

ENFORCING CONSTRAINTS IN MOLECULAR AND CRYSTALLINE GENERATIVE MODELS VIA PHYSICS-CONSTRAINED FLOW MATCHING

Pengfei Cai¹ Utkarsh Utkarsh^{1*} Nofit Segal^{1*} Akshay Subramanian^{1*}

Katharina Jager¹ Elton Pan¹ Alan Edelman¹ Christopher Vincent Rackauckas¹

Rafael Gomez-Bombarelli^{1†}

¹Massachusetts Institute of Technology

ABSTRACT

Pretrained flow-based generative models for molecules and crystals often violate hard geometric constraints at inference time, despite being trained on valid data. Herein, we extend Physics-Constrained Flow Matching (PCFM) to atomistic generative models and demonstrate it in two representative settings: (i) enforcing bond length, aromatic planarity, E/Z double bond stereochemistry, and R/S tetrahedral chirality constraints for conformer sampling, and (ii) enforcing lattice system constraints in crystal structure prediction. On GEOM-DRUGS, PCFM enforces bond length bounds and aromatic planarity while preserving ET-Flow recall and precision, and improves stereochemical pass rates by up to 14.5%. On MP-20, conditional lattice correction with PCFM reduces unit cell mismatches and increases FlowMM’s crystal structure match rate to 74.3%. Overall, PCFM turns pretrained flow matching models into constrained samplers for molecular and crystalline generation, without finetuning or architectural changes, with broader applications to molecular and materials design currently underway.

1 INTRODUCTION

Generative models have become promising methods for atomistic modeling, enabling continuous-time sampling of molecular conformers and crystalline structures. In molecules, diffusion- and flow-based approaches achieve high conformer quality and geometric fidelity, while in crystalline materials, related methods jointly model atom types, fractional coordinates, and lattice parameters for crystal structure prediction (CSP). While these models are trained to match data distributions, many such models do not explicitly enforce specific geometric constraints during sampling. In practice, it is often desirable to steer generation towards structures with prescribed geometries, such as conformers with stereochemistry compatible with a binding site, or crystal structures constrained to a given unit cell geometry. However, current samplers follow the learned vector field without enforcing constraints, allowing violations such as stereochemistry flips or inconsistent unit cell geometry to occur at inference time. This motivates methods that guide pretrained generative models towards constraint-satisfying regions of the configuration space at inference time, without finetuning the underlying model.

We build on Physics-Constrained Flow Matching (PCFM) (Utkarsh et al., 2025), which introduces inference-time correction for pretrained flow matching models via constraint-aware projections, originally applied to continuum systems governed by partial differential equations (PDEs). Our approach is complementary to symmetry-aware generators: rather than modifying training or architecture, PCFM enforces user-specified constraints at inference time on pretrained models. We

*These authors contributed equally.

†Corresponding author: rafagb@mit.edu

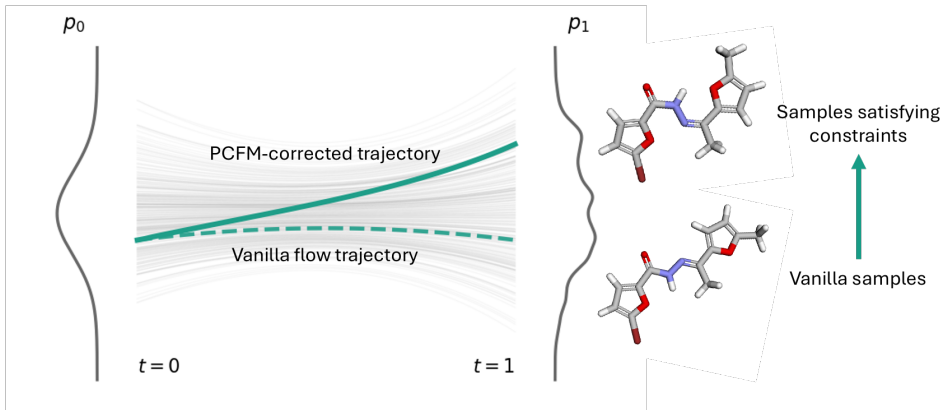


Figure 1: Steering samples generated by a pretrained flow matching model towards constraint-satisfying molecular or crystalline structures at inference time.

demonstrate PCFM on two atomistic generative settings as representative testbeds: molecular conformer generation with ET-Flow (Hassan et al., 2024) and CSP with FlowMM (Miller et al., 2024). Our focus is *constraint enforcement during inference*: PCFM does not finetune or modify the base model, and instead corrects the trajectory on-the-fly.

