
SYNTHESIS-CONSTRAINED DISCRETE DIFFUSION FOR IONIZABLE LIPID GENERATION

Rohin Maganti^{1,3,†}, **Rahul Maganti**^{2,†}, **Mohamad-Gabriel Alameh**^{3,4,*}

¹Department of Bioengineering, University of Pennsylvania, Philadelphia, PA, USA

²Children’s Hospital of Philadelphia, Philadelphia, PA, USA

³Penn Institute for RNA Innovation, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA, USA

⁴Department of Medicine, University of Pennsylvania, Philadelphia, PA, USA

†These authors contributed equally

*Correspondence: mg.alameh@pennmedicine.upenn.edu

ABSTRACT

Ionizable lipids are the critical component of lipid nanoparticles for mRNA delivery, yet their discovery remains bottlenecked by library enumeration. Existing machine learning approaches can rank pre-defined candidates but cannot generate novel structures. We introduce synthesis-constrained diffusion, the first deep generative model for ionizable lipids, embedding combinatorial chemistry constraints directly into scaffold-conditioned generation. Our proof of concept enforces Ugi scaffold integrity by construction: core bonds formed by the reaction mechanism are fixed throughout diffusion, while region-aware noise distributions capture the distinct chemistry of ionizable heads versus lipophilic tails. A three-stage curriculum (pretraining on drug-like molecules, domain adaptation on virtual lipids, and property-conditioned fine-tuning) enables learning from limited experimental data. Demonstrating this framework on Ugi-based lipid synthesis, 99% of generated samples are chemically valid with intact scaffolds and 62% are novel. The top candidate achieves $2\times$ higher predicted transfection than the training mean (in silico).

1 INTRODUCTION

Lipid nanoparticles (LNPs) are the leading delivery vehicle for messenger RNA (mRNA), enabling the rapid deployment of COVID-19 vaccines and clinical development of gene editing therapies (Polack et al., 2020; Baden et al., 2021; Hou et al., 2021; Buschmann et al., 2021). The ionizable lipid (iL) component is critical to LNP function, forming stable complexes with mRNA while evading immune clearance and enabling endosomal escape, yet discovering effective iLs remains a fundamental bottleneck (Paunovska et al., 2022; Cullis & Hope, 2017; Carrasco et al., 2021).

Current machine learning approaches for lipid discovery are exclusively discriminative: given a candidate molecule, they predict its activity (Xu et al., 2024; Witten et al., 2025; Chan et al., 2025). These models rank molecules from pre-enumerated libraries but cannot propose novel structures, and synthesizable lipid space vastly exceeds any library.

Generative models offer an alternative by learning molecular distributions and sampling novel structures (Gómez-Bombarelli et al., 2018; Jin et al., 2018; Vignac et al., 2023), but standard approaches fail on iLs. Their chemical grammar (tertiary amine heads, extended carbon chains, branched tails, ester/amide linkages) differs from drug-like molecules dominating public datasets. However, lipid synthesis via combinatorial reactions like the Ugi three-component reaction (Ugi, 1959; Xu et al., 2024), among others, imposes deterministic structural constraints: the reaction mechanism fixes specific core bonds. We exploit this determinism to embed synthesis constraints directly into the generative process.

Concurrent work explores automated lipid design through building-block combinatorics. Zhao et al. (2024) use Monte Carlo tree search to navigate fragment libraries, while Ou et al. (2024) learn poli-

cies over pre-defined building blocks. These approaches inherit the enumeration bottleneck: they can only propose combinations of known fragments. In contrast, we introduce the first deep generative model for iLs, learning the distribution of molecular structures directly rather than searching over fragment combinations.

We propose *synthesis-constrained diffusion*, a framework that embeds combinatorial chemistry constraints directly into discrete diffusion over molecular graphs. Core bonds formed by the reaction mechanism are fixed throughout the denoising process, guaranteeing that every output contains a valid synthetic scaffold. Region-aware noise distributions respect the distinct chemistry of ionizable heads and lipophilic tails, while classifier-free guidance enables property-conditioned generation. A three-stage curriculum addresses data scarcity: pretraining on drug-like molecules, domain adaptation on virtual lipids, and fine-tuning on $\sim 1,100$ experimental measurements.

On Ugi-based lipid synthesis, 99% of generated samples are chemically valid with intact scaffolds and 62% are novel, with the top candidate reaching $2\times$ higher predicted transfection than the training mean.

2 METHODS

2.1 PRELIMINARIES

Ionizable lipids for mRNA delivery. LNPs encapsulate mRNA for cellular delivery, with the iL governing efficacy. These molecules become cationic in endosomes (acidic pH) but neutral under physiological conditions (pH ~ 7.4), enabling mRNA complexation and endosomal escape while minimizing toxicity (Cullis & Hope, 2017). Structurally, iLs comprise a nitrogen-containing head group, lipophilic tails, and linker regions. Measured transfection potency (mTP) quantifies delivery efficiency. We evaluate candidates using LANTERN (Mehradfar et al., 2025), a graph neural network trained on experimental transfection data that predicts mTP from molecular structure ($R^2 \approx 0.7$ on held-out data). We define a generated iL as *valid* if it is chemically valid (correct valence, parseable structure) and contains an intact Ugi scaffold. Predicted mTP is then evaluated separately as a measure of expected efficacy.

Molecular graph representation. We represent molecules as graphs $G = (X, E)$ where $X \in \{1, \dots, K\}^n$ encodes atom types (n atoms, K atom categories) and $E \in \{0, \dots, B\}^{n \times n}$ encodes bond types (B bond categories plus "no bond"). Each atom is assigned to a functional region $r(i) \in \mathcal{R} = \{\text{HEAD}, \text{CORE}, \text{TAIL}_1, \text{TAIL}_2\}$ based on its relationship to the synthetic scaffold; we write TAIL when the distinction between the two tail chains is not relevant. Each edge inherits the region of its source node: $r_e(i, j) = r(i)$.

Discrete diffusion on graphs. Molecular graphs have categorical attributes (atom types and bond types) rather than continuous pixel values, requiring discrete diffusion (Austin et al., 2021; Vignac et al., 2023). The forward process independently corrupts each variable toward a prior distribution m :

$$q(G_t | G_{t-1}) = \prod_{i=1}^n \text{Cat}(x_t^{(i)}; x_{t-1}^{(i)} Q_t^{\text{atom}}) \prod_{(i,j) \in \mathcal{E}} \text{Cat}(e_t^{(ij)}; e_{t-1}^{(ij)} Q_t^{\text{bond}}) \quad (1)$$

where Q_t is the transition matrix, $\beta_t \in [0, 1]$ is the noise schedule, m is the stationary marginal distribution, and x, e are one-hot row vectors representing atom and bond types respectively. The reverse process trains a neural denoiser f_θ to predict clean types from corrupted inputs, with classifier-free guidance (Ho & Salimans, 2021) enabling conditional generation.

Ugi multicomponent reactions. The Ugi three-component reaction combines an amine, aldehyde, and isocyanide in a single pot to form a bis-amide product (Ugi et al., 2000). The reaction mechanism dictates specific bond connectivity in the product core; the same three reactants always yield the same scaffold topology. This determinism underpins combinatorial lipid libraries (Xu et al., 2024): varying the amine controls head group chemistry, varying the aldehyde and isocyanide controls tail properties, while the core bonds remain fixed by the reaction itself.

Algorithm 1 Reaction-Constrained Diffusion Sampling

Require: Denoiser f_θ , scaffold bonds E_C , core mask \mathcal{C} , region assignments $r(\cdot)$, target property y , guidance scale w , diffusion steps T

- 1: Sample G_T : atom $i \sim \text{Cat}(m_{\text{atom}}^{(r(i))})$, edge $(i, j) \sim \text{Cat}(m_{\text{bond}}^{(r_e(i,j))})$ {Region-aware init.}
 - 2: Fix core: $E_T[\mathcal{C}] \leftarrow E_C$
 - 3: **for** $t = T$ **to** 1 **do**
 - 4: $\hat{G}_0 \leftarrow f_\theta(G_t, t, \emptyset) + w \cdot (f_\theta(G_t, t, y) - f_\theta(G_t, t, \emptyset))$ {CFG in logit space}
 - 5: **for** each region $r \in \mathcal{R}$ **do**
 - 6: $G_{t-1}^{(r)} \sim q(G_{t-1}^{(r)} | G_t^{(r)}, \hat{G}_0^{(r)}; Q_t^{(r)})$ {Region-specific posterior}
 - 7: **end for**
 - 8: $E_{t-1}[\mathcal{C}] \leftarrow E_C$ {Restore scaffold}
 - 9: **end for**
 - 10: **return** G_0
-

2.2 REACTION-CONSTRAINED DIFFUSION

Combinatorial reactions like the Ugi three-component reaction impose *deterministic* constraints on product structure: the reaction mechanism fixes specific bonds in the core. We exploit this determinism through three complementary modifications to the diffusion process.

Fixing the core. The core bonds formed by Ugi chemistry are not statistical patterns to be learned but rather mechanistic invariants. We identify these bonds via substructure matching and remove them from the diffusion process entirely: during training, core positions are masked from the loss; during generation, core bonds are restored after each denoising step. Formally, let $\mathcal{C} \subset \{1, \dots, n\}$ denote the set of core atom indices. The training objective becomes:

$$\mathcal{L} = \mathbb{E}_{t, G_0, G_t} \left[\sum_{i \notin \mathcal{C}} \ell(f_\theta(G_t, t)_i, G_0^{(i)}) \right] \quad (2)$$

where ℓ is the cross-entropy loss, $f_\theta(G_t, t)_i$ is the predicted distribution over clean types for position i , $G_0^{(i)}$ is the ground-truth type, and the sum excludes core positions. We prove that scaffold restoration is exact under the factorized (per-node and per-edge) forward process of discrete diffusion (Proposition 2, Appendix D.4).

