SOLVMORPH: GENERATING SOLVENT-DEPENDENT CYCLIC PEPTIDE ENSEMBLES FROM LIMITED SOURCE CONFORMATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Cyclic peptides can adopt distinct conformational ensembles across solvents, but repeatedly sampling each molecule under every environment is computationally expensive. Inferring a target-solvent ensemble from conformations observed in another solvent requires recovering structures that may be absent from the source ensemble. Fixed-bank reweighting can improve agreement with target observables, but cannot recover missing conformational support. We introduce **SolvMorph**, a source- and target-conditioned Schrödinger bridge that generates target-specific conformations in periodic torsion space, followed by observable-guided probability refinement. On CycPeptMPDB-4D, source conformations improve unobserved target-solvent ensemble prediction, while target-aware generation expands conformational support beyond fixed candidate banks. On CONF192, a generator-naive confirmation panel, SolvMorph improves target-emergent mode recall by 0.26 and 0.35 over source copying for water-to-hexane and hexane-to-water inference, respectively. These gains persist after matching proposal counts and under swapped-target controls, demonstrating recovery of solvent-specific conformational modes from limited source observations. Code is available at `https://anonymous.4open.science/r/SolvMorph`.

1 INTRODUCTION

Cyclic peptides can reorganize their backbones and intramolecular interactions across solvent environments, changing the accessibility of polar groups to the surrounding solvent (Danelius et al., 2020; Ramelot et al., 2023; Hagemeyer et al., 2026). This behavior, known as molecular chameleonicity, can favor internally shielded conformations in nonpolar environments and influence passive membrane permeability (Poongavanam et al., 2024; Linker et al., 2022; 2023). Understanding these changes requires characterizing conformational ensembles under different solvent conditions (Figure 1). While molecular dynamics simulations could resolve solvent-dependent structure, repeated sampling becomes costly across molecules and environments (Damjanovic et al., 2021). We posit that cross-solvent inference could reduce this burden by using a limited ensemble observed in one solvent to predict the ensemble of the same molecule in another.

Recent diffusion and flow-based methods have accelerated conformational sampling of macro-

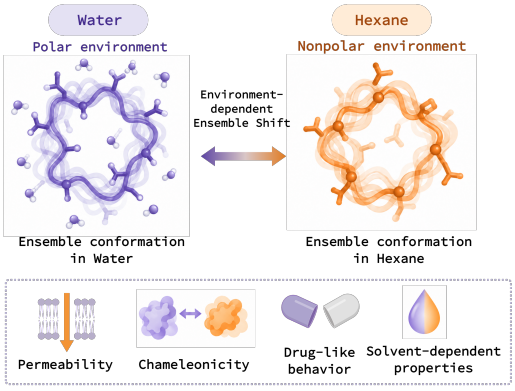


Figure 1: **Solvent-dependent conformational ensembles of cyclic peptides.** Water and hexane can favor different backbone geometries and patterns of polar-group accessibility in the same cyclic peptide. These changes contribute to molecular chameleonicity and are relevant to membrane permeability.

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

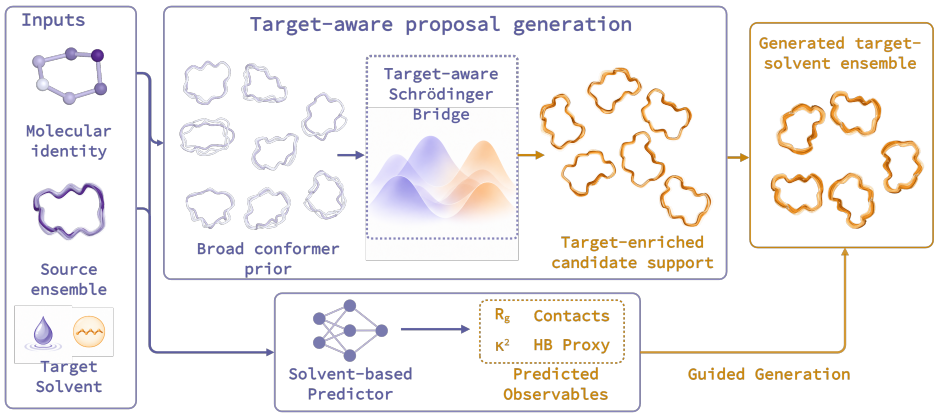


Figure 2: **SolvMorph generates target-specific conformations for cross-solvent ensemble inference.** Conditioned on molecular identity, the observed source ensemble, and the requested solvent, we transport a broad conformer prior through periodic torsion space using a Schrödinger bridge. We reconstruct cyclic conformations by combining generated torsions with non-generated internal coordinates from an observed source conformer, then refine candidate weights using predicted target observables.

cycles and flexible peptides (Grambow et al.,
2023; Abdin & Kim, 2024; Bayaraa et al., 2026; Schaufelberger et al., 2026; Secor et al., 2026). Schrödinger bridges further support conformer refinement (Lee et al., 2024), and environment-conditioned generation predicts structures across physical settings (Yan et al., 2026). For cyclic peptides, multisolvent ensemble studies have also connected structural changes to permeability (Liu et al., 2026; Yu et al., 2026; Wen et al., 2026). Building on these advances, we hypothesize that recovering target-specific structures using conformations already observed for the same molecule is possible. A natural starting point for this goal is to predict target observables and reweight available conformers to match them (Bottaro et al., 2020; Gama Lima Costa & Fushman, 2022). However, matching ensemble averages does not uniquely determine a conformational distribution, and reweighting cannot recover structures absent from the candidate bank. Recovering modes rarely sampled in the source environment therefore requires generating additional conformations conditioned on the requested solvent.

To generate this missing support, we introduce **SolvMorph**, a source- and target-conditioned Schrödinger bridge followed by observable-guided probability refinement (Figure 2). Using simulation-free bridge training (Tong et al., 2024), we learn transport from a broad conformer prior to target-solvent conformations in periodic torsion space, conditioned on molecular identity, the observed source ensemble, and the requested solvent transition. During cyclic reconstruction, we combine generated torsions with the remaining internal coordinates from one observed source conformer. Source observations thus guide both generation and reconstruction. We then refine valid candidate weights using predicted target observables, evaluating solvent specificity before refinement to isolate the contribution of generation.

Across 3,456 cyclic peptides from CycPeptMPDB-4D (Liu et al., 2026), we find that source conformations improve target-observable prediction beyond molecular identity alone and persistence of the source state. On 192 molecules excluded from generator development, we recover target-emergent modes with paired recall gains over source copying of 0.26 for water-to-hexane inference and 0.35 for hexane-to-water inference. These gains persist after matching proposal counts, and swapped-solvent controls confirmed target specificity before refinement. SolvMorph therefore recovers solvent-dependent conformational modes from limited source observations, extending structural coverage to regions nearly absent from the source ensemble.

2 Related Work

Cyclic-peptide conformational ensemble generation. Large-scale conformational resources, including CREMP and MacroConf, support the training and evaluation of cyclic-peptide structural models (Grambow et al., 2024; Crusius et al., 2023). MD-supervised approaches such as StrEAMM predict cyclic-peptide structural populations from sequence, with extensions to neural ensemble prediction and thioether-linked scaffolds (Miao et al., 2021; Hui et al., 2023; Ho et al., 2026). Explicit ensemble generation has advanced through internal-coordinate diffusion in RINGER, sequence-conditioned all-atom generation in PepFlow, and recent diffusion or flow-based models for heterochiral and constrained macrocycles (Grambow et al., 2023; Abdin & Kim, 2024; Bayaraa et al., 2026; Powers et al., 2025; Secor et al., 2026; Schaufelberger et al., 2026). These approaches enable efficient conformational sampling, motivating the use of observed conformations to infer environmental changes.

Solvent-dependent conformational modeling. Solvent-induced changes in backbone organization, hydrogen bonding, and polar-group exposure underlie cyclic-peptide chameleonicity and its connection to permeability (Danelius et al., 2020; Tang et al., 2023; Linker et al., 2022; 2023). CycPeptMPDB-4D provides water and hexane trajectories for large-scale solvent-dependent ensemble analysis (Liu et al., 2026). EnsembleCycPerm and Delta PSA further use solvent-specific ensembles to predict permeability-related properties and quantify conformational adaptation (Wen et al., 2026; Yu et al., 2026). Learned solvent-aware sampling methods, including implicit-solvent potentials and environment-conditioned generation, aim to reduce the cost of obtaining such ensembles (Katzberger & Riniker, 2024; Katzberger et al., 2025; Yan et al., 2026). Our setting differs by using a finite ensemble of the same molecule observed in one solvent to infer the ensemble in another.