Contributions. (i) We apply PCFM on ET-Flow to enforce bond length, aromatic planarity, E/Z double bond stereochemistry, and R/S tetrahedral chirality constraints derived from molecular graphs and stereochemical annotations, aligning with recent geometric validity evaluations (Buttenschoen et al., 2024; Ishitani & Moriwaki, 2025). (ii) We apply conditional lattice system corrections to FlowMM that projects the lattice metric to satisfy lattice system equalities while leaving atom types and fractional coordinates unaltered, complementing recent works in CSP that are symmetry-aware (Jiao et al., 2024; Puny et al., 2025; Levy et al., 2025; Kazeev et al., 2025; Cao et al., 2025). While PCFM can be extended to enforce full space group constraints, we focus on lattice system as a first step in this work. (iii) Empirically, PCFM improves constraint satisfaction on GEOM-DRUGS (Axelrod & Gómez-Bombarelli, 2022) and reduces lattice mismatches on MP-20 (Jain et al., 2013; Xie et al., 2022), while largely preserving or improving standard quality metrics (i.e. COV/AMR for conformers, match rate for CSP).

2 RESULTS

PCFM consistently improves constraint satisfaction. Tables 1–4 evaluate constraint satisfaction with two metrics: *Pass Rate*, the fraction of generated conformers satisfying all constraints of a given type, and *Residual*, the mean constraint violation, averaged over all eligible conformers. In our experiments herein, each PCFM variant enforces a single constraint type at inference time. Reference conformers correspond to the GEOM-DRUGS test set’s conformers used for CovMat evaluation. Across all constraint types, PCFM consistently increases Pass Rate and drives Residuals towards numerical precision. For bond length bounds, PCFM-Bond raises Pass Rate to 0.994. For aromatic planarity, PCFM-Planar achieves near-reference Pass Rate. For E/Z double bond stereochemistry, PCFM-EZ improves E/Z consistency pass rate from 0.764 to 0.909 without affecting recall and precision. For R/S stereochemistry, PCFM-RS improves the R/S pass rate without relying on post-hoc parity correction.

PoseBusters validity checks corroborate constraint evaluations. PoseBusters (Buttenschoen et al., 2024) validity checks (Table 7) independently confirm that PCFM-EZ and PCFM-RS improve E/Z double bond stereochemistry and R/S tetrahedral chirality respectively, while bond length and planarity checks remain near-saturated across methods.

PCFM maintains learned conformer distributions. Table 5 reports conformer quality on GEOM-DRUGS using the CovMat metrics used in ET-Flow (Coverage and AMR at $\delta = 0.75 \text{ \AA}$). To compare fairly between PCFM and vanilla baselines, we reran ET-Flow and ET-Flow-SS on the

Table 1: Bond length constraints satisfaction (means only, $N_{\text{bond}} = 998$).

| Method | Pass Rate \uparrow | Residual \downarrow |
|------------|----------------------|-----------------------|
| Reference | 0.981 | 3.27×10^{-5} |
| ET-Flow | 0.980 | 3.24×10^{-5} |
| ET-Flow-SS | 0.980 | 3.75×10^{-5} |
| PCFM-Bond | 0.994 | 1.15×10^{-9} |

Table 3: E/Z stereochemistry (means only, $N_{\text{EZ}} = 180$).

| Method | Pass Rate \uparrow | Residual \downarrow |
|------------|----------------------|-----------------------|
| Reference | 0.906 | 1.08×10^{-1} |
| ET-Flow-SS | 0.764 | 3.03×10^{-1} |
| ET-Flow | 0.764 | 3.03×10^{-1} |
| PCFM-EZ | 0.909 | 3.41×10^{-3} |

Table 2: Aromatic ring planarity satisfaction (means only, $N_{\text{pl}} = 977$).