Region-aware noise. Ionizable lipids are chemically heterogeneous: nitrogen-containing heads enable endosomal escape, carbon-dominated tails drive membrane partitioning, and the reaction-formed core bridges them. Standard diffusion corrupts all atoms toward a single global marginal m , creating a mismatch between the noise distribution and the data distribution that the denoiser must overcome. We instead define region-specific marginal distributions that match the chemical composition of each functional region $r \in \{\text{HEAD}, \text{CORE}, \text{TAIL}\}$:

$$Q_t^{(r)} = (1 - \beta_t)I + \beta_t \mathbf{1}m^{(r)\top} \quad (3)$$

HEAD atoms diffuse toward a nitrogen-enriched distribution (28% N vs. 8% globally), while TAIL atoms diffuse toward a carbon-dominated one (94% C). The cumulative transition and reverse posterior retain closed-form solutions with region-specific matrices (Proposition 1, derivation in Appendix D.3).

Property-conditioned generation. The modifications above ensure validity; property conditioning steers generation toward activity. We inject target properties (e.g., transfection potency) via Feature-wise Linear Modulation (FiLM) layers (Perez et al., 2018) that modulate intermediate representations, and apply classifier-free guidance (CFG) (Ho & Salimans, 2021) at inference to strengthen the conditioning signal. The guidance scale w interpolates between unconditional sampling ($w = 0$) and strong property steering ($w > 1$), trading diversity for targeted optimization.

Algorithm 1 summarizes the sampling procedure. Since Ugi lipids span hundreds of distinct scaffold topologies, we select the input scaffold via property-weighted sampling from a library extracted from training data (Appendix J).

Table 1: Generation summary.

Metric	Value
Generated samples	100
Chemically valid	99%
Ugi scaffold intact	100%
Unique structures	91%
Novel (not in training)	62%
Mean predicted mTP	5.37
Best predicted mTP	10.55
Training mean mTP	4.85

Table 2: Generation metrics across training seeds and guidance scales. Each cell aggregates 60 samples (10 per target mTP $\in \{4, 6, 8, 10, 12, 14\}$). Validity = chemically valid with intact Ugi scaffold.

Seed	CFG	Validity (%)	Novelty (%)	Mean mTP
0	1.0	96.7	60.4	5.54
0	2.0	93.3	60.2	5.58
0	3.0	96.7	65.4	5.61
1	1.0	96.7	70.2	5.04
1	2.0	93.3	63.4	5.69
1	3.0	76.7	74.2	4.77
2	1.0	98.3	70.6	4.62
2	2.0	98.3	60.9	5.41
2	3.0	96.7	72.6	5.33

2.3 THREE-STAGE CURRICULUM

Deep generative models typically require tens of thousands of training examples, yet the experimental dataset contains only $\sim 1,100$ Ugi-based ionizable lipids with transfection measurements. We address this through progressive transfer learning (Appendix A): (1) **pretraining** on ZINC250k (Sterling & Irwin, 2015) teaches chemical grammar; (2) **domain adaptation** on the AGILE virtual library (Xu et al., 2024) introduces lipid-specific patterns with region-aware diffusion and scaffold conditioning; (3) **property conditioning** on experimental data enables activity-guided generation via FiLM layers and classifier-free guidance with 15% dropout.

3 EXPERIMENTS

3.1 GENERATION AND EVALUATION

We generated 100 molecules targeting mTP=10 with classifier-free guidance (scale=2.0). Of these, 99% are chemically valid with intact Ugi scaffolds (Table 1). Across a sweep of training seeds and guidance scales (Table 2), all conditions achieved $>90\%$ validity and 100% scaffold integrity, with novelty 60–73%.

Of generated candidates, 62% are entirely novel, absent from both the 12,000-molecule AGILE virtual library and the 1,100-molecule experimental set. Using Atom-Pair fingerprints (which capture lipid topology better than Morgan fingerprints), 43% of novel molecules have nearest-neighbor similarity below 0.95, confirming they differ structurally from training examples.

3.2 ABLATION: REGION-AWARE DIFFUSION

To test region-aware noise, we trained a model with uniform (non-regional) atom-type distributions. The ablated model produced 8% invalid molecules (vs. 0% for full model) and generated **allenes/cumulenes** (C=C=C), chemically unstable motifs absent from training data, including extended chains up to C=C=C=C=C=C (Appendix L). Without the bias that tails should be lipophilic

Table 3: Top 12 generated candidates ranked by predicted mTP.

Rank	mTP	MW	LogP	TPSA	HBD	Novel
1	10.55	790	14.2	71	2	Yes
2	9.56	694	11.7	71	2	Yes
3	9.01	788	13.6	71	2	Yes
4	8.58	622	9.5	71	2	Yes
5	8.18	638	10.0	71	2	Yes
6	8.00	708	11.9	71	2	Yes
7	7.89	706	11.9	71	2	Yes
8	7.87	547	8.2	96	3	Yes
9	7.83	743	13.7	96	3	Yes
10	7.81	662	11.1	71	2	Yes
11	7.70	706	11.8	71	2	No
12	7.64	741	13.5	96	3	Yes

Table 4: Comparison to discriminative approaches.

Capability	AGILE	LiON	This work
Predict activity	Yes	Yes	Yes
Generate novel structures	No	No	Yes
Guaranteed synthesizability	Library	Library	By construction
Property-conditioned	No	No	Yes
Explore beyond enumeration	No	No	Yes

alkyl chains, the model generates implausible chemistry that scaffold conditioning alone cannot prevent.

3.3 VARIANCE DECOMPOSITION

ANOVA-style variance decomposition on training data shows scaffold identity explains **65%** of mTP variance, while within-scaffold chemical variation explains **35%** (Appendix N). This 35% represents a large chemical space that enumeration samples sparsely. Our diffusion process targets this space directly, with median within-scaffold mTP range of 6.2 points.

3.4 CONVERGENT DESIGN PRINCIPLES

Despite novelty, top candidates share common features: the top three have dibutylamine heads (NCCCN(CCCC)CCCC), a minority motif in training data, which the model appears to favor for high activity. The top 7 share TPSA=71 Å² and HBD=2 (Yang et al., 2024). Tails are C12–C16 chains, and the top candidate has an ω -unsaturated C18 tail.

3.5 COMPARISON TO DISCRIMINATIVE APPROACHES

Our generative approach addresses fundamental limitations of existing discriminative methods (Table 4). While AGILE and LiON (Witten et al., 2025) achieve strong predictive performance, they cannot propose molecules outside their pre-enumerated libraries. Our synthesis-constrained diffusion framework generates novel structures while guaranteeing synthesizability by construction, a capability absent from all prior approaches.

4 DISCUSSION

Reaction-constrained diffusion generates novel, synthesizable ionizable lipids with predicted transfection up to 2 \times the training mean. The 62% novelty rate supports the hypothesis that generative models can discover chemical matter beyond the training library, not merely interpolate within known examples. The convergence of top candidates on specific structural motifs (dibutylamine

heads, C12–C16 tails with terminal unsaturation) despite this novelty suggests the model has learned meaningful structure-activity relationships that generalize to unexplored regions of chemical space.

A critical distinction of our approach is that the generative model and the property predictor (LANTERN) are entirely independent. The diffusion model learns the distribution of molecular structures during Stages 1–2 without any property labels; only Stage 3 fine-tuning uses LANTERN-corrected mTP values as conditioning targets. At generation time, LANTERN is used solely for post-hoc evaluation of generated candidates. The model does not optimize against LANTERN’s loss surface. This decoupling means the model cannot exploit blind spots in the predictor: it proposes molecules from a learned chemical distribution, and LANTERN merely scores them after the fact. Any predictor, including future experimental measurements, can replace LANTERN without retraining the generative model.

Region-aware diffusion addresses a key failure mode of standard generative models on specialized molecular classes. Without region-specific noise distributions, the model generates chemically unstable motifs such as allenes and cumulenes, because the global marginal (78% carbon, 8% nitrogen) cannot distinguish the nitrogen-rich chemistry of ionizable heads from the carbon-dominated tails. Separate marginals for each functional region resolve this, ensuring the denoising network learns appropriate chemistry for each molecular substructure. This approach generalizes naturally to other molecular classes with compositionally distinct regions.

Scaffold conditioning represents a paradigm shift from “generate then filter” to “constrain then generate.” By fixing Ugi-reaction bonds throughout diffusion, we provide a structural guarantee of synthesizability rather than relying on post-hoc retrosynthetic analysis or synthetic accessibility scores (Coley et al., 2019). While this constrains the model to Ugi-accessible chemical space, this is a feature for practical drug discovery: every generated molecule can be synthesized via established one-pot protocols from commercial building blocks. The principle extends naturally to other multicomponent reactions (Passerini, Biginelli, Hantzsch) that impose deterministic core connectivity, each opening a distinct region of synthesizable chemical space.

Several limitations merit discussion. Our evaluation relies on LANTERN predictions ($R^2 \approx 0.7$); reported gains are *in silico* and must be validated experimentally. The fixed Ugi scaffold ensures synthesizability but inherently restricts structural diversity, as the model cannot discover fundamentally new reaction scaffolds. This is an intentional trade-off: guaranteeing synthetic accessibility requires constraining the generative space. Additionally, *in vitro* transfection correlates imperfectly with *in vivo* efficacy, which depends on formulation parameters, biodistribution, and immunogenicity factors not captured by our model.

Looking forward, active learning cycles that iterate between generating candidates, synthesizing top hits, measuring activity, and retraining could enable rapid experimental optimization. Supporting multiple scaffold types simultaneously would expand accessible chemical space while preserving synthesizability guarantees. Multi-property conditioning incorporating predicted pKa, logP, and clearance would enable more comprehensive optimization beyond transfection alone. We are currently synthesizing lead candidates for experimental validation in hepatocyte transfection and mouse biodistribution studies.