Generation and transport across physical conditions. Boltzmann generators and temperature-conditioned flows connect learned sampling with physical distributions across thermodynamic states (Noé et al., 2019; Dibak et al., 2022). Thermodynamic interpolation and latent thermodynamic flows further extend generation across changing environments (Moqvist et al., 2025; Qiu et al., 2026). Schrödinger-bridge methods learn transport between distributions through score and flow matching, with extensions to geometric spaces and molecular conformer refinement (Shi et al., 2023; Tong et al., 2024; Thornton et al., 2022; Chen & Lipman, 2024; Lee et al., 2024). Bayesian and maximum-entropy reweighting can further adjust probabilities of existing conformations to match structural observables (Bottaro et al., 2020; Gama Lima Costa & Fushman, 2022). SolvMorph combines these ideas by conditioning transport on observed source conformations and the requested solvent, then refining generated ensembles using predicted target observables.

3 Methods

3.1 Cross-Solvent Ensemble Inference

Let G denote the molecular graph of a cyclic peptide, including atom identities, connectivity, and stereochemistry, and let s and t denote the source and target solvents. The corresponding conformational distributions are

$$p_s(C \mid G) \qquad \text{and} \qquad p_t(C \mid G), \tag{1}$$

where C denotes a molecular conformation.

At inference, we observe molecular identity, the source and requested target solvents, and a finite ensemble of K source conformations,

$$E_s = \{C_i^{(s)}\}_{i=1}^{K}. \tag{2}$$

The empirical source distribution assigns equal probability to each observed conformation,

$$\widehat{p}_s = \frac{1}{K} \sum_{i=1}^{K} \delta_{C_i^{(s)}}, \tag{3}$$

where δ_C denotes a point mass at conformation C. Our goal is to construct a target ensemble $\widehat{E}_t$ that approximates the target-solvent distribution using these source observations.

For a bank $\mathcal{B}_t = \{C_j\}_{j=1}^N$ of N valid candidate conformations, the weighted target approximation is

$$\widehat{p}_t = \sum_{j=1}^{N} w_j \, \delta_{C_j}, \qquad w_j \geq 0, \qquad \sum_{j=1}^{N} w_j = 1. \tag{4}$$

Each weight w_j specifies the probability assigned to candidate C_j. Reconstruction therefore depends on both the conformational regions covered by $\mathcal{B}_t$ and the probability assigned within those regions.

We construct this bank through target-conditioned generation from a broad conformer prior. Source observations condition generation through a permutation-invariant ensemble representation and determine the non-generated internal coordinates used during cyclic reconstruction. A separate predictor estimates target structural observables for subsequent probability refinement. Dataset partitions, source and reference blocks, and cohort construction are described in Appendix A.

3.2 Cross-Solvent Observable Prediction

To guide probability refinement, we first estimate a compact set of target structural averages from the observed source ensemble. Let $\phi(C) \in \mathbb{R}^d$ denote the structural measurements of conformation C. We use $d = 4$ measurements: radius of gyration, relative shape anisotropy, intramolecular contact density, and a geometric hydrogen-bond proxy. Their ensemble averages under a conformational distribution p are

$$\mathbf{m}(p) = \mathbb{E}_{C\sim p}[\phi(C)]. \tag{5}$$

Thus, $\mathbf{m}(p)$ summarizes the expected structural measurements across the ensemble.

Using a predictor f_θ with fitted parameters θ, we estimate the target averages from molecular identity, the observed source ensemble, and the solvent transition,

$$\widehat{\mathbf{m}}_t = f_\theta(G, E_s, s, t). \tag{6}$$

We parameterize this prediction through the change from the observed source averages,

$$\widehat{\mathbf{m}}_t = \mathbf{m}(\widehat{p}_s) + \widehat{\Delta\mathbf{m}}_{s\to t}, \qquad \mathbf{m}(\widehat{p}_s) = \frac{1}{K}\sum_{i=1}^{K} \phi(C_i^{(s)}). \tag{7}$$

Here, $\widehat{\Delta\mathbf{m}}_{s\to t}$ is the predicted solvent-induced change. Separate gradient-boosted regressors estimate each component for each solvent transition using molecular descriptors and permutation-invariant source-ensemble statistics.

However, agreement between structural averages does not uniquely determine the conformational distribution. For two distributions p and q,

$$\mathbf{m}(p) = \mathbf{m}(q) \not\Rightarrow p = q. \tag{8}$$

We therefore use the predicted averages to refine candidate probabilities after generating target-specific conformations. Predictor features, fitting procedures, and hyperparameters are detailed in Appendix C.

3.3 Broad Prior and Candidate Support

To generate conformations beyond those observed in the source ensemble, we begin from a broad prior that depends on molecular identity,

$$C_0 \sim q_0(C \mid G). \tag{9}$$

We construct this prior using ETKDGv3 followed by MMFF94 relaxation (Wang et al., 2020). The resulting conformer bank is independent of the requested solvent.

A conditional generator with parameters ψ transforms prior conformations using the observed source ensemble and solvent transition,

$$C_t \sim q_\psi(C \mid C_0, G, E_s, s, t). \tag{10}$$

For M attempted proposals, we construct

$$\widetilde{\mathcal{B}}_t = \{C_j^{(t)}\}_{j=1}^{M}, \qquad C_j^{(t)} \sim q_\psi(\cdot \mid C_{0,j}, G, E_s, s, t), \tag{11}$$

with each $C_{0,j}$ sampled from $q_0(\cdot \mid G)$. We retain the N proposals that pass chemical quality control to obtain the valid bank $\mathcal{B}_t \subseteq \widetilde{\mathcal{B}}_t$. Raw valid yield is N/M. For matched-support comparisons, we also construct a baseline bank that combines observed source conformers with generic prior conformations at the same total candidate budget.

To quantify coverage independently of candidate weights, let $d_{\mathrm{str}}(C, C')$ denote a structural distance and let $\tau > 0$ denote a coverage threshold. The uncovered target mass is

$$u_\tau(\mathcal{B}_t, p_t) = \Pr_{C \sim p_t(\cdot|G)} \left[\min_{C_j \in \mathcal{B}_t} d_{\mathrm{str}}(C, C_j) > \tau \right]. \tag{12}$$

This quantity is the probability that a target conformation lies farther than τ from every candidate in the bank. Because u_τ depends only on candidate locations, changing their probability weights leaves it unchanged. Reducing uncovered mass therefore requires additional conformations near the missing target regions. We estimate this quantity using held-out target references after generation.

3.4 Target-Aware Schrödinger-Bridge Generation

We restrict generation to flexible periodic torsions and retain molecular identity, connectivity, and stereochemical labels as fixed conditioning information. Let $x(C) \in \mathbb{T}^m$ collect the m generated torsions of conformation C, where $\mathbb{T}^m$ denotes the space of m angles defined modulo 2π. During reconstruction, we combine these torsions with the remaining internal coordinates copied from one observed source conformer. This construction uses the source ensemble both for learned conditioning and for the geometry required to reconstruct complete cyclic conformations.

During training, we sample conformations independently from the broad prior and the target-solvent ensemble. For a fixed molecule and solvent transition, let $\Pi(q_0, p_t)$ denote the set of joint distributions with these prior and target marginals. We couple the samples through an entropy-regularized transport objective,

$$\pi^* = \arg \min_{\pi \in \Pi(q_0, p_t)} \left[\mathbb{E}_{(C_0, C_1) \sim \pi}[c(C_0, C_1)] + \varepsilon D_{\mathrm{KL}}(\pi \| q_0 \otimes p_t) \right]. \tag{13}$$

Here, $c(C_0, C_1)$ combines periodic torsional and intramolecular-distance discrepancies, $\varepsilon > 0$ controls regularization, and D_{KL} denotes Kullback–Leibler divergence. The product distribution $q_0 \otimes p_t$ corresponds to independent prior and target sampling. Minimizing this objective favors structurally similar endpoint pairs while penalizing departures from independent coupling. The pairing defines training paths between conformations without requiring correspondence between trajectory frames.