| Method | Pass Rate \uparrow | Residual \downarrow |
|-------------|----------------------|-----------------------|
| Reference | 0.886 | 2.41×10^{-4} |
| ET-Flow | 0.849 | 2.44×10^{-4} |
| ET-Flow-SS | 0.869 | 2.23×10^{-4} |
| PCFM-Planar | 0.882 | 4.13×10^{-6} |

Table 4: R/S stereochemistry (means only, $N_{\text{RS}} = 564$). PCFM-RS omits ET-Flow’s post-hoc chirality correction to isolate the effect of PCFM.

| Method | Pass Rate \uparrow | Residual \downarrow |
|------------|----------------------|-----------------------|
| Reference | 0.858 | 1.62×10^{-3} |
| ET-Flow | 0.727 | 1.22×10^{-1} |
| ET-Flow-SS | 0.729 | 1.22×10^{-1} |
| PCFM-RS | 0.860 | 4.41×10^{-4} |

same seed. Across planarity, E/Z, and bond length constraints, PCFM maintains both recall and precision COV/AMR at levels comparable to or slightly better than the ET-Flow baselines, indicating that constraint enforcement does not distort the learned conformer distribution. For bond-length constraints, the selective PCFM variant most closely matches the baseline metrics, consistent with bond violations occurring primarily in the tail of the distribution. For R/S tetrahedral chirality, we disable ET-Flow’s post-hoc parity correction; in this stricter setting, PCFM-RS improves chirality (Table 4) at the cost of lower COV and higher AMR.

Constraining lattice system increases match rate by reducing lattice mismatch. On MP-20, Table 6 shows that applying lattice-only PCFM to FlowMM increases match rate from 0.688 to 0.743 and similarly improves METRe and cRMSE metrics (Martirosyan et al., 2025). PCFM corrects the unit cell geometry while leaving atom types and fractional coordinates unchanged, converting a subset of previously lattice-mismatch dominated failures into site-dominated failures (Tables 6, 8).

3 METHODS

Physics-Constrained Flow Matching. We consider a pretrained flow matching model that generates samples by integrating a learned time-dependent vector field $\frac{dx(t)}{dt} = v_{\theta}(x(t), t)$, $t \in [0, 1]$, from a prior distribution $x(0) \sim p_0$, using an explicit Euler ODE solver with $\Delta t = 1/N$. Here $x(t)$ denotes the intermediate denoised state (e.g. molecular coordinates or a crystal state on a product manifold). We aim to generate samples that satisfy constraint residuals $h(x(1)) = 0$ at inference time. Many atomistic constraints are naturally expressed as inequalities; we rewrite these as hinge (ReLU) residuals, extending the PCFM framework that was demonstrated for equality constraints in previous work. PCFM (Utkarsh et al., 2025) augments flow-based sampling with a lightweight inference-time correction based on endpoint look-ahead and projection. At each discrete integration step t_k , with state $x(t_k)$, an endpoint estimate is obtained as $\hat{x}(1) = x(t_k) + (1 - t_k) v_{\theta}(x(t_k), t_k)$ and a single Gauss–Newton projection is applied to reduce the constraint residual, yielding a corrected endpoint $\tilde{x}(1)$. Let $x \in \mathbb{R}^{3n}$ denote the flattened coordinate vector and let $J = \nabla_x h(x)$ be the Jacobian of the constraint residual. The projection step involves $x \leftarrow x - J^{\top} (JJ^{\top} + \varepsilon I)^{-1} h(x)$. We refer to *always-on* PCFM when this projection is applied at every integration step, and to *selective* PCFM when projection is skipped whenever $h(\hat{x}(1)) \leq \varepsilon_{\text{skip}}$.