5 FIGURES

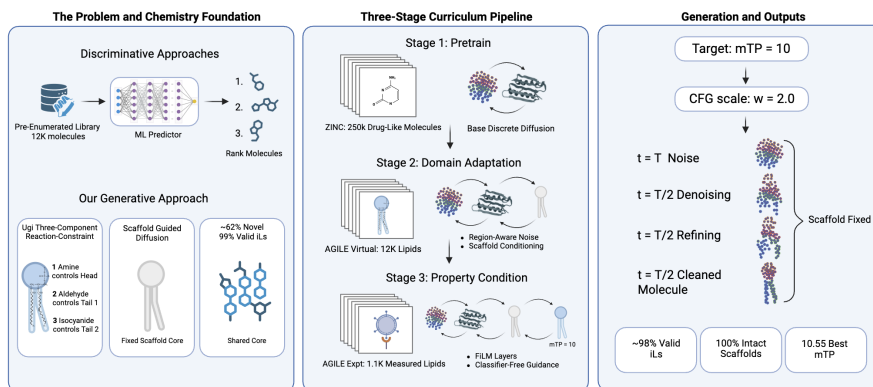


Figure 1: Reaction-constrained diffusion framework. **(A)** Problem formulation: discriminative approaches rank pre-enumerated libraries; our generative approach discovers novel structures via Ugi scaffold constraints. **(B)** Three-stage curriculum: ZINC250k \rightarrow AGILE Virtual \rightarrow AGILE Experimental. **(C)** Generation outputs: 62% novel molecules with intact scaffolds, top candidates achieving $2\times$ training mean mTP.

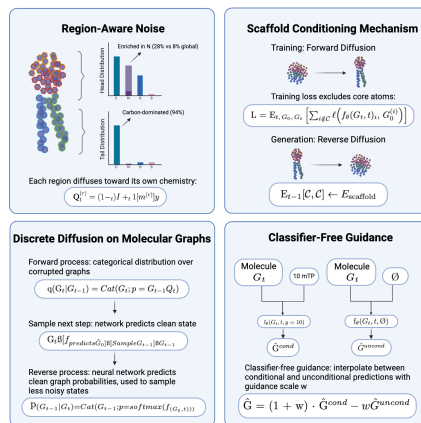


Figure 2: Technical innovations. **(A)** Region-aware noise: nitrogen-enriched heads (28% N), carbon-dominated tails (94% C). **(B)** Scaffold conditioning: core bonds fixed throughout diffusion. **(C)** Discrete diffusion on graphs with GNN denoiser. **(D)** Classifier-free guidance with 15% property dropout.

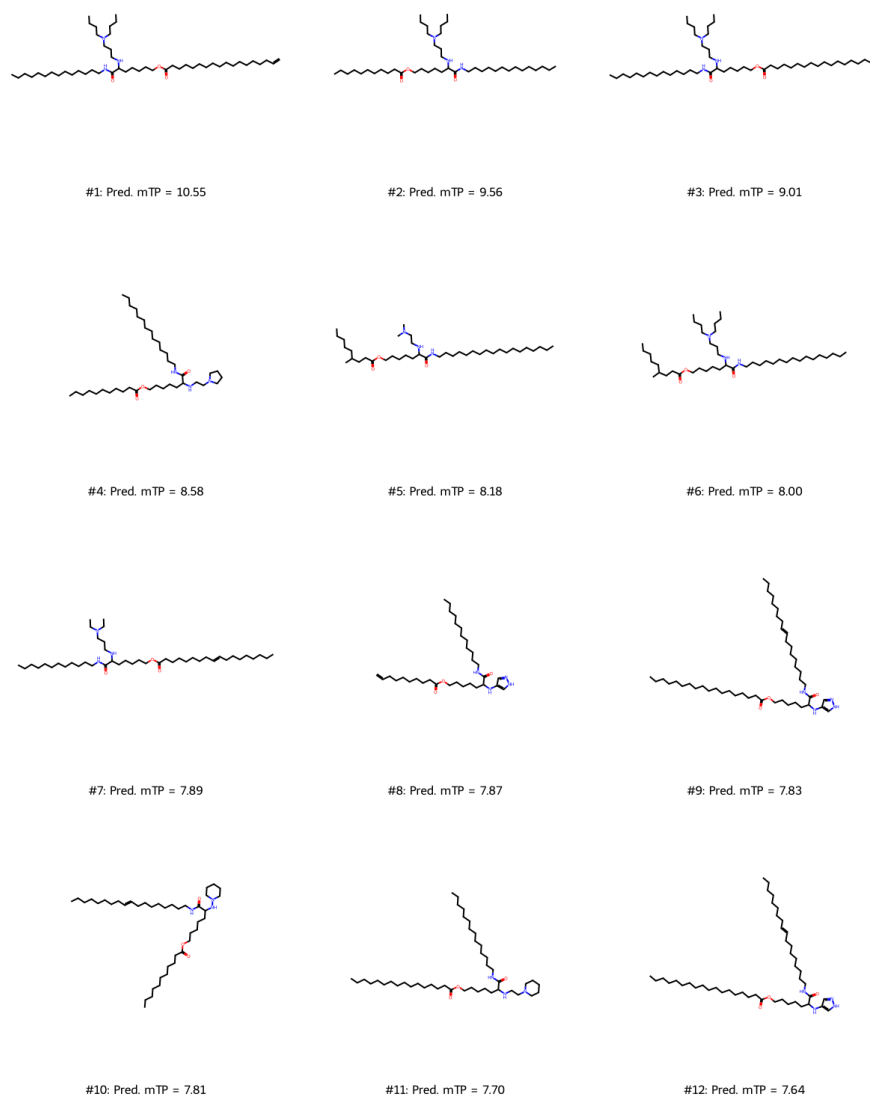


Figure 3: Top 12 candidates by predicted mTP. All contain intact Ugi scaffolds. Top three share dibutylamine heads (NCCCN (CCCC) CCCC); other head groups include pyrrolidine, piperidine, pyrazole. Best candidate: mTP=10.55 ($2\times$ training mean).

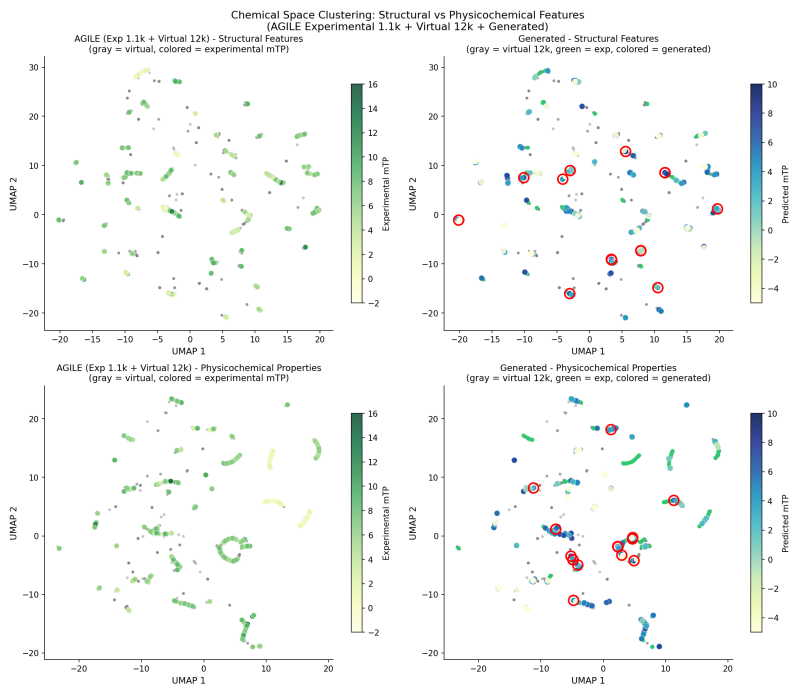


Figure 4: Chemical space analysis. UMAP embeddings using Morgan fingerprints (top) and physicochemical descriptors (bottom). Left: AGILE experimental (colored by mTP) and virtual (gray) libraries. Right: Generated molecules overlaid on training data. Red circles = highest predicted mTP. Generated candidates occupy training space while achieving high predicted activity.

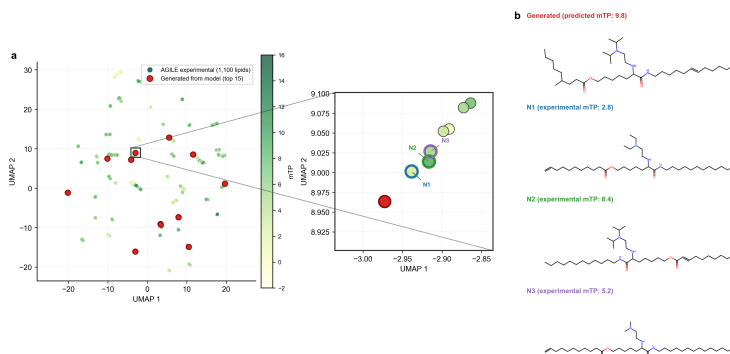


Figure 5: Nearest-neighbor analysis. **(a)** Generated molecules (red) on AGILE data; inset shows one candidate and its three nearest neighbors. **(b)** Structures of generated candidate (pred. mTP: 9.8) vs. neighbors (exp. mTP: 2.8, 8.4, 5.2). Despite structural similarity, neighbors span a 5.6-point mTP range, suggesting the model interpolates toward high-activity regions.