For each coupled pair $(C_0, C_1) \sim \pi^*$, we write $x_0 = x(C_0)$ and $x_1 = x(C_1)$. At bridge time $r \in (0, 1)$, we sample an intermediate torsional state x_r from the wrapped Brownian-bridge path between these endpoints. Let $u_r(x_r \mid x_0, x_1)$ denote the analytical conditional velocity of this path. Following simulation-free Schrödinger-bridge flow matching (Tong et al., 2024), we fit a velocity field v_ψ by minimizing its mean squared difference from the conditional velocity,

$$\mathcal{L}_{\mathrm{SB}} = \mathbb{E}_{(C_0, C_1), r, x_r} \left[\| v_\psi(x_r, r \mid G, E_s, s, t) - u_r(x_r \mid x_0, x_1) \|_2^2 \right]. \tag{14}$$

The expectation averages over coupled endpoint samples, bridge times, and intermediate states. Conditioning includes molecular identity, a permutation-invariant source-ensemble encoding, and the requested solvent transition.

To penalize steric clashes during generation, we evaluate a differentiable feasibility loss $\mathcal{L}_{\mathrm{clash}}$ after combining generated torsions with the source-derived reconstruction geometry. The full training objective is

$$\mathcal{L} = \mathcal{L}_{\mathrm{SB}} + \lambda_{\mathrm{clash}} \mathcal{L}_{\mathrm{clash}}, \tag{15}$$

where $\lambda_{\mathrm{clash}} \geq 0$ controls the contribution of clash penalties. We use the same source-derived geometry during training and inference.

At inference, we integrate the learned probability-flow ODE on the torsional space,

$$\frac{\mathrm{d}x_r}{\mathrm{d}r} = v_\psi(x_r, r \mid G, E_s, s, t), \qquad x_0 = x(C_0), \qquad C_0 \sim q_0(\cdot \mid G). \tag{16}$$

Integration from $r = 0$ to $r = 1$ transforms prior torsions into target-conditioned torsions. We then reconstruct complete conformations using the source-derived internal coordinates and retain candidates that pass quality control. The wrapped path, architecture, reconstruction procedure, clash loss, and sampling settings are detailed in Appendices C.4 and H.

3.5 OBSERVABLE-GUIDED ENSEMBLE REFINEMENT

Given a nonempty valid bank $\mathcal{B}_t = \{C_j\}_{j=1}^N$, we refine candidate probabilities through regularized moment matching (Bottaro et al., 2020). The weight vector belongs to the probability simplex,

$$\mathbf{w} \in \Delta_N, \qquad \Delta_N = \left\{ \mathbf{w} \in \mathbb{R}^N_{\geq 0} : \sum_{j=1}^N w_j = 1 \right\}. \tag{17}$$

The corresponding ensemble averages are

$$\mathbf{m}(\mathcal{B}_t, \mathbf{w}) = \sum_{j=1}^N w_j \phi(C_j). \tag{18}$$

Each component is therefore a weighted average of the structural measurements across valid candidates.

To compare observables with different scales, let $\sigma_o > 0$ denote the standard deviation of target observable o estimated from the training fold. We regularize the weights toward the uniform distribution $\mathbf{u} = (1/N, \ldots, 1/N)$ and limit concentration through the effective sample size,

$$\mathrm{ESS}(\mathbf{w}) = \frac{1}{\sum_{j=1}^N w_j^2}. \tag{19}$$

The refined weights minimize standardized observable error and deviation from uniform weighting, subject to an effective sample size of at least $N/4$,

$$\mathbf{w}^* = \arg \min_{\substack{\mathbf{w} \in \Delta_N \\ \mathrm{ESS}(\mathbf{w}) \geq N/4}} \left[\frac{1}{d} \sum_{o=1}^d \left(\frac{\sum_{j=1}^N w_j \phi_o(C_j) - \widehat{m}_{t,o}}{\sigma_o} \right)^2 + \lambda D_{\mathrm{KL}}(\mathbf{w} \| \mathbf{u}) \right]. \tag{20}$$

Here, $\phi_o(C_j)$ and $\widehat{m}_{t,o}$ are the candidate measurement and predicted target average for observable o, respectively, and we use $\lambda = 0.1$. The predicted target vector is computed once from the original source ensemble and remains fixed during optimization. We then sample $\widehat{E}_t$ from the weighted distribution in equation 4 using $\mathbf{w}^*$.

To assess generation independently of probability refinement, we measure target coverage, proposal precision, chemical validity, novelty, and correct-versus-swapped solvent specificity before optimizing the weights. We then apply the same frozen refinement procedure to each valid proposal bank and evaluate the resulting ensembles. Optimization settings, scaling, and oracle controls are specified in Appendices C and F.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETTINGS

We evaluated cross-solvent inference in both directions between water (W) and hexane (H) using stereochemistry-corrected CycPeptMPDB-4D (Liu et al., 2026). After quality control, we retained 5,057 molecules and assessed observable prediction on 3,456 molecules from 162 structural families under grouped out-of-fold evaluation. For generator development, we used 96 molecules from distinct structural groups (PILOT96), then evaluated the frozen model on 192 molecules from 35 groups excluded from development (CONF192).

Our evaluation separates prediction, conformational support, and probability assignment. We measured target-observable prediction error, structural agreement, proposal coverage, chemical validity, and solvent specificity with grouped bootstrap uncertainty. Appendices A, B, and C describe the cohorts, baselines, metrics, and evaluation procedures.

4.2 SOURCE CONFORMATIONS PREDICT UNOBSERVED TARGET-SOLVENT ENSEMBLE STRUCTURE

We first tested whether source conformations contain transferable information about the unseen solvent beyond molecular identity and source-state persistence. Incorporating source geometry reduced

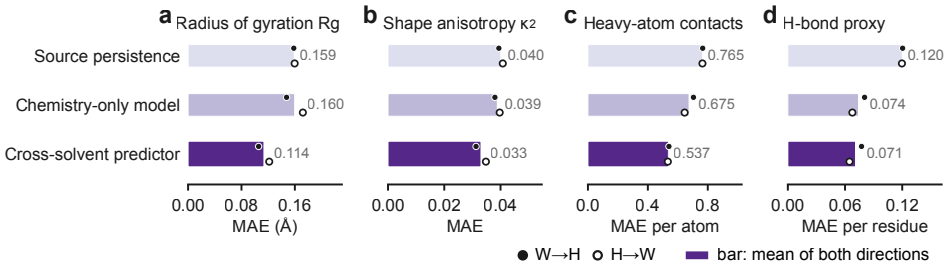


Figure 3: **Cross-solvent prediction from limited source conformations.** MAE comparison across four target observables. Bars show the mean over both solvent directions and dots show each direction separately.

MAE for all four observables, including approximately 29% and 30% reductions for R_g and heavy-atom contacts relative to source persistence (Figure 3). These results indicate that limited source observations preserve structural information about solvent-dependent conformational changes. Additional encoder and source-size analyses are reported in Appendix E.

4.3 Conformational Coverage Limits Fixed-Support Reconstruction

Observable-guided reweighting improves ensemble averages but cannot recover conformations absent from the candidate bank. This limitation arises because reweighting only changes probability mass over existing candidates. Even accurate target-observable prediction cannot recover states missing from the candidate support, motivating target-aware generation as a mechanism for support expansion rather than only probability refinement (Appendix G and Appendix F).

4.4 Target-Aware Generation Recovers Missing Conformational Support

We evaluated target-aware generation before probability refinement to isolate the contribution of conformational support generation from subsequent probability assignment (Figure 4). These experiments test whether target conditioning can recover conformational regions unavailable from the observed source ensemble. On PILOT96, mixed support already improved over the generic prior, confirming that observed source conformers provide useful candidate coverage. Target-aware generation further expanded this support, reducing the uncovered target fraction while preserving chemical validity.

We next evaluated the frozen generator on CONF192, where molecules were excluded from generator development. As shown in Figure 4, SolvMorph maintains improved target coverage on this generator-naive confirmation panel, indicating that target-aware support expansion generalizes beyond the molecules used for development.

We then quantified recovery of target-emergent modes, defined as target-reference modes with at least 10% occupancy and source-reference occupancy below 2%. These modes directly test whether generation reaches conformational regions unavailable to fixed-support reconstruction. Figure 5 shows strong recall gains before refinement. On PILOT96, SolvMorph improved target-emergent mode recall over the source-ensemble baseline by 0.36 in both solvent directions. On CONF192, the paired gains remained 0.26 for W→H and 0.35 for H→W. The gains persisted after matching the 20-proposal source budget and decreased under wrong-target conditioning, supporting solvent-specific recovery rather than generic conformational expansion.