Molecular geometry constraints. We use ET-Flow (Hassan et al., 2024) as the base molecular conformer sampler, which learns a conditional equivariant vector field over Cartesian atom coordinates. For ODE experiments we integrate with explicit Euler using $N = 50$ steps, we also reran and

Table 5: Conformer generation quality of ET-Flow and ET-Flow with PCFM on GEOM-DRUGS (means only, $\delta = 0.75 \text{ \AA}$). All methods, including PCFM inference, use Euler integration over 50 steps except for ET-Flow-SS which includes a stochastic sampler introduced in the original paper.

| Method | Recall COV \uparrow | Recall AMR \downarrow | Precision COV \uparrow | Precision AMR \downarrow |
|---|-----------------------|-------------------------|--------------------------|----------------------------|
| <i>Reported in ET-Flow paper</i> | | | | |
| ET-Flow | 0.795 | 0.452 | 0.744 | 0.541 |
| ET-Flow-SS | 0.796 | 0.439 | 0.752 | 0.517 |
| <i>Our re-runs (same checkpoint and seed)</i> | | | | |
| ET-Flow | <u>0.799</u> | 0.458 | 0.734 | 0.562 |
| ET-Flow-SS | <u>0.799</u> | <u>0.456</u> | 0.735 | 0.559 |
| PCFM-Planar | 0.795 | 0.457 | <u>0.742</u> | <u>0.549</u> |
| PCFM-BOND (always-on) | 0.792 | 0.461 | 0.750 | 0.542 |
| PCFM-BOND (selective) | 0.800 | 0.453 | <u>0.742</u> | 0.553 |
| PCFM-EZ | <u>0.799</u> | 0.458 | 0.734 | 0.561 |
| PCFM-RS | 0.771 | 0.487 | 0.644 | 0.672 |

Table 6: Performance on the MP-20 test split ($N = 9046$). All methods use ODE sampling with 1000 integration steps and annealed vector fields (anneal slope = 10). For PCFM, lattice-only PCFM is applied as a correction every 10 steps. **Lattice-fail** denotes failure to match unit cell geometry under StructureMatcher. **Site-fail (cond.)** denotes match failure *given that the lattice is matched*.

| Method | Match rate \uparrow | METRe \uparrow | RMSD \downarrow | cRMSE \downarrow | Lattice-fail \downarrow | Site-fail (cond.) \downarrow |
|---------------|-----------------------|------------------|-------------------|--------------------|---------------------------|--------------------------------|
| FlowMM | 0.688 | 0.695 | 0.054 | 0.190 | 0.213 | 0.101 |
| FlowMM + PCFM | 0.743 | 0.750 | 0.066 | 0.175 | 0.112 | 0.145 |

report the ET-Flow-SS stochastic sampler from the original work. PCFM is only applied on top of the ET-Flow ODE sampler, noting that further work to integrate stochastic sampling is also possible. All molecular constraints are expressed as residuals $h(x)$ constructed from the current intermediate molecular graphs and the true stereochemical annotations.

Bond length bounds. For each bonded atom pair (i, j) , RDKit’s DistanceGeometry provides relaxed lower and upper bounds $[L_h, U_h]$ defining a chemically admissible bond length interval. Let $d_{ij} = \|x_i - x_j\|$, violations outside the bounds are penalized using hinge residuals:

$$h_{ij}^{\text{bond}}(x) = \left[\max\left(0, \frac{L_h - d_{ij}}{L_h}\right), \max\left(0, \frac{d_{ij} - U_h}{U_h}\right) \right].$$

Aromatic planarity. For each aromatic ring R (we consider 5- and 6-membered aromatic rings), we compute the best-fit plane with unit normal n_R and centroid \bar{x}_R . Planarity is enforced by requiring each ring atom to lie within a tolerance ε_{pl} of the plane:

$$h_k^{\text{pl}}(x) = \max\left(0, |(x_k - \bar{x}_R)^\top n_R| - \varepsilon_{\text{pl}}\right)$$

with $\varepsilon_{\text{pl}} = 0.021 \text{ \AA}$.

E/Z double bond stereochemistry. For each labeled double bond quadruple (a, i, j, b) with target configuration $\text{tgt} \in \{\pm 1\}$, we compute the cosine of the dihedral angle $\cos \phi(a, i, j, b)$. E/Z stereochemistry is then enforced via the residual

$$h^{\text{EZ}}(x) = \max\left(0, m_{\text{EZ}} - \text{tgt} \cdot \cos \phi(a, i, j, b)\right),$$

with margin $m_{\text{EZ}} = 0.5$.