REFERENCES

- Jacob Austin et al. Structured denoising diffusion models in discrete state-spaces. *Proceedings of NeurIPS*, 2021.
- Lindsey R. Baden et al. Efficacy and safety of the mrna-1273 sars-cov-2 vaccine. *New England Journal of Medicine*, 384:403–416, 2021.
- Michael D. Buschmann, Manuel J. Carrasco, Suman Alishetty, Mikell Paige, Mohamad-Gabriel Alameh, and Drew Weissman. Nanomaterial delivery systems for mrna vaccines. *Vaccines*, 9(1): 65, 2021.
- Manuel J. Carrasco, Suman Alishetty, Mohamad-Gabriel Alameh, Hooda Said, Lacey Wright, Mikell Paige, Ousamah Soliman, Drew Weissman, Thomas E. Cleveland, Alexander Grishaev, and Michael D. Buschmann. Ionization and structural properties of mrna lipid nanoparticles influence expression in intramuscular and intravascular administration. *Communications Biology*, 4:956, 2021.
- Lawrence Chan, Ameya R. Kirtane, Qiaobing R. Qu, Xinhong Huang, Jaehwan Woo, Dhruva A. Subramanian, Roni Dey, Rishabh Semalty, Joshua D. Bernstock, et al. Designing lipid nanoparticles using a transformer-based neural network. *Nature Nanotechnology*, 20:1491–1501, 2025.
- C.W. Coley et al. A robotic platform for flow synthesis of organic compounds informed by ai planning. *Science*, 365:eaax1566, 2019.
- Pieter R. Cullis and Michael J. Hope. Lipid nanoparticle systems for enabling gene therapies. *Molecular Therapy*, 25:1467–1475, 2017.
- Rafael Gómez-Bombarelli, Jennifer N. Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D. Hirzel, Ryan P. Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4:268–276, 2018.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *NeurIPS Workshop*, 2021.
- Xucheng Hou, Tal Zaks, Robert Langer, and Yizhou Dong. Lipid nanoparticles for mrna delivery. *Nature Reviews Materials*, 6:1078–1094, 2021.
- Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *Proceedings of ICML*, 2018.
- Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. *Proceedings of CVPR*, 2022.
- Asal Mehradfar, Mohammad Shahab Sepehri, Jose Miguel Hernandez-Lobato, Glen S. Kwon, Mahdi Soltanolkotabi, Salman Avestimehr, and Morteza Rasoulianboroujeni. Lantern: A machine learning framework for lipid nanoparticle transfection efficiency prediction. *arXiv preprint arXiv:2507.03209*, 2025.
- Yuxuan Ou, Jingyi Zhao, Austin Tripp, Morteza Rasoulianboroujeni, and José Miguel Hernández-Lobato. A deep generative model for the design of synthesizable ionizable lipids. *NeurIPS Workshop on AI for New Drug Modalities*, 2024. arXiv:2412.00928.
- Kalina Paunovska, David Loughrey, and James E. Dahlman. Drug delivery systems for rna therapeutics. *Nature Reviews Genetics*, 23:265–280, 2022.
- Ethan Perez et al. Film: Visual reasoning with a general conditioning layer. *Proceedings of AAAI*, 2018.
- Fernando P. Polack et al. Safety and efficacy of the bnt162b2 mrna covid-19 vaccine. *New England Journal of Medicine*, 383:2603–2615, 2020.
- Teague Sterling and John J. Irwin. Zinc 15—ligand discovery for everyone. *Journal of Chemical Information and Modeling*, 55:2324–2337, 2015.

-
- I. Ugi et al. Multicomponent reactions in organic chemistry. *Angewandte Chemie International Edition*, 39:3168–3210, 2000.
- Ivar Ugi. Versuche mit isonitrilen. *Angewandte Chemie*, 71:386, 1959.
- Clement Vignac et al. Digress: Discrete denoising diffusion for graph generation. *Proceedings of ICLR*, 2023.
- Jacob Witten, Idris Raji, Rajith S. Manan, Emily Beyer, Susanna Bartlett, Yingying Tang, Mohamad Ebadi, Junhan Lei, David Nguyen, Favour Oladimeji, et al. Artificial intelligence-guided design of lipid nanoparticles for pulmonary gene therapy. *Nature Biotechnology*, 43:1790–1799, 2025.
- Yue Xu, Shihao Ma, Haotian Cui, Jingan Chen, Shufen Xu, Bo Wang, Bowen Li, et al. Agile platform: a deep learning powered approach to accelerate lnp development for mrna delivery. *Nature Communications*, 15:6305, 2024.
- Z. Yang et al. Uni-pka: An accurate and physically consistent pka prediction through protonation ensemble representation. *Journal of Chemical Information and Modeling*, 2024.
- Jingyi Zhao, Yuxuan Ou, Austin Tripp, Morteza Rasoulianboroujeni, and José Miguel Hernández-Lobato. Generative model for synthesizing ionizable lipids: A monte carlo tree search approach. *NeurIPS Workshop on AI for New Drug Modalities*, 2024. arXiv:2412.00807.

A CURRICULUM LEARNING DETAILS

Table 5: Three-stage curriculum learning.

Stage	Dataset	Size	Epochs	Features enabled
1. Pretrain	ZINC250k	250K	500	Molecular graph grammar
2. Adapt	AGILE Virtual	12K	300	Region-aware, scaffold cond.
3. Condition	AGILE Expt.	1.1K	300	FiLM, CFG

B SAMPLING OVERVIEW

We provide a complete description of the scaffold-conditioned generation procedure. Understanding this procedure is essential for reproducing our results and for designing ablation experiments.

B.1 WHAT IS A SCAFFOLD?

A *scaffold* defines the fixed structural template for generation. Formally, a scaffold consists of:

- **Number of atoms** n : The molecular size (typically 40–49 atoms for Ugi lipids)
- **Edge connectivity**: The bond topology as an adjacency list, specifying which atoms are bonded to which
- **Region assignments** $r : \{1, \dots, n\} \rightarrow \mathcal{R}$: Each atom labeled as HEAD, CORE, TAIL1, or TAIL2
- **Core bond types**: The specific bond types (single, double, etc.) for bonds involving core atoms

Two molecules share a scaffold if and only if they have identical values for all four components. The experimental training set (1,100 molecules) contains 527 unique scaffolds; the virtual library (12,276 molecules) contains 7,620 unique scaffolds.

B.2 TWO-STAGE GENERATION PIPELINE

Generation proceeds in two stages:

Stage 1: Scaffold Selection. Given a target property value (e.g., mTP=10), we select a scaffold from the library using property-weighted sampling (Section J). Scaffolds associated with training molecules whose properties are close to the target receive higher sampling probability. This stage determines the molecular topology.

Stage 2: Diffusion. Given the selected scaffold, we run the discrete diffusion denoising process to fill in:

- **Atom types**: What element each atom is (C, N, O, S, etc.)
- **Non-core bond types**: Single, double, triple, or aromatic bonds for edges not involving core atoms

Core bonds are *never* modified and remain fixed to their ground-truth values throughout diffusion (see Section D.4 for the mathematical justification).

B.3 WHAT DIFFUSION CONTROLS VS. WHAT IS FIXED

Component	Fixed by Scaffold	Sampled by Diffusion
Number of atoms	✓	
Which atoms are bonded (topology)	✓	
Region assignments	✓	
Core bond types (~15% of bonds)	✓	
Atom types (C/N/O/S/etc.)		✓
Non-core bond types (~85% of bonds)		✓

B.4 WHY SCAFFOLD SELECTION MATTERS

Variance decomposition on the training data reveals that scaffold identity explains 65% of mTP variance, while the remaining 35% is explained by within-scaffold chemical variation; i.e., different atom/bond type assignments on the same topology (see Appendix N for details). This motivates property-weighted scaffold selection: by biasing toward scaffolds associated with high-mTP molecules, we leverage the dominant source of variance. The diffusion process then optimizes within the selected scaffold’s chemical space.

C SAMPLING IMPLEMENTATION DETAILS

Algorithm 1 (Section 2.2) summarizes the sampling procedure. We provide additional implementation details below.

Initialization. At $t = T$, atom and bond types are sampled independently from region-specific marginal distributions $m^{(r)}$ (Table 9). This ensures the initial noise matches the chemical composition expected in each functional region. For example, HEAD atoms are initialized with 28% probability of nitrogen versus 8% globally.

Classifier-free guidance. Both conditional and unconditional forward passes share identical architecture; the unconditional pass sets the property input to a learned null token. Guidance is applied in logit space before softmax normalization, which was empirically more stable than guidance in probability space in our experiments. The guidance scale w interpolates: $w=0$ recovers unconditional sampling, $w=1$ doubles the conditional signal, and $w>1$ amplifies it further at the cost of diversity.

Posterior sampling. At each step, the posterior $q(G_{t-1} | G_t, \hat{G}_0)$ is computed per-region using region-specific transition matrices $Q_t^{(r)}$. This ensures that the reverse process respects the same region-aware noise schedule used during training.

Scaffold restoration. After each posterior sample, core bond entries are overwritten with the ground-truth Ugi scaffold. This hard constraint is applied at every timestep, not just the final step, preventing scaffold corruption from accumulating through the reverse process.

Final step. At $t=1$, we sample from \hat{G}_0 directly via multinomial sampling (rather than argmax) to preserve diversity in the generated set.

D MATHEMATICAL DETAILS

We present the mathematical foundations underlying synthesis-constrained diffusion. We first state standard discrete diffusion results (Austin et al., 2021), then derive the region-aware and scaffold-constrained extensions that constitute our contribution.

D.1 DISCRETE DIFFUSION ON CATEGORICAL VARIABLES

Consider a categorical variable $x \in \{1, \dots, K\}$ represented as a one-hot row vector. The single-step forward transition (Austin et al., 2021) is:

$$q(x_t | x_{t-1}) = \text{Cat}(x_t; x_{t-1}Q_t), \quad Q_t = (1 - \beta_t)I + \beta_t \mathbf{1}m^\top \quad (4)$$

where $\beta_t \in [0, 1]$ is the noise rate at step t , $\mathbf{1} \in \mathbb{R}^K$ is the all-ones column vector, and $m \in \Delta^{K-1}$ is the stationary (marginal) distribution. Entry $Q_t[i, j]$ gives the probability of transitioning from category i to category j : with probability $1 - \beta_t$ the value is unchanged, and with probability β_t it is redrawn from m .

The cumulative transition matrix $\bar{Q}_t = Q_1 Q_2 \cdots Q_t$ admits a closed form. By induction on t (using the idempotency $\mathbf{1}m^\top \mathbf{1}m^\top = \mathbf{1}m^\top$, which holds because $m^\top \mathbf{1} = 1$):

$$\bar{Q}_t = \bar{\alpha}_t I + (1 - \bar{\alpha}_t) \mathbf{1}m^\top, \quad \bar{\alpha}_t = \prod_{s=1}^t (1 - \beta_s) \quad (5)$$

This gives a direct sampling formula: $q(x_t = j | x_0 = i) = \bar{\alpha}_t \delta_{ij} + (1 - \bar{\alpha}_t) m_j$. At each timestep, the variable either retains its original value (with probability $\bar{\alpha}_t$) or is replaced by a sample from m (with probability $1 - \bar{\alpha}_t$). As $t \rightarrow T$, $\bar{\alpha}_T \rightarrow 0$ and $q(x_T) \rightarrow m$.