4.5 Case Studies of Target-Emergent Mode Recovery

Figure 6 shows two representative target-emergent modes, one from each solvent direction. In both examples, the observed source ensemble and fixed-support mixed-support baseline miss the target-emergent mode, whereas SolvMorph generates conformations assigned to the same target region. The held-out target ensemble confirms that these recovered conformations correspond to solvent-

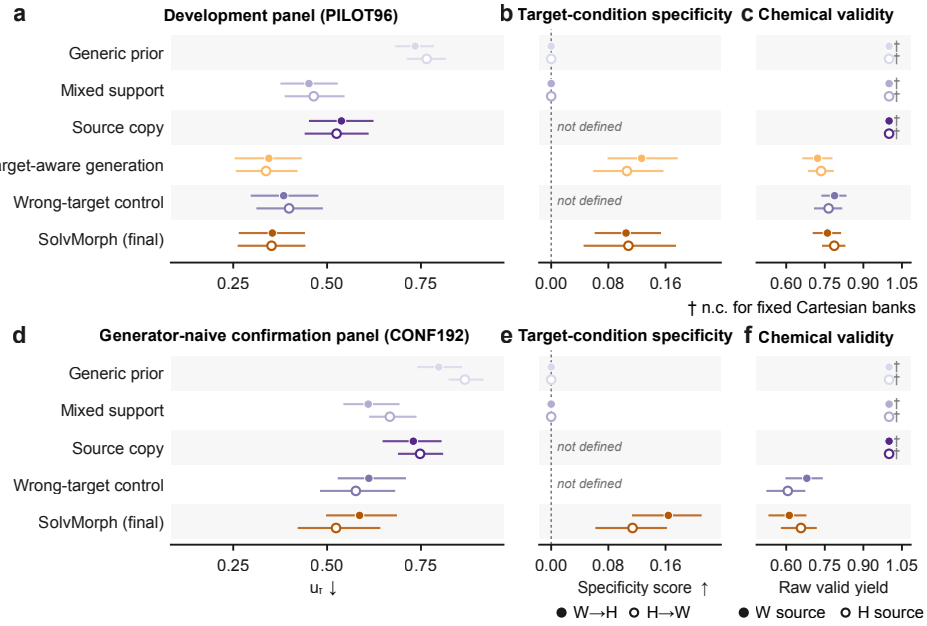


Figure 4: **Proposal-level confirmation of target-aware generation on two panels.** Lower u_τ indicates better target coverage, while specificity compares correct and swapped target conditions.

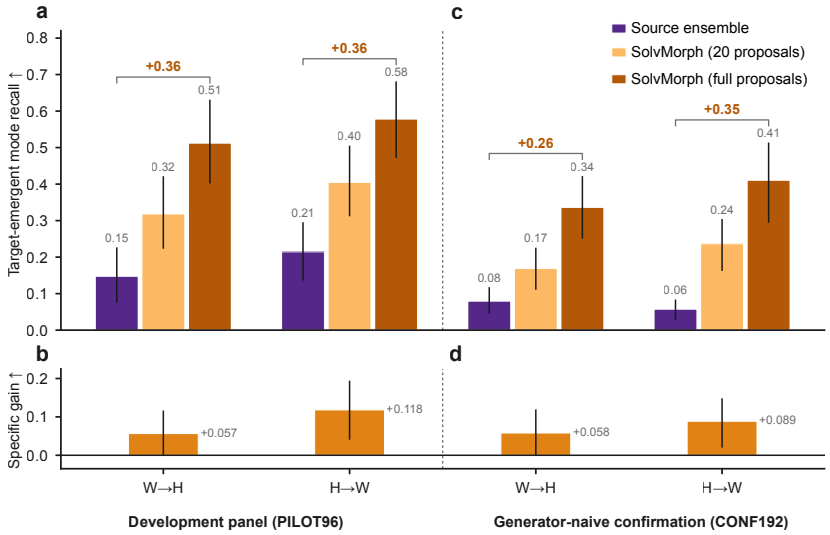


Figure 5: **Recovery of target-emergent modes.** SolvMorph improves recall of target-reference modes that are absent or rarely sampled in the source ensemble. Error bars denote 95% grouped-bootstrap confidence intervals.

specific structures rather than generic novelty. These examples qualitatively illustrate the same support-expansion phenomenon quantified by the held-out mode-recovery analysis.

After observable-guided refinement, the generated ensembles further improve ensemble agreement beyond support generation alone. These results show that observable-guided refinement complements target-aware generation by improving probability assignment after candidate support has been expanded, rather than replacing the need for support generation. Additional held-out structural response analyses are reported in Appendix E.3.

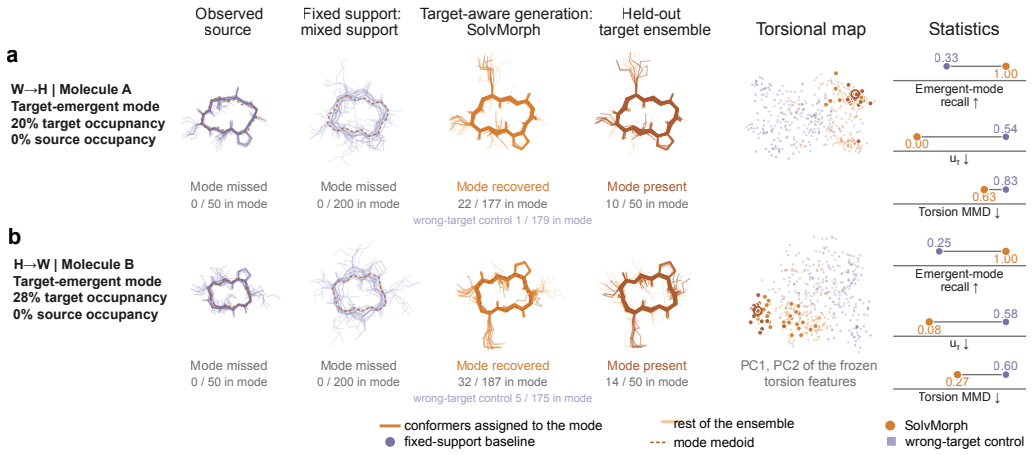


Figure 6: **Case studies of target-emergent mode recovery.** Two held-out examples illustrate structural recovery of target-emergent modes. SolvMorph recovers conformational regions missed by observed source ensembles and fixed-support baselines, consistent with the mode-recall analysis.

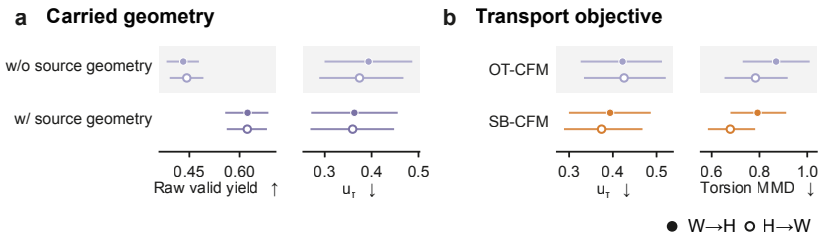


Figure 7: **Generation ablations on source geometry and transport objectives.** Dots denote water→hexane and hexane→water.

4.6 Ablation Studies

Figure 7 isolates two design choices. Source-carried geometry substantially improves chemical validity, while Schrödinger-bridge training improves target support and structural agreement over optimal-transport matching under comparable conditions. Additional ablations on support construction, conditioning, and chemical feasibility are reported in Appendix H.

5 Conclusion

We introduce SolvMorph, a source- and target-conditioned framework for reconstructing solvent-dependent cyclic-peptide ensembles from limited source observations. By combining target-aware Schrödinger-bridge generation with observable-guided refinement, SolvMorph recovers target-emergent conformational modes absent from fixed-support reconstruction on held-out molecules. Accurate recovery of target populations remains unresolved, and broader validation requires more diverse solvents and longer trajectory references. Future work will focus on improving proposal validity and extending cross-environment ensemble generation to larger molecular systems.

Ethics Statement

This work uses publicly available molecular simulation datasets and does not involve human subjects, personal data, animal experiments, or newly collected biological samples. The study focuses

on computational modeling of molecular conformational ensembles. No additional ethical concerns related to data collection, privacy, or experimental intervention arise from the reported analyses.

AI USAGE STATEMENT

AI-assisted tools were used only to aid manuscript organization, language editing, and writing polish. They were not used to generate experimental data or execute scientific analyses. All scientific claims, results, interpretations, and final manuscript text were reviewed and approved by the authors.

REPRODUCIBILITY STATEMENT

All reported experiments use fixed dataset partitions, frozen evaluation definitions, and the source/target trajectory blocks described in Appendix A. Prediction results use grouped out-of-fold evaluation, while generation results use fixed development and generator-naive confirmation panels. The Appendix provides dataset construction details, metric definitions, implementation settings, ablation configurations, and supplementary analyses required to reproduce the reported comparisons. Reported uncertainty intervals are computed with grouped bootstrap resampling, and all comparisons use the same structural metrics, quality-control criteria, and proposal budgets described in the appendix.

