopher M Collins, Dmytro Antypov, Vladimir V Gusev, Rahul Savani, Matthew S Dyer, George R Darling, Igor Potapov, Matthew J Rosseinsky, and Paul G Spirakis. Reinforcement learning in crystal structure prediction. *Digital Discovery*, 2(6):1831–1840, 2023.

Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Zilong Wang, Aliaksandra Shysheya, Jonathan Crabbé, Shoko Ueda, et al. A generative model for inorganic materials design. *Nature*, 639(8055):624–632, 2025.

Gongbo Zhang, Yanting Li, Renqian Luo, Pipi Hu, Zeru Zhao, Lingbo Li, Guoqing Liu, Zun Wang, Ran Bi, Kaiyuan Gao, et al. Unigenx: Unified generation of sequence and structure with autoregressive diffusion. *arXiv preprint arXiv:2503.06687*, 2025a.

Shiyuan Zhang, Weitong Zhang, and Quanquan Gu. Energy-weighted flow matching for offline reinforcement learning. *arXiv preprint arXiv:2503.04975*, 2025b.

A Overall algorithm

Algorithm 1 Online EWFM fine-tuning for a single property k

Require: Base field $v_0(t, x)$; initialize $v_{\theta_k} \leftarrow v_0$.
Require: Property-specific reward $r_k(\cdot)$; batch size B; steps N_{iter}; temperature β_k; ϵ.
for $n = 1$ to N_{iter} **do**
  Sample $x_0^{1:B} \sim p_{\text{prior}}, \;\; t^{1:B} \sim p(t)$
  Roll out $\dot{x} = v_0(t, x)$ from $0 \to 1$ for each x_0^i; collect $(x_{t_i}^i, \; x_1^i)$
  $r_i \leftarrow r_k(x_1^i)$ ▷ property-k reward, see App. B
  $\tilde{r} \leftarrow \dfrac{r - \text{mean}(r)}{\text{std}(r) + \epsilon}$
  $\tilde{w} \leftarrow \text{softmax}(\beta_k \tilde{r})$
  $v_i^\star \leftarrow v^\star(t_i, x_{t_i}^i; x_0^i, x_1^i)$
  $\mathcal{L} \leftarrow \sum_{i=1}^B \tilde{w}_i \left\| v_{\theta_k}(t_i, x_{t_i}^i) - v_i^\star \right\|_2^2$
  Update θ_k by one optimizer step to minimize $\mathcal{L}$
end for
return fine-tuned field v_{θ_k}

Algorithm 2 Multi-property sampling via orthogonalized composition

Require: Base field $v_0(t, x)$; residuals $\{g_k(t, x)\}_{k \in \mathcal{K}}$; weights $\{\lambda_k\}$; ϵ.
Sample $x_0 \sim p_{\text{prior}}$
$g_k^\perp(t, x) \leftarrow g_k(t, x) - \dfrac{\langle g_k(t, x), v_0(t, x) \rangle}{\langle v_0(t, x), v_0(t, x) \rangle + \epsilon} \, v_0(t, x)$
$\tilde{v}(t, x) \leftarrow v_0(t, x) + \sum_{k \in \mathcal{K}} \lambda_k \, g_k^\perp(t, x)$
Integrate $\dot{x} = \tilde{v}(t, x)$ from $t = 0 \to 1$ with initial x_0 to obtain x_1
return x_1

B Implementation detail

Property evaluators. We use three fast surrogates to score generated crystals x: **(i) Formation energy (FE).** We use ORB's(Neumann et al., 2024) pretrained force field (`pretrained.orb_v2`) to predict total energy $E_{\text{tot}}(x)$ and compute formation energy per atom via tabulated reference energies (`REFERENCE_ENERGIES["vasp"]`):

$$\widehat{E}_{\text{form}}(x) = \frac{E_{\text{tot}}(x) - \sum_{a=1}^N E_{\text{ref}}(Z_a)}{N}.$$

(ii) Band gap (BG). We use MatGL's pretrained MEGNet(Chen et al., 2019) band-gap model (`MEGNet-MP-2019.4.1-BandGap-mfi`) to predict $\widehat{\text{BG}}(x)$ in eV; structures containing elements outside the model's supported element set are marked as NaN and dropped. **(iii) HHI.** We compute $\widehat{\text{HHI}}(x)$ using `pymatgen.analysis.hhi.HHIModel` (reserve-mode in our runs) applied to the composition.