Proof of Eq. 5. The base case $\bar{Q}_1 = Q_1$ holds by definition with $\bar{\alpha}_1 = 1 - \beta_1$. Assuming $\bar{Q}_{t-1} = \bar{\alpha}_{t-1} I + (1 - \bar{\alpha}_{t-1}) \mathbf{1}m^\top$:

$$\begin{aligned} \bar{Q}_t &= \bar{Q}_{t-1} Q_t = [\bar{\alpha}_{t-1} I + (1 - \bar{\alpha}_{t-1}) \mathbf{1}m^\top] [(1 - \beta_t) I + \beta_t \mathbf{1}m^\top] \\ &= \underbrace{(1 - \beta_t) \bar{\alpha}_{t-1}}_{\bar{\alpha}_t} I + [(1 - \beta_t)(1 - \bar{\alpha}_{t-1}) + \beta_t \bar{\alpha}_{t-1} + \beta_t (1 - \bar{\alpha}_{t-1})] \mathbf{1}m^\top \end{aligned} \quad (6)$$

where we used $\mathbf{1}m^\top \mathbf{1}m^\top = \mathbf{1}m^\top$ and $I \cdot \mathbf{1}m^\top = \mathbf{1}m^\top \cdot I = \mathbf{1}m^\top$. The coefficient of $\mathbf{1}m^\top$ simplifies to $1 - (1 - \beta_t) \bar{\alpha}_{t-1} = 1 - \bar{\alpha}_t$. \square

D.2 REGION-AWARE TRANSITION MATRICES

Ionizable lipids are chemically heterogeneous: nitrogen-containing heads, carbon-dominated tails, and a heteroatom-containing core. Standard discrete diffusion corrupts all atoms toward a single global marginal m , forcing the denoising network to invert a distribution mismatch. We instead define per-region transition matrices, with each region $r \in \mathcal{R} = \{\text{HEAD}, \text{CORE}, \text{TAIL1}, \text{TAIL2}\}$ having its own stationary distribution $m^{(r)}$.

Definition 1 (Region-aware forward process). *Let $r : \{1, \dots, n\} \rightarrow \mathcal{R}$ assign each atom to a functional region. Each edge inherits the region of its source node: $r_e(i, j) = r(i)$. The region-aware single-step transitions are:*

$$q(x_t^{(i)} | x_{t-1}^{(i)}) = \text{Cat}\left(x_t^{(i)}; x_{t-1}^{(i)} Q_t^{(r(i))}\right), \quad Q_t^{(r)} = (1 - \beta_t) I + \beta_t \mathbf{1} m^{(r)\top} \quad (7)$$

and analogously $q(e_t^{(ij)} | e_{t-1}^{(ij)}) = \text{Cat}(e_t^{(ij)}; e_{t-1}^{(ij)} Q_t^{(r_e(i,j))})$ with bond-type marginals $m_{\text{bond}}^{(r)}$.

Two properties make this well-defined:

(i) Per-variable independence. Nodes and edges are corrupted independently in the forward process, so the joint distribution factorizes:

$$q(G_t | G_{t-1}) = \prod_{i=1}^n q(x_t^{(i)} | x_{t-1}^{(i)}; r(i)) \cdot \prod_{(i,j) \in \mathcal{E}} q(e_t^{(ij)} | e_{t-1}^{(ij)}; r_e(i, j)) \quad (8)$$

This factorization is preserved regardless of whether variables share a common marginal or use region-specific ones.

(ii) Region-specific closed form. Since each $Q_t^{(r)}$ has the marginal form of Eq. 4 with m replaced by $m^{(r)}$, the proof of Eq. 5 applies identically:

$$\bar{Q}_t^{(r)} = \bar{\alpha}_t I + (1 - \bar{\alpha}_t) \mathbf{1} m^{(r)\top} \quad (9)$$

Note that $\bar{\alpha}_t$ is shared across regions, and only the stationary distribution differs. This means all regions follow the same noise schedule but converge to different limiting distributions: HEAD atoms converge to the nitrogen-enriched $m^{(\text{HEAD})}$ (28% N), TAIL atoms to the carbon-dominated $m^{(\text{TAIL})}$ (94% C).

D.3 POSTERIOR WITH REGION-SPECIFIC MATRICES

The reverse process requires the posterior $q(x_{t-1} | x_t, x_0)$. In the standard (non-region-aware) case, this is given by Bayes' rule (Austin et al., 2021):

$$q(x_{t-1} = k | x_t = j, x_0 = i) = \frac{q(x_t = j | x_{t-1} = k) q(x_{t-1} = k | x_0 = i)}{q(x_t = j | x_0 = i)} \quad (10)$$

We now show that this posterior remains tractable with region-specific matrices and derive its explicit form.

Proposition 1 (Region-aware reverse posterior). *For an atom in region r , the reverse posterior with region-specific transition matrices is:*

$$q(x_{t-1}^{(i)}=k | x_t^{(i)}=j, x_0^{(i)}=\ell; r) \propto Q_t^{(r)}[k, j] \cdot \bar{Q}_{t-1}^{(r)}[\ell, k] \quad (11)$$

The posterior has the same algebraic structure as the standard (non-region-aware) posterior, with m replaced by $m^{(r)}$, and introduces no additional computational cost.

Proof. Substituting the closed forms from Eqs. 7 and 9 into Eq. 11:

$$q(x_{t-1}=k | x_t=j, x_0=\ell; r) \propto \underbrace{[(1-\beta_t) \delta_{kj} + \beta_t m_j^{(r)}]}_{q(x_t=j|x_{t-1}=k)} \cdot \underbrace{[\bar{\alpha}_{t-1} \delta_{\ell k} + (1-\bar{\alpha}_{t-1}) m_k^{(r)}]}_{q(x_{t-1}=k|x_0=\ell)} \quad (12)$$

We expand this product into four cases.

Case 1: $k = j$ and $k = \ell$ (token unchanged through both steps).

$$[(1-\beta_t) + \beta_t m_j^{(r)}] \cdot [\bar{\alpha}_{t-1} + (1-\bar{\alpha}_{t-1}) m_k^{(r)}] \quad (13)$$

This is the dominant term when the atom retains its original type throughout, weighted by both the probability of no corruption at step t and the probability of no corruption through steps $1, \dots, t-1$.

Case 2: $k = j$ but $k \neq \ell$ (corrupted before $t-1$, stable at step t).

$$[(1-\beta_t) + \beta_t m_j^{(r)}] \cdot (1-\bar{\alpha}_{t-1}) m_k^{(r)} \quad (14)$$

The atom was replaced by a sample from $m^{(r)}$ before $t-1$ (landing on k), and then persisted unchanged at step t .

Case 3: $k \neq j$ but $k = \ell$ (unchanged until step t , then corrupted).

$$\beta_t m_j^{(r)} \cdot [\bar{\alpha}_{t-1} + (1-\bar{\alpha}_{t-1}) m_k^{(r)}] \quad (15)$$

The atom retained its original type $\ell = k$ through step $t-1$, then was redrawn from $m^{(r)}$ at step t (landing on $j \neq k$).

Case 4: $k \neq j$ and $k \neq \ell$ (corrupted at both stages).

$$\beta_t m_j^{(r)} \cdot (1-\bar{\alpha}_{t-1}) m_k^{(r)} \quad (16)$$

The atom was replaced by a sample from $m^{(r)}$ before $t-1$, then replaced again at step t .

The normalizing constant is:

$$q(x_t=j | x_0=\ell; r) = \bar{Q}_t^{(r)}[\ell, j] = \bar{\alpha}_t \delta_{\ell j} + (1 - \bar{\alpha}_t) m_j^{(r)} \quad (17)$$

In all four cases, the region marginal $m^{(r)}$ appears in exactly the same positions where m would appear in the standard posterior. The algebraic structure is identical; only the values of the marginal vector change. This means: (a) no new approximations are introduced, (b) the computational cost is $O(K)$ per atom (same as standard), and (c) the posterior is guaranteed to be a valid probability distribution after normalization. \square

Model-predicted posterior. In practice, the denoiser predicts a distribution $p_\theta(x_0 | G_t, t)$ rather than a point estimate. The posterior marginalizes over this prediction:

$$q(x_{t-1}^{(i)}=k | x_t^{(i)}, p_\theta; r) \propto Q_t^{(r)}[k, x_t^{(i)}] \cdot \sum_{\ell=1}^K \bar{Q}_{t-1}^{(r)}[\ell, k] p_\theta(x_0^{(i)}=\ell) \quad (18)$$

Substituting the closed forms:

$$\begin{aligned} q(x_{t-1}^{(i)}=k | x_t^{(i)}=j, p_\theta; r) &\propto [(1-\beta_t) \delta_{kj} + \beta_t m_j^{(r)}] \\ &\cdot \sum_{\ell=1}^K [\bar{\alpha}_{t-1} \delta_{\ell k} + (1-\bar{\alpha}_{t-1}) m_k^{(r)}] p_\theta(x_0^{(i)}=\ell) \\ &= [(1-\beta_t) \delta_{kj} + \beta_t m_j^{(r)}] \cdot [\bar{\alpha}_{t-1} p_\theta(x_0^{(i)}=k) + (1-\bar{\alpha}_{t-1}) m_k^{(r)}] \end{aligned} \quad (19)$$

where in the last step we used $\sum_{\ell} \delta_{\ell k} p_\theta(\ell) = p_\theta(k)$ and $\sum_{\ell} p_\theta(\ell) = 1$. In vectorized form, let $\hat{p} = p_\theta(x_0^{(i)}) \in \mathbb{R}^K$. The unnormalized posterior over x_{t-1} is:

$$\tilde{q}(x_{t-1}^{(i)}) = Q_t^{(r)}[:, x_t^{(i)}] \odot (\bar{Q}_{t-1}^{(r)\top} \hat{p}) \quad (20)$$

where \odot denotes element-wise multiplication. While this vectorized form suggests an $O(K^2)$ matrix-vector product, substituting the closed form of $\bar{Q}_{t-1}^{(r)}$ (Eq. 9) reduces it to $O(K)$: the product $\bar{Q}_{t-1}^{(r)\top} \hat{p} = \bar{\alpha}_{t-1} \hat{p} + (1-\bar{\alpha}_{t-1}) m^{(r)} (\hat{p}^\top \mathbf{1})$ requires only a scalar-vector product. Region-awareness adds no computational overhead.