REFERENCES

Osama Abdin and Philip M Kim. Direct conformational sampling from peptide energy landscapes through hypernetwork-conditioned diffusion. *Nature Machine Intelligence*, 6(7):775–786, 2024.

Nomindari Bayaraa, Maxim Secor, Marc L Descoteaux, and Yu-Shan Lin. Fast generation of simulation-quality structural ensembles of mixed-chirality cyclic peptides via diffusion models. *Journal of Chemical Theory and Computation*, 22(6):3103, 2026.

Sandro Bottaro, Tone Bengtsen, and Kresten Lindorff-Larsen. Integrating Molecular Simulation and Experimental Data: A Bayesian/Maximum Entropy Reweighting Approach. In *Structural Bioinformatics*, Methods in Molecular Biology, pp. 219–240. Springer, 2020. doi: 10.1007/978-1-0716-0270-6_15. URL https://doi.org/10.1007/978-1-0716-0270-6_15.

Ricky TQ Chen and Yaron Lipman. Flow matching on general geometries. In *International Conference on Learning Representations*, volume 2024, pp. 47922–47945, 2024.

Daniel Crusius, Jason R Schnell, Flaviu Cipcigan, and Philip C Biggin. Macroconf–dataset & workflows to assess cyclic peptide solution structures. *Digital Discovery*, 2(4):1163–1177, 2023.

Jovan Damjanovic, Jiayuan Miao, He Huang, and Yu-Shan Lin. Elucidating Solution Structures of Cyclic Peptides Using Molecular Dynamics Simulations. *Chemical Reviews*, 121(4):2292–2324, 2021. doi: 10.1021/acs.chemrev.0c01087. URL https://doi.org/10.1021/acs.chemrev.0c01087.

Emma Danelius, Vasanthanathan Poongavanam, Stefan Peintner, Lianne HE Wieske, Máté Erdélyi, and Jan Kihlberg. Solution conformations explain the chameleonic behaviour of macrocyclic drugs. *Chemistry–A European Journal*, 26(23):5231–5244, 2020.

Manuel Dibak, Leon Klein, Andreas Krämer, and Frank Noé. Temperature steerable flows and boltzmann generators. *Physical Review Research*, 4(4):L042005, 2022.

Raquel Gama Lima Costa and David Fushman. Reweighting methods for elucidation of conformation ensembles of proteins. *Current Opinion in Structural Biology*, 77:102470, 2022. doi: 10.1016/j.sbi.2022.102470. URL https://doi.org/10.1016/j.sbi.2022.102470.

Colin A Grambow, Hayley Weir, Nathaniel L Diamant, Gabriele Scalia, Tommaso Biancalani, and Kangway V Chuang. Accurate and efficient structural ensemble generation of macrocyclic peptides using internal coordinate diffusion. *arXiv preprint arXiv:2305.19800*, 2023.

Colin A Grambow, Hayley Weir, Christian N Cunningham, Tommaso Biancalani, and Kangway V Chuang. Cremp: Conformer-rotamer ensembles of macrocyclic peptides for machine learning. *Scientific Data*, 11(1):859, 2024.

Luca J Hagemeyer, William TP Darling, Nele-Johanna Hempel-Højgaard, Jan Kihlberg, and Mate Erdelyi. Solution ensembles of griselimycin expand the definition of molecular chameleonicity. *Journal of Medicinal Chemistry*, 2026.

Minh Ngoc Ho, Jiayuan Miao, Yi Shan, Choi Yi Li, Hiroaki Suga, James D Baleja, and Yu-Shan Lin. Streamm-thioether: Efficient structure prediction for thioether-linked cyclic peptides. *The Journal of Physical Chemistry B*, 130(7):2031–2043, 2026.

Tiffani Hui, Marc L Descoteaux, Jiayuan Miao, and Yu-Shan Lin. Training neural network models using molecular dynamics simulation results to efficiently predict cyclic hexapeptide structural ensembles. *Journal of Chemical Theory and Computation*, 19(14):4757, 2023.

Paul Katzberger and Sereina Riniker. A general graph neural network based implicit solvation model for organic molecules in water. *Chemical Science*, 15(28):10794–10802, 2024. doi: 10.1039/d4sc02432j. URL https://doi.org/10.1039/d4sc02432j.

Paul Katzberger, Lea M Hauswirth, Antonia S Kuhn, Gregory A Landrum, and Sereina Riniker. Rapid access to small molecule conformational ensembles in organic solvents enabled by graph neural network-based implicit solvent model. *Journal of the American Chemical Society*, 147 (16):13264, 2025.

Danyeong Lee, Dohoon Lee, Dongmin Bang, and Sun Kim. DiSCO: Diffusion Schrödinger Bridge for Molecular Conformer Optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 13365–13373, 2024. doi: 10.1609/aaai.v38i12.29238. URL https://doi.org/10.1609/aaai.v38i12.29238.

Stephanie M. Linker, Christian Schellhaas, Benjamin Ries, Hans-Jörg Roth, Marianne Fouché, Stephane Rodde, and Sereina Riniker. Polar/apolar interfaces modulate the conformational behavior of cyclic peptides with impact on their passive membrane permeability. *RSC Advances*, 12(10):5782–5796, 2022. doi: 10.1039/d1ra09025a. URL https://doi.org/10.1039/d1ra09025a.

Stephanie M. Linker, Christian Schellhaas, Anna S. Kamenik, Mac M. Veldhuizen, Franz Waibl, Hans-Jörg Roth, Marianne Fouché, Stephane Rodde, and Sereina Riniker. Lessons for Oral Bioavailability: How Conformationally Flexible Cyclic Peptides Enter and Cross Lipid Membranes. *Journal of Medicinal Chemistry*, 66(4):2773–2788, 2023. doi: 10.1021/acs.jmedchem.2c01837. URL https://doi.org/10.1021/acs.jmedchem.2c01837.

Wei Liu, Nguyen Hung Pham, Chandra S Verma, Hwee Kuan Lee, and Jianguo Li. Multi-solvent conformational ensembles for predicting cyclic peptide permeability. *Scientific Data*, 2026.

Jiayuan Miao, Marc L Descoteaux, and Yu-Shan Lin. Structure prediction of cyclic peptides by molecular dynamics+ machine learning. *Chemical science*, 12(44):14927–14936, 2021.

Selma Moqvist, Weilong Chen, Mathias Schreiner, Feliks N"uske, and Simon Olsson. Thermodynamic interpolation: A generative approach to molecular thermodynamics and kinetics. *Journal of Chemical Theory and Computation*, 21(5):2535–2545, 2025.

Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019.

Vasanthanathan Poongavanam, Lianne H. E. Wieske, Stefan Peintner, Máté Erdélyi, and Jan Kihlberg. Molecular chameleons in drug discovery. *Nature Reviews Chemistry*, 8(1):45–60, 2024. doi: 10.1038/s41570-023-00563-1. URL https://doi.org/10.1038/s41570-023-00563-1.

Andrew Powers, P. Douglas Renfrew, Parisa Hosseinzadeh, and Vikram Khipple Mulligan. Cyclic-CAE: A Conformational Autoencoder for Efficient Heterochiral Macrocyclic Conformational Sampling. *bioRxiv*, 2025. doi: 10.1101/2025.02.21.639569. URL https://doi.org/10.1101/2025.02.21.639569. Preprint.

Yunrui Qiu, Richard John, Lukas Herron, and Pratyush Tiwary. Latent thermodynamic flows: unified representation learning and generative modeling of temperature-dependent behaviors from limited data. *Chemical Science*, 17(7):3743–3757, 2026.

Theresa A Ramelot, Jonathan Palmer, Gaetano T Montelione, and Gaurav Bhardwaj. Cell-permeable chameleonic peptides: Exploiting conformational dynamics in de novo cyclic peptide design. *Current opinion in structural biology*, 80:102603, 2023.

Luca Schaufelberger, Aline Hartgers, and Kjell Jorner. Generating cyclic conformers with flow matching in cremer-pople coordinates. *arXiv preprint arXiv:2601.12859*, 2026.

Maxim Secor, Jovan Damjanovic, Veronika Thost, Sebastian Swanson, Kristine Deibler, and Jesper Ferkinghoff-Borg. CPgen: Heterochiral Cyclic Peptide Ensemble Generation and Ensemble-Based Sequence Design. In *ICML 2026 Workshop on Generative and Agentic AI for Biology (GenBio)*, 2026. URL https://openreview.net/forum?id=eKUwFNn6Ku.

Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36:62183–62223, 2023.

Xuechen Tang, Janik Kokot, Franz Waibl, Monica L Fernández-Quintero, Anna S Kamenik, and Klaus R Liedl. Addressing challenges of macrocyclic conformational sampling in polar and apolar solvents: Lessons for chameleonicity. *Journal of Chemical Information and Modeling*, 63(22): 7107, 2023.

James Thornton, Michael Hutchinson, Emile Mathieu, Valentin De Bortoli, Yee Whye Teh, and Arnaud Doucet. Riemannian Diffusion Schrödinger Bridge, 2022. URL https://arxiv.org/abs/2207.03024. Preprint.

Alexander Y Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free schrödinger bridges via score and flow matching. In *International Conference on Artificial Intelligence and Statistics*, pp. 1279–1287. PMLR, 2024.

Shuzhe Wang, Jagna Witek, Gregory A. Landrum, and Sereina Riniker. Improving Conformer Generation for Small Rings and Macrocycles Based on Distance Geometry and Experimental Torsional-Angle Preferences. *Journal of Chemical Information and Modeling*, 60(4):2044–2058, 2020. doi: 10.1021/acs.jcim.0c00025. URL https://doi.org/10.1021/acs.jcim.0c00025.

Sicheng Wen, Yang Wang, and Yue Qian. Ensemblecycperm: Interpretable modeling of cyclic peptide permeability through solvent-dependent conformational ensembles. *Journal of Chemical Information and Modeling*, 2026.

Bing Yan, Anuroop Sriram, Zachary Ward Ulissi, Kyunghyun Cho, Ricky T. Q. Chen, and Benjamin Kurt Miller. EVA-Flow: Learning Shared Chemistry for Unified Environment-Aware Molecular Conformation Generation. *Transactions on Machine Learning Research*, 2026. URL https://bingyan.me/papers/EVA_Flow.pdf.

Yiping Yu, Qilong Wu, Junwei Chen, Yuli She, Lizhe Zhu, and Zhiye Guo. Delta psa: A new metric for conformational dynamics underlying macrocyclic peptide permeability. *bioRxiv*, pp. 2026–01, 2026.

Appendix

A Dataset and Evaluation Protocol

A.1 Dataset construction

Experiments use the stereochemistry-corrected release of CycPeptMPDB-4D (Liu et al., 2026), which contains solvent-specific molecular-dynamics conformations of cyclic peptides in water and hexane. The deposited collection contains 5,160 molecular entries. Coordinate atoms are matched to the molecular graph through graph isomorphism between the deposited coordinates and the corresponding molecular representation. This procedure succeeds for 5,057 molecules.

A total of 103 entries fail the mapping procedure. One hundred do not admit a complete graph isomorphism under the tested bond reconstruction settings, while three have inconsistent atom counts. These molecules are excluded before any model fitting or evaluation. The excluded subset is not assumed to be random, and all reported results therefore refer to the successfully mapped cohort.

Each successfully mapped molecule contains 100 deposited conformations in water and 100 in hexane. We use the deposited water and hexane trajectories from the stereochemistry-corrected CycPeptMPDB-4D release and follow the dataset publication for simulation details. The retained local provenance and the published force-field description are not fully consistent, so we do not restate a force field here. Several other simulation details are also unavailable in the retained analysis metadata. We therefore use the deposited conformations as empirical solvent-specific references and do not infer equilibrium populations or free energies from their observed frequencies.

A.2 Source and target blocks

The source ensemble contains frames 0–19 from the observed solvent. The held-out source reference and target reference contain frames 50–99 from the corresponding solvent. These two source-side quantities remain distinct throughout evaluation. The source input statistic is computed from frames 0–19, whereas response labels use the held-out source reference from frames 50–99.

For observable o, the target endpoint is

$$y_{t,o} = \frac{1}{|E_t|} \sum_{C \in E_t} \phi_o(C). \tag{21}$$

The canonical solvent response uses a fixed hexane-minus-water convention,

$$\Delta_o = y_{H,o} - y_{W,o}. \tag{22}$$

The same sign convention is used regardless of which solvent is observed at inference.

A.3 Grouped cross-validation

The primary prediction benchmark contains 3,456 molecules in 162 structural groups. Structural groups are constructed from exact structural identity, stereochemical scaffold, and analogue-series relationships followed by transitive closure. Both solvent trajectories and all frames from a molecule remain in the same group and outer fold.

Three grouped outer folds are used. Model fitting, feature normalization, early stopping, and training-set constants are computed only from the training portion of each fold. Every reported prediction on the primary benchmark is therefore out of fold. The final model fitted on the complete development cohort is used for deployment only and has no held-out score.

The generation development panel contains 96 molecules from 96 distinct structural groups. Panel membership is fixed before comparative generation experiments. Results on this panel are development results. CONF192 contains 192 molecules from 35 structural groups drawn from molecules that had not appeared in any generator run during development. Its selection rule was fixed before its generated candidates were opened. CONF192 therefore serves as a generator-naive confirmation within the same benchmark rather than an external dataset.

A.4 Development exposure and evaluation status

The grouped out-of-fold benchmark and the 96-molecule generation panel were used during method development. CONF192 was evaluated only after the final generator and evaluation definitions were frozen. The development and confirmation panels use the same source/reference blocks, structural thresholds, proposal budget, quality-control rules, and evaluation metrics.

A.5 Aggregation and uncertainty

Water-to-hexane and hexane-to-water results are reported separately. Equal-weighted summaries first average the two directions within each molecule and then average across molecules.

Confidence intervals use grouped bootstrap resampling. Structural groups are resampled with replacement, and paired effects are computed per molecule before resampling. Mean-based intervals use 4,000 bootstrap samples. Correlation intervals use 1,000 samples. An interval is reported as unavailable when fewer than four independent structural groups are present.

Frames and sampling replicates are never counted as independent molecules.

B Structural Observables and Evaluation Metrics

B.1 Target observables

Four ensemble observables are used for cross-solvent prediction and ensemble selection.

Radius of gyration. The radius of gyration is computed over the covered heavy atoms with uniform atom weights,

$$R_g = \sqrt{\frac{1}{n}\sum_i \|\mathbf{x}_i - \bar{\mathbf{x}}\|_2^2}. \tag{23}$$

Relative shape anisotropy. Let $\lambda_1, \lambda_2, \lambda_3$ denote the eigenvalues of the gyration tensor. Relative shape anisotropy is

$$\kappa^2 = \frac{3}{2}\frac{\sum_i \lambda_i^2}{\left(\sum_i \lambda_i\right)^2} - \frac{1}{2}. \tag{24}$$

Intramolecular contacts. A contact is counted when two covered heavy atoms are separated by at least two residues along the cyclic sequence and lie within 6.5 Å. Counts are normalized by the number of covered heavy atoms.

Hydrogen-bond proxy. Backbone donor-acceptor pairs separated by at least two residues are counted when the constructed H–O distance is below 2.5 Å and the N–H–O angle exceeds 120°. The count is normalized by residue number. This quantity is a geometric proxy rather than an experimental hydrogen-bond measurement.

B.2 Structural distribution metrics

Observable agreement and conformational-distribution agreement are evaluated separately.

The primary torsional metric uses a multi-bandwidth Gaussian maximum mean discrepancy over periodic internal-coordinate features represented by their sine and cosine values. The reported quantity is the biased squared MMD. This feature space does not contain the four structural observables used during ensemble selection.

A second structural readout compares intramolecular heavy-atom distance patterns. It is used as a complementary metric because it captures structural organization not directly optimized by the observable-matching objective.

Coverage and precision are computed through nearest-neighbor structural distances. Coverage measures the distance from each target-reference conformation to its nearest generated conformation. Precision reverses this direction.

For support analysis, the uncovered target fraction is

$$u_\tau(\mathcal{B}_t, E_t) = \frac{1}{|E_t|} \sum_{C \in E_t} \mathbb{1}\left[\min_{\widetilde{C} \in \mathcal{B}_t} d(C, \widetilde{C}) > \tau\right], \tag{25}$$

where τ is estimated from training data and fixed before evaluation.

Chemical quality control evaluates molecular graph consistency, stereochemistry, bond geometry, and steric clashes. Failed generations remain in the generation-yield denominator.