R/S tetrahedral chirality. For each tetrahedral stereocenter c with neighbors (p_1, p_2, p_3) and stereochemical label $\text{tag} \in \{\pm 1\}$, we define the normalized signed volume

$$\tau(x) = \frac{(x_{p_1} - x_c)^\top ((x_{p_2} - x_c) \times (x_{p_3} - x_c))}{\|x_{p_1} - x_c\| \|x_{p_2} - x_c\| \|x_{p_3} - x_c\|}.$$

R/S chirality is then enforced using the residual

$$h^{\text{RS}}(x) = \max(0, m_{\text{RS}} - \text{tag} \cdot \tau(x)),$$

with margin $m_{\text{RS}} = 0.62$.

Lattice system constraints. We build PCFM as an inference-time correction method on top of FlowMM (Miller et al., 2024) and treat the lattice system (i.e., the unit cell geometry) as a provided condition. A crystal is represented as (a, f, L) , where a denotes atom types, $f \in [0, 1)^{n \times 3}$ are fractional coordinates, and $L \in \mathbb{R}^{3 \times 3}$ is the lattice matrix with basis vectors as rows. The unit cell geometry is fully characterized by the lattice metric tensor $G := LL^T \in \mathbb{R}_{\text{sym}}^{3 \times 3}$ which encodes unit cell lengths and angles. Constraints associated with a lattice system can therefore be expressed as linear equalities on G . FlowMM evolves a product-manifold state $x(t) = (a(t), f(t), \ell(t)) \in \mathcal{M}$ by integrating a learned time-dependent vector field. At selected steps during inference, we apply *lattice-only* PCFM by mapping the intermediate predicted lattice to a conventional setting, forming the corresponding metric \hat{G} , and projecting it onto the constraint set associated with its lattice system to obtain G_{corr} . The corrected lattice is mapped back and converted to lattice parameters used in FlowMM, while $a(t), f(t)$ remain unchanged.

4 CONCLUSION

We demonstrated that Physics-Constrained Flow Matching enables zero-shot enforcement of molecular and crystalline constraints in pretrained flow-based generative models. By correcting sampling trajectories, PCFM improves bond length validity, aromatic planarity, E/Z double bond stereochemistry, and R/S tetrahedral chirality in molecules, and improves lattice system match in crystals, while preserving generative quality. More broadly, PCFM provides a modular framework for inference-time constraint enforcement. In this work, we enforce and evaluate one constraint family at a time. Simultaneous multi-constraint enforcement and richer crystallographic symmetry constraints are important directions in our ongoing work.

5 ACKNOWLEDGMENTS

We acknowledge the MIT SuperCloud and Lincoln Laboratory Supercomputing Center for providing HPC resources. P.C. is supported by the Regenerative Energy-Efficient Manufacturing of Thermoset Polymeric Materials (REMAT) Energy Frontier Research Center, funded by the U.S. Department of Energy, Office of Science, Basic Energy Sciences under award DE-SC0023457.