D.4 SCAFFOLD CONDITIONING AS CONSTRAINED DIFFUSION

Let $\mathcal{C} \subset \{1, \dots, n\}$ denote core atom indices and let $\mathcal{E}_{\mathcal{C}} = \{(i, j) \in \mathcal{E} : i \in \mathcal{C} \text{ or } j \in \mathcal{C}\}$ denote the protected edge set, comprising all bonds incident to at least one core atom. Scaffold conditioning decomposes the graph into constrained and free components: $G = (G^{\mathcal{C}}, G^{-\mathcal{C}})$, where $G^{\mathcal{C}}$ is fixed at the ground-truth scaffold values throughout diffusion.

Proposition 2 (Exactness of scaffold conditioning). *If core positions \mathcal{C} are held fixed throughout the forward process ($x_t^{(i)} = x_0^{(i)}$ for all $i \in \mathcal{C}$, all t), then restoring core values after each reverse step is equivalent to exact sampling from the conditional posterior $q(G_{t-1}^{-\mathcal{C}} | G_t, G_0^{\mathcal{C}})$.*

Proof. We show this in two parts: the forward process and the reverse process.

Forward process. For core positions, the forward corruption is replaced by the identity: $q(x_t^{(i)} | x_0^{(i)}) = \delta_{x_t^{(i)}, x_0^{(i)}}$ for $i \in \mathcal{C}$. Since nodes and edges are corrupted independently (Section D.2), the joint forward process factorizes as:

$$q(G_t | G_0) = \underbrace{\prod_{i \in \mathcal{C}} \delta_{x_t^{(i)}, x_0^{(i)}}}_{\text{core: det.}} \cdot \underbrace{\prod_{i \notin \mathcal{C}} q(x_t^{(i)} | x_0^{(i)}; r(i))}_{\text{non-core: region-aware}} \cdot \underbrace{\prod_{(i,j) \in \mathcal{E}_{\mathcal{C}}} \delta_{e_t^{(ij)}, e_0^{(ij)}}}_{\text{core bonds: det.}} \cdot \underbrace{\prod_{(i,j) \notin \mathcal{E}_{\mathcal{C}}} q(e_t^{(ij)} | e_0^{(ij)}; r(i))}_{\text{non-core bonds}} \quad (21)$$

The independence structure means the marginal distribution over non-core positions $q(G_t^{-\mathcal{C}} | G_0)$ does not depend on how core positions are handled, since they are always at their true values.

Reverse process. By the same per-variable independence, the reverse posterior factorizes:

$$q(G_{t-1} | G_t, G_0) = \prod_{i \in \mathcal{C}} \delta_{x_{t-1}^{(i)}, x_0^{(i)}} \cdot \prod_{i \notin \mathcal{C}} q(x_{t-1}^{(i)} | x_t^{(i)}, x_0^{(i)}; r(i)) \cdot (\text{bond terms}) \quad (22)$$

Sampling from this posterior and then overwriting $G_{t-1}^{\mathcal{C}} \leftarrow G_0^{\mathcal{C}}$ is therefore identical to sampling directly from $q(G_{t-1}^{-\mathcal{C}} | G_t, G_0^{\mathcal{C}})$. No approximation is introduced. \square

Remark (learned vs. true posterior). Proposition 2 holds for the *true* posterior $q(G_{t-1} | G_t, G_0)$, which factorizes across variables. The learned denoiser f_θ couples all atoms through message passing, so $p_\theta(G_{t-1} | G_t)$ does not factorize. However, scaffold overwrite remains valid: core values are

deterministic regardless of the denoiser’s predictions, and overwriting them after each step simply replaces whatever the model predicted for core positions with the known ground truth. The exactness guarantee is a property of the constrained forward process, not of the learned reverse model.

Training loss. The loss excludes core positions:

$$\mathcal{L} = \mathbb{E}_{t, G_0, G_t} \left[\sum_{i \notin \mathcal{C}} \ell(f_\theta(G_t, t)_i, G_0^{(i)}) + \sum_{(i,j) \notin \mathcal{E}_c} \ell(f_\theta(G_t, t)_{ij}, E_0^{(ij)}) \right] \quad (23)$$

where ℓ is the cross-entropy loss. Core bonds are held fixed during forward corruption ($e_t^{(ij)} = e_0^{(ij)}$ for $(i, j) \in \mathcal{E}_c$), so the denoiser always observes the correct scaffold. Excluding core positions from the loss is consistent: since these positions are never corrupted, the denoiser’s predictions for them carry no training signal and should not contribute to the gradient.

Comparison to continuous inpainting. In continuous diffusion, inpainting (Lugmayr et al., 2022) requires re-noising the known region and blending it with the generated region at each step, because the known pixels must follow the same noise level as the generated ones for the posterior to be valid. In our discrete setting, this complication does not arise: core positions are deterministic (they never enter the noise process), so the hard overwrite is exact at every timestep. This is a structural advantage of our constrained formulation over analogous approaches in continuous diffusion.

E ARCHITECTURE DETAILS

E.1 GRAPH NEURAL NETWORK

We use a message-passing neural network (MPNN) based on the DiGress architecture (Vignac et al., 2023). The network operates on molecular graphs with the following feature dimensions:

- **Atom features** (26-dim): One-hot encoding of atom type (10 types: C, N, O, S, F, Cl, Br, P, I, other) + formal charge (5) + hybridization (4) + aromaticity (1) + ring membership (1) + region encoding (3: HEAD/CORE/TAIL) + hydrogen count (2)
- **Bond features** (10-dim): One-hot encoding of bond type (5: none, single, double, triple, aromatic) + ring membership (1) + conjugation (1) + stereo (3)
- **Global features** (8-dim): Number of atoms, number of bonds, molecular weight (binned), target mTP (continuous), timestep embedding (4-dim sinusoidal)

E.2 NETWORK ARCHITECTURE

Table 6: Network architecture specifications.

Component	Specification
Message passing layers	8
Hidden dimension	256
Attention heads	8
Dropout	0.1
Normalization	LayerNorm
Activation	SiLU
FiLM layers	After layers 2, 4, 6, 8
Total parameters	$\sim 7.2\text{M}$

E.3 FiLM CONDITIONING

Property conditioning is implemented via Feature-wise Linear Modulation (FiLM) (Perez et al., 2018). The target mTP value y is encoded through a 2-layer MLP to produce scale (γ) and shift (β) parameters:

$$\text{FiLM}(h, y) = \gamma(y) \odot h + \beta(y) \quad (24)$$

where h is the hidden representation at each FiLM layer. The output layer of each FiLM module is initialized to $\gamma = 1, \beta = 0$ (identity transform), so that at the start of Stage 3 the model behaves identically to its Stage 2 checkpoint before property conditioning takes effect. During training with classifier-free guidance, we replace y with a null token with probability 0.15.

F DATASET DETAILS

Table 7: Dataset statistics and train/validation splits.

Stage	Dataset	Total	Train	Val
1	ZINC250k	250,000	237,500 (95%)	12,500 (5%)
2	AGILE Virtual	12,000	10,800 (90%)	1,200 (10%)
3	AGILE Experimental	1,100	990 (90%)	110 (10%)

ZINC250k (Sterling & Irwin, 2015) contains 250,000 drug-like molecules from the ZINC database, filtered for molecular weight 200–500 Da and $\text{LogP} < 5$. This dataset teaches general chemical grammar without lipid-specific patterns.

AGILE Virtual Library (Xu et al., 2024) is a publicly available dataset containing 12,000 ionizable lipids enumerated via Ugi three-component reactions from curated building block sets. All molecules share the characteristic Ugi scaffold and span diverse head groups and tail chemistries.

AGILE Experimental (Xu et al., 2024) is a publicly available dataset containing $\sim 1,100$ ionizable lipids with measured transfection potency (mTP) from high-throughput screening. mTP values are log-transformed luminescence measurements; higher values indicate more effective mRNA delivery.

All splits use a fixed random seed (42) for reproducibility. Validation sets are used for early stopping and model selection; no separate test set is held out since our goal is generation rather than prediction.

G TRAINING HYPERPARAMETERS

Table 8: Training hyperparameters for each curriculum stage.

Hyperparameter	Stage 1	Stage 2	Stage 3
Dataset	ZINC250k	AGILE Virtual	AGILE Expt.
Dataset size	250,000	12,000	1,100
Epochs	500	300	109
Batch size	128	64	32
Learning rate	1×10^{-4}	5×10^{-5}	1×10^{-5}
LR scheduler	Cosine	Cosine	Cosine
Warmup epochs	10	5	3
Weight decay	1×10^{-5}	1×10^{-5}	1×10^{-5}
Gradient clipping	1.0	1.0	1.0
Diffusion timesteps	500	500	500
Noise schedule	Cosine	Cosine	Cosine
Region-aware noise	No	Yes	Yes
Scaffold conditioning	No	Yes	Yes
FiLM conditioning	No	No	Yes
CFG dropout	–	–	0.15

Compute requirements. All experiments were conducted on NVIDIA A10G GPUs via Modal cloud compute. Stage 1 training requires approximately 8 hours per seed; Stage 2 requires 4 hours; Stage 3 requires 2 hours. Total compute for the main experiments (5 seeds \times 3 stages) is approximately 70 GPU-hours.