B.3 Empirical target-mode recovery

Empirical modes are defined only from held-out trajectory references. Target frames are covered by greedy medoids in the frozen torsional distance at radius τ. The mode scale is fixed on training molecules before evaluation. A major target mode contains at least 10% of the 50 target-reference frames. A target-emergent mode is major in the target reference and has at most 2% occupancy in the held-out source reference. Generated conformers are assigned to the nearest mode within the same radius, with all other conformers retained in an unassigned state.

We report proposal-level major-mode and emergent-mode recall, $K = 20$ major-mode recall, total variation over mode occupancies plus the unassigned state, Jensen–Shannon divergence, and target-mode precision. These are empirical state summaries rather than kinetic or metastable-state estimates. Population metrics are interpreted cautiously because the 50-frame reference is finite and shows non-negligible variation across internal trajectory splits.

B.4 Coverage, precision, and reach

Target uncovered fraction, proposal precision, and outside-source reach measure different properties. The uncovered fraction asks whether target-reference structures have nearby candidates in the proposal bank. Proposal precision asks whether proposed structures lie near the target reference. Outside-source reach asks whether selected conformations extend beyond the observed source ensemble. A proposal bank can improve target coverage without increasing outside-source reach after selection, so these quantities are reported separately.

C Implementation Details

C.1 Cross-solvent predictor

The deployed cross-solvent predictor uses molecular descriptors together with permutation-invariant statistics computed from the source ensemble.

Molecular identity is represented by 12 physicochemical descriptors and a 1,024-bit radius-2 Morgan fingerprint with chirality, giving 1,036 molecular features. The source ensemble contributes 36 geometry features. Twenty correspond to the mean, standard deviation, 10th percentile, and 90th percentile of five per-conformation structural quantities. The remaining features describe molecular size, global geometry, and conformational dispersion.

Prediction uses histogram gradient-boosted regression. A separate estimator is fitted for each solvent transition and output observable. The regression target is the residual between the target endpoint and the source-block value,

$$r_{s \to t, o} = y_{t,o} - o(E_s). \tag{26}$$

The predicted target endpoint is

$$\widehat{y}_{t,o} = o(E_s) + \widehat{r}_{s \to t, o}. \tag{27}$$

Each estimator uses a maximum depth of 4, learning rate 0.06, and at most 300 boosting iterations. Early stopping uses a training-side validation split. Geometry features are standardized using training-fold statistics.

C.2 FINITE-SUPPORT ENSEMBLE SELECTION

For a candidate bank $\mathcal{B}_t = \{C_j\}_{j=1}^N$, weights minimize

$$\frac{1}{d}\sum_{o=1}^{d}\left(\frac{\sum_j w_j \phi_o(C_j) - \widehat{y}_{t,o}}{\sigma_o}\right)^2 + \lambda D_{\mathrm{KL}}\left(\mathbf{w}\|\mathbf{u}\right), \tag{28}$$

subject to

$$w_j \geq 0, \qquad \sum_j w_j = 1, \qquad \mathrm{ESS}(\mathbf{w}) = \frac{1}{\sum_j w_j^2} \geq \frac{N}{4}. \tag{29}$$

Here σ_o is estimated from target endpoints in the training fold and $\mathbf{u}$ is the uniform distribution over candidates. We use $\lambda = 0.1$. The predicted target vector is computed once from the original source ensemble and remains fixed during selection.

The target-observing oracle uses the same optimization with measured target observables replacing predicted values. It is used only as a diagnostic and does not define an attainable physical or statistical ceiling.

C.3 BROAD CONFORMER PRIOR

The broad source-independent proposal prior is generated with ETKDGv3 followed by MMFF94 relaxation. The resulting bank contains up to 200 quality-controlled conformations per molecule and is independent of the requested solvent. This prior supplies broad conformational support for the target-aware bridge rather than a predicted target ensemble.

C.4 TARGET-AWARE SCHRÖDINGER BRIDGE

The bridge operates on cyclic internal coordinates. One token represents each cyclic ring position. Flexible torsional channels are encoded periodically with sine/cosine coordinates, while molecular graph, topology, atom identity, stereochemistry, D/L identity, N-methylation, bond orders, and ring size remain fixed conditioning variables.

The bridge uses a cyclic Transformer with width 256, six pre-norm blocks, eight attention heads, cyclic relative-position bias, GEGLU feed-forward layers, adaptive layer normalization, and a zero-initialized velocity head. A separate source-set encoder uses two cyclic Transformer blocks with width 128. Source conformations are summarized through permutation-invariant pooling and projected to the source context representation. Training samples multiple source-set sizes to support variable numbers of observed conformations.

Conditioning combines diffusion time, molecular identity, the source-set representation, and an explicit source-to-target solvent-transition embedding. The transition embedding distinguishes the four water/hexane source-target combinations and a null condition.

Training target conformers are sampled independently from the target-solvent trajectory. Source and target trajectory frames are not paired by frame index. Broad-prior and target samples are coupled through an entropy-regularized structural transport plan using a frozen cost that combines periodic torsional and intramolecular-distance discrepancies. The four observable targets used by the selector are not part of this bridge coupling cost.

The conditional path follows a wrapped Brownian bridge in periodic torsion space. Let $\Delta_{\mathbb{T}}(x_0, x_1)$ denote the shortest-arc displacement on the torus. A bridge sample can be written schematically as

$$x_\tau = \mathrm{wrap}\left(x_0 + \tau\Delta_{\mathbb{T}}(x_0, x_1) + \sigma_\tau \epsilon\right), \qquad \epsilon \sim \mathcal{N}(0, I), \tag{30}$$

where σ_τ follows the Brownian-bridge schedule and vanishes at the endpoints. The conditional velocity target $u_\tau(x_\tau \mid x_0, x_1)$ is determined by this path. The velocity field minimizes

$$\mathcal{L}_{\mathrm{SB}} = \mathbb{E}\left[\left\|v_\psi(x_\tau, \tau \mid G, E_s, s, t) - u_\tau(x_\tau \mid x_0, x_1)\right\|_2^2\right], \tag{31}$$

following the simulation-free Schrödinger-bridge flow-matching formulation of Tong et al. (2024). The final training recipe evaluates a differentiable steric-clash feasibility term on the same source-carried geometry used at deployment. The training objective is

$$\mathcal{L} = \mathcal{L}_{\mathrm{SB}} + \lambda_{\mathrm{clash}}\mathcal{L}_{\mathrm{clash}}. \tag{32}$$

Inference uses probability-flow ODE sampling. The evaluated bridge therefore uses stochastic bridge training without adding stochastic SDE noise during final sampling.

The bridge generates the flexible torsional channels. Non-generated internal-coordinate channels are supplied by one legal observed source conformer before reconstruction. Canonical reconstruction uses the same cyclic closure and frozen quality-control procedure for all bridge variants. This source-consistent donor is target blind and uses only legal inference-time source structures.

D Training and Inference Algorithms

We summarize the training and inference procedures of SolvMorph. The training procedure learns a source- and target-conditioned Schrödinger bridge in periodic torsion space, while inference generates target-specific candidate conformations followed by observable-guided ensemble refinement.

Algorithm 1 Training target-aware Schrödinger bridge

Require: Molecule graph G, source ensemble $\mathcal{E}_s$, solvents (s, t)
Require: Prior $q_0(C|G)$, target distribution $p_t(C|G)$
Require: Transport cost $c(\cdot, \cdot)$

Compute entropy-regularized coupling

$$\pi^* = \arg \min_{\pi \in \Pi(q_0, p_t)} \mathbb{E}_{(C_0, C_1) \sim \pi}[c(C_0, C_1)] + \epsilon D_{\mathrm{KL}}(\pi \| q_0 \otimes p_t)$$

while not converged **do**
  Sample $(C_0, C_1) \sim \pi^*$
  Let $x_0 = x(C_0), x_1 = x(C_1)$
  Sample bridge time $r \sim U(0, 1)$
  Sample wrapped bridge state

$$x_r \sim p(x_r | x_0, x_1)$$

  Compute conditional velocity target

$$u_r(x_r | x_0, x_1)$$

  Update v_ψ by minimizing

$$\mathbb{E} \| v_\psi(x_r, r | G, \mathcal{E}_s, s, t) - u_r(x_r | x_0, x_1) \|_2^2 + \lambda_{\mathrm{clash}} \mathcal{L}_{\mathrm{clash}}$$

end while

Ensure: Learned velocity field v_ψ

E Additional Cross-Solvent Prediction Analyses

The following analyses examine whether source conformations provide transferable information for predicting unobserved target-solvent ensemble structural changes.

E.1 Feature and model baselines

As shown in Table 1, source-informed predictors improve over chemistry-only baselines, demonstrating that observed source conformations contain transferable information about solvent-dependent structural changes.