REFERENCES

- Simon Axelrod and Rafael Gómez-Bombarelli. Geom, energy-annotated molecular conformations for property prediction and molecular generation. *Scientific Data*, 9(1):185, 2022. doi: 10.1038/s41597-022-01288-4. URL <https://doi.org/10.1038/s41597-022-01288-4>.
- Martin Buttenschoen, Garrett M. Morris, and Charlotte M. Deane. Posebusters: Ai-based docking methods fail to generate physically valid poses or generalise to novel sequences. *Chemical Science*, 15(9):3130–3139, 2024. doi: 10.1039/D3SC04185A.
- Zhendong Cao, Xiaoshan Luo, Jian Lv, and Lei Wang. Space group informed transformer for crystalline materials generation. *Science Bulletin*, 2025.
- Majdi Hassan, Nikhil Shenoy, Jungyoon Lee, Hannes Stark, Stephan Thaler, and Dominique Beaini. Et-flow: Equivariant flow-matching for molecular conformer generation, 2024. URL <https://arxiv.org/abs/2410.22388>.
- Ryuichiro Ishitani and Yoshitaka Moriwaki. Improving stereochemical limitations in protein–ligand complex structure prediction. *ACS Omega*, 10(46):56075–56084, 2025. doi: 10.1021/acsomega.5c07675.
- Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, and Kristin A. Persson. The materials project: A materials genome approach to accelerating materials innovation. *APL Materials*, 1(1):011002, 2013. doi: 10.1063/1.4812323. URL <https://doi.org/10.1063/1.4812323>.
- Rui Jiao, Wenbing Huang, Yu Liu, Deli Zhao, and Yang Liu. Space group constrained crystal generation. *arXiv preprint arXiv:2402.03992*, 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.
- Daniel Levy, Siba Smarak Panigrahi, Sékou-Oumar Kaba, Qiang Zhu, Kin Long Kelvin Lee, Mikhail Galkin, Santiago Miret, and Siamak Ravanbakhsh. Symmcd: Symmetry-preserving crystal generation with diffusion models. *arXiv preprint arXiv:2502.03638*, 2025.
- Maya M. Martirosyan, Thomas Egg, Philipp Hoellmer, George Karypis, Mark Transtrum, Adrian Roitberg, Mingjie Liu, Richard G. Hennig, Ellad B. Tadmor, and Stefano Martiniani. All that structure matches does not glitter, 2025. URL <https://arxiv.org/abs/2509.12178>.
- Benjamin Kurt Miller, Ricky T. Q. Chen, Anuroop Sriram, and Brandon M. Wood. Flowmm: Generating materials with riemannian flow matching, 2024. URL <https://arxiv.org/abs/2406.04713>.
- Omri Puny, Yaron Lipman, and Benjamin Kurt Miller. Space group conditional flow matching, 2025. URL <https://arxiv.org/abs/2509.23822>.
- Utkarsh Utkarsh, Pengfei Cai, Alan Edelman, Rafael Gomez-Bombarelli, and Christopher Vincent Rackauckas. Physics-constrained flow matching: Sampling generative models with hard constraints, 2025. URL <https://arxiv.org/abs/2506.04171>.
- Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi Jaakkola. Crystal diffusion variational autoencoder for periodic material generation, 2022. URL <https://arxiv.org/abs/2110.06197>.

6 APPENDIX

Evaluation metrics. For molecular conformer generation on GEOM-DRUGS, we report recall and precision metrics (COV and AMR at $\delta = 0.75 \text{ \AA}$; Table 5) in accordance to ET-Flow’s implementation, together with constraint-specific Pass Rate and Residual metrics (Tables 1–4). We evaluate molecular validity with PoseBusters and evaluation results are shown in Table 7. For crystal structure prediction on MP-20, we evaluate structural validity using pymatgen StructureMatcher’s match rate and RMSD, and additionally report METRe and cRMSE introduced recently (Martirosyan et al., 2025).

Table 7: PoseBusters evaluation on GEOM-DRUGS conformers ($K \leq 32$ per molecule). Values are mean pass rates over molecules. Molecular geometry checks are computed with PoseBusters ($N = 997$), while stereochemistry identities (E/Z and R/S) are evaluated using PoseBusters’ InChI-based identity checks with ETKDG-generated conformers as reference structures (E/Z: $N = 128$, R/S: $N = 301$).

| Method | Mol checks (mean pass rate) | | | Identity (mean pass rate) | |
|----------------|-----------------------------|------------------------|----------------------|-----------------------------|-----------------------|
| | Bond length | Aromatic ring flatness | Double bond flatness | Double bond stereochemistry | Tetrahedral chirality |
| ET-Flow-SS | 0.9992 | 1.0000 | 0.9944 | 0.8421 | 0.8032 |
| PCFM-Bond | 1.0000 | 1.0000 | 0.9943 | 0.8462 | 0.8038 |
| PCFM-Planarity | 0.9908 | 1.0000 | 0.9937 | 0.8462 | 0.8034 |
| PCFM-EZ | 0.9975 | 0.9997 | 0.9801 | 0.8848 | 0.8032 |
| PCFM-RS | 0.8727 | 0.9957 | 0.9842 | 0.8388 | 0.9556 |

Table 8: Failure mode decomposition under StructureMatcher for MP-20 dataset. **Site-only fail** indicates atomic mismatch given a matched lattice. **Both fail** indicates cases where lattice mismatch prevents site-level matching. Rates are reported over all test samples.

| Method | Site-only fail | Both lattice + site fail | Total fail rate |
|---------------|----------------|--------------------------|-----------------|
| FlowMM | 0.101 | 0.211 | 0.312 |
| FlowMM + PCFM | 0.145 | 0.111 | 0.257 |