H REGION-SPECIFIC MARGINAL DISTRIBUTIONS

Table 9 shows the atom type distributions computed separately for each functional region. These distributions are used as the noise target in region-aware diffusion.

Table 9: Atom type marginal distributions by region (%).

Atom	HEAD	CORE	TAIL	Global
C	52.3	71.8	94.2	78.4
N	28.1	12.4	0.3	8.2
O	15.2	14.6	4.8	10.1
S	2.1	0.4	0.2	0.8
Other	2.3	0.8	0.5	2.5

The HEAD region is significantly enriched in nitrogen (28.1% vs 8.2% globally), reflecting the ionizable amine functionality. TAIL regions are carbon-dominated (94.2%), consistent with lipophilic alkyl chains. Using these region-specific marginals rather than a single global marginal allows the denoising network to learn appropriate reconstructions for each functional region.

I SMARTS PATTERNS FOR SCAFFOLD DETECTION

I.1 UGI CORE DETECTION

The Ugi three-component reaction produces a characteristic bis-amide core structure. We detect this scaffold using the following SMARTS pattern, validated on all 1,100 experimental lipids (100% match rate):

```
[NX3][CX3](=[OX1])[CX4]([#6])[NX3]
```

This pattern matches five core atoms:

- [NX3]: Amide nitrogen (from isocyanide, connects to TAIL1)
- [CX3](=[OX1]): Carbonyl carbon and oxygen (amide bond formed during Ugi mechanism)
- [CX4]: Central sp³ carbon (the Ugi product center)
- [NX3]: Head nitrogen (connects toward ionizable amine HEAD)

I.2 REGION ASSIGNMENT ALGORITHM

Given a molecule with a detected Ugi scaffold, regions are assigned by breadth-first search (BFS) outward from the core:

1. **CORE**: The five matched atoms plus the carbonyl oxygen (5–6 atoms)
2. **HEAD**: BFS from the head nitrogen, excluding core atoms. Captures the ionizable amine and its substituents.
3. **TAIL1**: BFS from the amide nitrogen, excluding core and HEAD atoms. Captures the isocyanide-derived chain.
4. **TAIL2**: BFS from the carbon attached to the central C, excluding all previously assigned atoms. Captures the ester-linked chain.
5. **Fallback**: Any remaining unassigned atoms are assigned to their nearest assigned neighbor via BFS

Core bonds are defined as all bonds incident to at least one core atom (not just bonds between two core atoms). This broader definition ensures that the bonds connecting the core to each functional region, which are also formed by the Ugi mechanism, are protected during diffusion.

J PROPERTY-WEIGHTED SCAFFOLD SAMPLING

Ugi-based lipids span hundreds of distinct scaffold topologies (unique combinations of atom count, connectivity, and region assignments). At generation time, Algorithm 1 requires a scaffold as input. The choice of scaffold affects the space of molecules the model can generate, so scaffold selection is a meaningful design lever; our variance decomposition (Appendix N) shows that scaffold identity alone explains 65% of mTP variance (among scaffolds with 2+ molecules).

Scaffold library. We extract all unique scaffolds from the training data, recording for each scaffold the list of associated property values (mTP). Scaffolds are grouped into *composition classes* based on coarse molecular features (atom counts, ring count, double bond count). This grouping defines a hierarchy: each composition class contains one or more exact topologies.

Two-stage property-weighted sampling. Given a target property value y^* , scaffold selection proceeds in two stages:

1. **Composition class.** For each composition class c , compute the mean property \bar{y}_c of all training molecules in that class. Define weights $w_c = \exp(-|y^* - \bar{y}_c|/\tau)$, where τ is a temperature parameter. Sample a class from $\text{Cat}(w/\sum w)$.
2. **Exact topology.** Within the selected class, repeat the same property-weighted sampling over individual scaffolds to select an exact topology.

This biases generation toward scaffolds whose training molecules had properties close to the target, while maintaining stochasticity through the temperature τ .

Composition granularity. The grouping granularity trades off specificity against coverage (Table 10). Exact topology matching yields 527 classes but only 27.4% of molecules share a class with at least one other molecule (low within-class coverage). Coarser groupings aggregate more molecules per class, providing better property statistics for weighted sampling. We use `atoms_rings_db` (atom count + ring count + double bonds, 121 classes) as the default, balancing specificity with sufficient within-class coverage (85.7%) for reliable property estimates.

Table 10: Scaffold composition modes: granularity vs. within-class coverage.

Composition mode	Classes	Coverage (%)
Exact topology	527	27.4
Atoms + rings + double bonds	121	85.7
Atoms + ring type (5/6)	60	89.9
Atoms (binned) + rings + DB	34	90.5
Carbons + double bonds	68	90.9

K CHEMICAL SPACE VISUALIZATION

L EXTENDED RESULTS

L.1 FULL CANDIDATE LIST

Table 11 lists the top 20 generated candidates ranked by LANTERN-predicted mTP. All candidates are chemically valid with intact Ugi scaffolds. The top three candidates share dibutylamine head groups (NCCCN(CCCC)CCCC), suggesting the model learned this motif’s association with high transfection. Novel indicates the molecule does not appear in the AGILE training data.

L.2 REGION-AWARE ABLATION SAMPLES

Table 12 shows representative samples from the region-aware ablation experiment (model trained without regional noise distributions). Notable failures include allenes/cumulenes (C=C=C patterns, highlighted) and invalid molecules that failed RDKit sanitization.

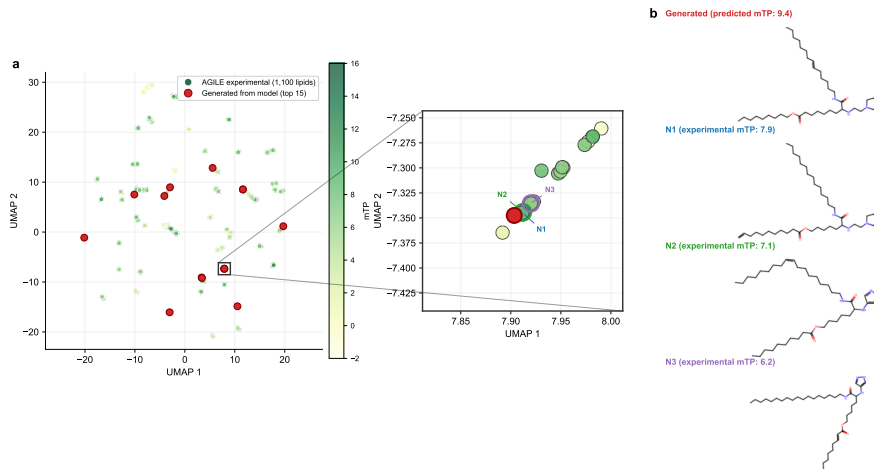


Figure 6: Nearest-neighbor interpolation analysis. **Left:** UMAP embedding of AGILE experimental data (colored by mTP) with a generated candidate (red star) and its three nearest training neighbors highlighted. **Right:** Molecular structures showing the generated candidate (predicted mTP = 9.8) alongside its nearest neighbors with experimental mTP values of 2.8, 8.4, and 5.2. Despite structural similarity in chemical space, the nearest neighbors span a 5.6-point mTP range, demonstrating that the model learns to interpolate toward high-activity regions rather than simply copying training molecules.

Table 11: Top 20 generated candidates by predicted mTP.

Rank	SMILES	Pred mTP	Novel
1	<chem>C=CCCCCCCCCCCCCCCCC (=O) OCCCCC (NCCCN (CCCC) CCCC) C (=O) NCCCCCCCCCCCCCCC</chem>	10.55	Yes
2	<chem>CCCCCCCCCCCCCCCCNC (=O) C (CCCCCOC (=O) CCCCCCCCCC) NCCCN (CCCC) CCCC</chem>	9.56	Yes
3	<chem>C#CCCCCCCCCCCCCCCCC (=O) OCCCCC (NCCCN (CCCC) CCCC) C (=O) NCCCCCCCCCCCCCCC</chem>	9.01	Yes
4	<chem>CCCCCCCCCCCCCCCCNC (=O) C (CCCCCOC (=O) CCCCCCCCCC) NCCN1CCCC1</chem>	8.58	Yes
5	<chem>CCCCCCCCCCCCCCCCCCCCNC (=O) C (CCCCCOC (=O) CCC (C) CCCCC) NCCN (C) C</chem>	8.18	Yes
6	<chem>CCCCCCCCCCCCCCCCCCCCNC (=O) C (CCCCCOC (=O) CCC (C) CCCCC) NCCCN (CCCC) CCCC</chem>	8.00	Yes
7	<chem>CCCCCCCCC=CCCCCCCCC (=O) OCCCCC (NCCCN (CC) CC) C (=O) NCCCCCCCCCCCCC</chem>	7.89	Yes
8	<chem>C=CCCCCCCCC (=O) OCCCCC (NC1=CCN=C1) C (=O) NCCCCCCCCCCCCC</chem>	7.87	Yes
9	<chem>CCCCCCCCC=CCCCCCCCCNC (=O) C (CCCCCOC (=O) CCCCCCCCCCCCCCCC) NC1=CCN=C1</chem>	7.83	Yes
10	<chem>CCCCCCCCC=CCCCCCCCC (=O) C (CCCCCOC (=O) CCCCCCCCCC) NN1CCCC1</chem>	7.81	Yes
11	<chem>CCCCCCCCC=CCCCCCCCC (=O) OCCCCC (NCCN1CCCC1) C (=O) NCCCCCCCCCCCCC</chem>	7.70	No
12	<chem>CCCCCCCCC=CCCCCCCCCNC (=O) C (CCCCCOC (=O) CCCCCCCCC=CCCCCCCC) NC1=CCN=C1</chem>	7.64	Yes
13	<chem>CCCCCCCCC=CCCCCCCCC (=O) OCCCCC (NCCN1CCCC1) C (=O) NCCCCCCCCCCCCCCC</chem>	7.57	No
14	<chem>CCCCCCCCC=CCCCCCCCCNC (=O) C (CCCCCOC (=O) CCCCCCCCCC) NC1CCN (C) CC1</chem>	7.44	Yes
15	<chem>CCCCCCCCC=CCCCCCCCCNC (=O) C (CCCCCOC (=O) CCCCCCCCCC) NCCCN1CCOCC1</chem>	7.41	Yes
16	<chem>CCCCCCCCC=CCCCCCCCCNC (=O) C (CCCCCOC (=O) CCC (C) CCCCC) NCCN (CC) CC</chem>	7.38	No
17	<chem>CCCCCCCCC=CC (=O) OCCCCC (NCCCN1CCOCC1) C (=O) NCCCCCCCCC=CCCCCCCCC</chem>	7.29	Yes
18	<chem>CCCCCCCCC=CCCCCCCCCNC (=O) C (CCCCCOC (=O) CCCCCCCCC=CCCCCCCC) NCCCN (CC) CC</chem>	7.24	Yes
19	<chem>C#CCCCCCCCCCCCCCCCC (=O) OCCCCC (NN1CCCC1) C (=O) NCCCCCCCCC=CCCCCCCCC</chem>	7.21	Yes
20	<chem>CCCCCCCCC=CCCCCCCCCNC (=O) C (CCCCCOC (=O) CCCCCCCCCCCCCCCC) NCCN (C) C</chem>	7.18	Yes