Reward shaping (used by EWFM). We convert each property value into a scalar reward $r(x) \in [0, 1]$:

$$r_{\text{BG}}(x) = \sigma\left(\frac{\widehat{\text{BG}}(x) - E_{\text{cut}}}{\Delta_E}\right), \quad E_{\text{cut}} = 3.0, \; \Delta_E = 1.0,$$

$$r_{\text{HHI}}(x) = \exp\left(-\frac{\max(\widehat{\text{HHI}}(x) - H_{\text{cut}}, 0)}{\Delta_H}\right), \quad H_{\text{cut}} = 1500, \; \Delta_H = 300,$$

$$r_{\text{FE}}(x) = \text{clip}\left(\frac{E_{\max} - \widehat{E}_{\text{form}}(x)}{E_{\max} - E_{\min}}, 0, 1\right), \quad E_{\min} = -5.5, \; E_{\max} = -2.0,$$

and we scale FE rewards by a constant factor in code ($r_{\mathrm{FE}} \leftarrow 1.5\, r_{\mathrm{FE}}$) before computing EWFM weights.

Online EWFM fine-tuning details. Each outer iteration performs: (1) **Sampling:** draw $M =$ `batch_size` $\times$ `batch_num` structures from the current flow model using the same ODE sampler as CrystalFlow (`ode_int_steps`=100 in our runs). (2) **Validity filtering:** keep only structures passing CrystalFlow's validity check. (3) **(Optional) ORB relaxation:** for band-gap training we additionally run an ASE–ORB relaxation (`FIRE` with ionic + optional cell relaxation; we use `fmax`=0.05, `max_steps`=100, stress cutoff `smax`=5 GPa; non-converged structures are discarded). (4) **Reward computation:** evaluate the property model, compute $r(x)$, and attach it to each graph. (5) **EWFM reweighting:** in each training mini-batch, standardize rewards and exponentiate:

$$\tilde{r}_i = \frac{r_i - \mu_B}{s_B + \epsilon}, \qquad w_i = \exp(\beta\, \tilde{r}_i), \qquad \hat{w}_i = \frac{w_i}{\frac{1}{B}\sum_{j=1}^{B} w_j + \epsilon},$$

with $\beta = 1.0$ and $\epsilon = 10^{-8}$. We reweight CrystalFlow's per-graph flow-matching loss as

$$\mathcal{L}_{\mathrm{EWFM}} = \frac{\sum_{i=1}^{B} \hat{w}_i\, \mathcal{L}_{\mathrm{FM},i}}{\sum_{i=1}^{B} \hat{w}_i + \epsilon} \;+\; \mathcal{L}_{\mathrm{reg}} \;+\; \lambda_{\mathrm{KL}} \sum_{p} \|p - p^{(0)}\|_2^2,$$

where $\mathcal{L}_{\mathrm{reg}}$ is the model's regularization term returned by the CrystalFlow forward pass, and the last term is an L2 penalty to keep parameters p close to their initialization $p^{(0)}$. We optimize with AdamW (weight decay 0), gradient clipping at 1.0, and train for 10 epochs per outer iteration. In our notebooks, FE/HHI use $\mathtt{lr} = 10^{-4},\ \lambda_{\mathrm{KL}} = 10^{-3}$; BG uses $\mathtt{lr} \approx 3 \times 10^{-5},\ \lambda_{\mathrm{KL}} \approx 3 \times 10^{-3}$ (or resumed from checkpoint).

Orthogonalized vector-field composition (code-level). For multi-property sampling we treat each fine-tuned model as a separate vector field and compose them at sampling time. Given a base model v_0 and a property-tuned model v_k, we form the residual $\Delta v_k = v_k - v_0$ and optionally remove its component parallel to v_0:

$$\Delta v_k^{\perp} = \Delta v_k - \frac{\langle \Delta v_k,\ v_0 \rangle}{\langle v_0,\ v_0 \rangle + \epsilon}\, v_0.$$

In our implementation we apply this projection separately to (a) lattice updates (per-crystal inner products) and (b) coordinate/type updates (inner products aggregated per crystal by summing over atoms). At each integration step we use

$$v_{\mathrm{mix}} = v_0 + \sum_k \lambda_k\, \Delta v_k^{\perp},$$

and generate the ODE forward from $t = 0$ to $t = 1$ to generate the final crystal sample.