L.3 DETAILED GENERATION SWEEP

Tables 13–15 report generation metrics for each combination of training seed, guidance scale, and target mTP. Each row represents 10 generated samples. Validity indicates chemically valid molecules with intact Ugi scaffolds. Novelty is computed against both the AGILE virtual library (12k molecules) and experimental training set (1.1k molecules).

M MOLECULE STRUCTURES

[TBD: Grid of top 16 generated molecules with:

- 2D structure depictions
- Region coloring (HEAD=blue, CORE=gray, TAIL=orange)
- Predicted mTP annotation

Table 12: Region-aware ablation samples. Allene/cumulene motifs underlined>.

Idx	SMILES
9	<chem>CCCCCCCC=C=C=C=CCCCC (=O) OCCCCC (NCCCN (CCCC) CCCC) C (=O) N . . .</chem>
12	<chem>C=CCCCCCCC (=O) OCCCCC (NCCN (C) C) C (=O) NCCCCCCC=C=CCCCCCC</chem>
14	<chem>CCCCC=CCC=CCCC (=O) OCCCCC (NCCCN1CCOCC1) C (=O) NCCCCCCC=C=CC . . .</chem>
15	<chem>CCCCCCCC=C=CCCCCCCCNC (=O) C (CCCCOC (=O) CCCCCC) NN1CCCC1</chem>
19	<chem>CCCCCCCC=C=CCCCCCCCNC (=O) C (CCCCOC (=O) CCCCCCCCCCCCCC) N . . .</chem>
43	<chem>CCCCCCC=C=C=C=CCCCC (=O) OCCCCC (NC1CCN (C) CC1) C (=O) N . . .</chem>
6	INVALID (failed sanitization)
7	INVALID (failed sanitization)
35	INVALID (failed sanitization)
37	INVALID (failed sanitization)

Table 13: Seed 0: Generation metrics by guidance scale and target mTP.

CFG	Target	Valid (%)	Unique (%)	Novel (%)	Pred mTP	Pred mTP Std
1.0	4	100.0	100.0	70.0	4.84	3.43
1.0	6	100.0	100.0	70.0	5.53	2.16
1.0	8	80.0	100.0	62.5	6.36	3.07
1.0	10	100.0	100.0	50.0	5.60	2.00
1.0	12	100.0	100.0	60.0	4.81	2.51
1.0	14	100.0	100.0	50.0	6.10	2.66
2.0	4	90.0	100.0	66.7	5.76	2.43
2.0	6	100.0	90.0	77.8	6.59	3.32
2.0	8	90.0	100.0	44.4	5.61	2.33
2.0	10	100.0	100.0	50.0	5.78	1.66
2.0	12	80.0	100.0	62.5	4.50	2.13
2.0	14	100.0	100.0	60.0	5.22	1.63
3.0	4	100.0	100.0	70.0	5.42	1.82
3.0	6	100.0	100.0	70.0	6.46	2.41
3.0	8	100.0	100.0	90.0	5.09	3.97
3.0	10	90.0	100.0	55.6	6.63	3.73
3.0	12	90.0	100.0	66.7	5.16	2.14
3.0	14	100.0	100.0	40.0	4.90	1.95

- Key molecular properties (MW, LogP, TPSA)

]

N VARIANCE DECOMPOSITION

To quantify the relative contributions of scaffold selection versus diffusion-based chemical optimization, we performed ANOVA-style variance decomposition on predicted mTP values.

Among the 527 unique scaffolds in the training data, 272 contain two or more molecules (845 molecules total). Variance decomposition on these scaffolds reveals that scaffold identity explains **65%** of mTP variance, while within-scaffold chemical variation explains **35%**. The median within-scaffold mTP range is 6.2 points, which is substantial given the training data spans approximately 15 points total.

This decomposition motivates our two-stage generation approach: property-weighted scaffold selection leverages the dominant source of variance (65%), while diffusion-based generation optimizes the remaining chemical degrees of freedom (35%). See Section B for details on how these components interact in the full pipeline.

Table 14: Seed 1: Generation metrics by guidance scale and target mTP.

CFG	Target	Valid (%)	Unique (%)	Novel (%)	Pred mTP	Pred mTP Std
1.0	4	100.0	100.0	90.0	5.65	1.53
1.0	6	90.0	100.0	33.3	4.27	2.44
1.0	8	100.0	100.0	70.0	6.12	3.01
1.0	10	90.0	100.0	77.8	5.20	1.60
1.0	12	100.0	100.0	70.0	4.61	2.85
1.0	14	100.0	100.0	80.0	4.37	1.85
2.0	4	100.0	100.0	50.0	6.71	2.23
2.0	6	90.0	100.0	77.8	5.72	2.60
2.0	8	100.0	100.0	50.0	3.77	2.54
2.0	10	100.0	100.0	50.0	6.42	2.58
2.0	12	80.0	100.0	75.0	6.58	3.07
2.0	14	90.0	100.0	77.8	4.92	2.41
3.0	4	100.0	100.0	70.0	4.50	2.44
3.0	6	90.0	100.0	66.7	5.35	1.94
3.0	8	90.0	100.0	66.7	4.39	2.81
3.0	10	80.0	100.0	75.0	5.52	2.03
3.0	12	40.0	100.0	100.0	4.39	2.23
3.0	14	60.0	100.0	66.7	4.49	4.07

Table 15: Seed 2: Generation metrics by guidance scale and target mTP.

CFG	Target	Valid (%)	Unique (%)	Novel (%)	Pred mTP	Pred mTP Std
1.0	4	100.0	100.0	90.0	5.06	2.21
1.0	6	100.0	90.0	77.8	3.78	2.53
1.0	8	100.0	100.0	80.0	5.27	2.91
1.0	10	100.0	100.0	70.0	4.43	1.86
1.0	12	100.0	100.0	50.0	5.32	1.96
1.0	14	90.0	100.0	55.6	3.87	2.38
2.0	4	100.0	100.0	80.0	4.97	1.88
2.0	6	100.0	100.0	40.0	5.67	1.03
2.0	8	100.0	100.0	80.0	5.76	3.23
2.0	10	100.0	100.0	60.0	5.63	2.00
2.0	12	100.0	100.0	50.0	5.59	2.66
2.0	14	90.0	100.0	55.6	4.84	1.91
3.0	4	100.0	100.0	70.0	5.42	1.94
3.0	6	90.0	100.0	77.8	5.09	1.55
3.0	8	100.0	100.0	70.0	5.76	2.31
3.0	10	100.0	100.0	70.0	5.74	2.18
3.0	12	90.0	100.0	77.8	4.00	1.90
3.0	14	100.0	100.0	70.0	5.97	1.76

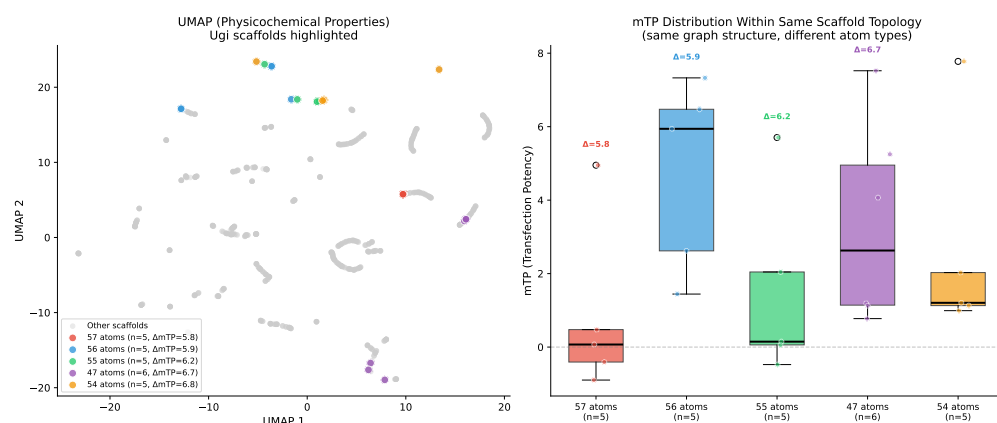


Figure 7: Within-scaffold mTP variance in Ugi lipids. **Left:** UMAP embedding of 1,100 AGILE lipids based on physicochemical properties (MW, LogP, TPSA, HBD, HBA), with five representative scaffolds highlighted. Scaffolds are defined by molecular topology (graph structure); molecules sharing a scaffold have identical atom count and bond connectivity but may differ in atom/bond types. **Right:** mTP distributions within each highlighted scaffold. Despite sharing identical graph topology, molecules within the same scaffold span 5–7 mTP points, demonstrating that atom and bond type variation; the degrees of freedom controlled by diffusion; drives meaningful activity differences. Scaffolds shown are median-variance (not outliers): 46 scaffolds have ≥ 5 molecules with ≥ 2 -point mTP range; these five are centered at the median.