The combined geometry and chemistry model achieves the strongest overall endpoint prediction, supporting the use of source ensemble information in cross-solvent inference.

E.2 Neural set encoders

As shown in Table 2, more expressive set encoders do not consistently improve over the simpler source-statistics predictor, motivating the compact predictor used in the main experiments.

The main text therefore uses the simpler gradient-boosted model as the cross-solvent predictor.

Algorithm 2 Inference with target-aware generation and observable-guided refinement

Require: Molecule graph G, source ensemble $\mathcal{E}_s$, solvents (s, t)
Require: Trained bridge v_ψ, predictor f_θ

1: Sample initial conformers

$$C_{0,j} \sim q_0(C|G), \quad j = 1, \ldots, M$$

2: Generate target torsions by integrating

$$\frac{dx_r}{dr} = v_\psi(x_r, r|G, \mathcal{E}_s, s, t)$$

3: Reconstruct conformations

$$B_t = \{C_j\}_{j=1}^N$$

using generated torsions and source-derived internal coordinates
4: Apply chemical quality control and retain valid candidates
5: Predict target observables

$$\hat{m}_t = f_\theta(G, \mathcal{E}_s, s, t)$$

6: Optimize candidate weights

$$w^* = \arg \min_{w \in \Delta^N,\ ESS(w) \geq N/4} \left[\frac{1}{d} \sum_o \left(\frac{\sum_j w_j \phi_o(C_j) - \hat{m}_{t,o}}{\sigma_o} \right)^2 + \lambda D_{\mathrm{KL}}(w||u) \right]$$

7: Sample final ensemble

$$\hat{\mathcal{E}}_t \sim \sum_j w_j^* \delta_{C_j}$$

Ensure: Target-solvent ensemble $\hat{\mathcal{E}}_t$

Table 1: Target-endpoint MAE for representative prediction baselines. Values are equal-weighted over the two cross-solvent directions.

Method	R_g	κ^2	Contacts	HB proxy
Persistence	0.1590	0.0401	0.7652	0.1201
Train-mean response	0.1404	0.0381	0.6767	0.0965
Train-median response	0.1384	0.0378	0.6570	0.0955
Chemistry-only endpoint	0.1598	0.0389	0.6751	0.0741
Source observables, Ridge	0.1219	0.0344	0.5771	0.0808
Chemistry, Ridge	0.1394	0.0380	0.6643	0.0876
Source geometry, Ridge	0.1166	0.0339	0.5778	0.0766
Source geometry, boosting	0.1207	0.0333	0.5663	0.0753
Geometry + chemistry, Ridge	0.1238	0.0346	0.5471	0.0730
Geometry + chemistry, boosting	**0.1137**	**0.0330**	**0.5371**	**0.0710**

E.3 Response correlations

As shown in Table 3, predicted solvent responses show positive rank correlations with measured changes across both solvent directions. We therefore report correlation as a complementary diagnostic, while retaining endpoint MAE as the primary metric.

E.4 Number of source conformations

We additionally examine sensitivity to source ensemble size. Matched refitting suggests that a single source conformation retains much of the predictive signal, while additional source conformations provide modest improvements. This analysis does not imply that a fixed number of source structures is required.

Table 2: Target-endpoint MAE for neural prediction variants. These models use the same grouped evaluation protocol as the final predictor.

Method	R_g	κ^2	Contacts	HB proxy
MLP on ensemble statistics	0.1192	0.0331	0.5413	0.0717
Graph only	0.2119	0.0404	0.6788	0.0806
Graph + source scalars	0.1169	0.0325	0.5518	0.0705
DeepSets	0.1179	0.0328	0.5521	0.0707
Set Transformer	0.1177	0.0328	0.5466	0.0707
DeepSets + geometry statistics	0.1176	0.0334	0.5359	0.0725
Set Transformer + geometry statistics	0.1165	0.0332	0.5312	0.0708
Validation-selected linear blend	0.1152	0.0325	0.5437	0.0713

Table 3: Spearman correlation between predicted and measured solvent responses.

Direction	R_g	κ^2	Contacts	HB proxy
Water $\rightarrow$ Hexane	0.519	0.441	0.438	0.480
Hexane $\rightarrow$ Water	0.427	0.385	0.520	0.679

F Fixed-Support Ensemble Selection

As shown in Table 4, observable-guided selection improves target endpoint observables while leaving the candidate support unchanged.

Table 4: Equal-weighted target-endpoint MAE after reweighting the unchanged source conformations.

Selection	R_g	κ^2	Contacts	HB proxy
Uniform source copy	0.1590	0.0401	0.7652	0.1201
Target-blind selection	0.1388	0.0359	0.6723	0.0962
Cross-solvent guided selection	**0.1340**	**0.0351**	**0.6490**	**0.0890**

The largest improvement is observed for compactness-related observables, supporting that probability refinement can improve ensemble averages without expanding structural support.

As shown in Table 5, fixed-support reweighting does not recover distribution-level agreement because it cannot generate conformations outside the original candidate bank.

These results distinguish probability refinement from support expansion, motivating target-aware generation.

G Matched-Support and Distribution-Learning Ablations

As shown in Table 6, adding observed source conformations improves reconstruction under matched candidate budgets, indicating that source observations contribute structural coverage rather than only probability refinement. The main improvement comes from expanding the candidate support: adding source conformers substantially reduces uncovered target regions compared with a source-independent candidate bank.

Direct distribution supervision and conditioning controls further support the separation between support generation and probability refinement. Directly predicting target distributions does not improve held-out reconstruction, and target conditioning alone is insufficient without appropriate candidate support.

Table 5: Structural-distribution results for fixed-support selection on the generation panel.

Selection	Torsion MMD^2	Useful distinct	Outside-source reach
Uniform source copy	0.985	11.38	0.000
Target-blind selection	0.990	11.44	0.000
Cross-solvent guided selection	1.017	10.61	0.000
Target-observing oracle	1.021	10.15	0.000

Table 6: Matched-support structural reconstruction. OBS denotes observable-guided selection and DIST denotes direct target-distribution supervision.

	Torsion MMD^2		Distance MMD^2		R_g MAE	
Method	W→H	H→W	W→H	H→W	W→H	H→W
Generic + OBS	0.9173	0.8540	1.0837	0.9766	0.0758	0.1030
Mixed + OBS	**0.7951**	**0.7323**	**0.9580**	**0.8727**	0.0785	0.1040
Generic + DIST	0.9724	0.8837	1.1010	1.0227	0.0969	0.1267
Mixed + DIST	0.9010	0.8595	1.1551	1.1295	0.1036	0.1133

H TARGET-AWARE GENERATION

H.1 GENERATION PROTOCOL AND HELD-OUT TARGET EVALUATION

Generation development uses the same 96-molecule panel as the reconstruction experiments. CONF192 provides a separate generator-naive confirmation. Source frames are available to the generator, while held-out target-reference frames are evaluator only. The broad prior uses a fixed candidate budget, and failed generation attempts remain in the raw-yield denominator. Both cross-solvent directions are evaluated independently.

H.2 PROPOSAL-LEVEL COMPARISON

As shown in Table 7, target-aware generation expands conformational coverage beyond fixed candidate banks before observable refinement.

Table 7: Proposal-bank results before observable-guided selection for the target-aware bridge snapshot used to establish the support-generation result. Lower uncovered fraction and higher precision are better. Raw yield is reported only for generated proposals; fixed Cartesian banks are marked n.c. because they bypass cyclic reconstruction.

Proposal bank	Raw yield W→H / H→W	u_τ W→H / H→W	Precision@τ W→H / H→W
Generic prior	n.c.	0.7348 / 0.7652	0.0124 / 0.0119
Mixed source + generic	n.c.	0.4517 / 0.4644	0.0624 / 0.0659
Wrong-target bridge	0.7398 / 0.7302	0.3867 / 0.3953	0.4993 / 0.5263
Target-aware bridge	0.7221 / 0.7356	**0.3449 / 0.3374**	**0.5507 / 0.5431**
Real target MD reference	n.c.	0.0417 / 0.0954	0.975 / 0.943

Compared with generic and mixed-support proposals, the target-aware bridge substantially reduces uncovered target regions, demonstrating recovery of conformational support unavailable to fixed banks.

H.3 GENERATION SPECIFICITY AND REFINEMENT

Requested-target specificity is evaluated before selection. The primary paired statistic compares the correct requested target condition with a swapped target condition for the same source molecule and proposal budget, which isolates sensitivity to the requested solvent condition. The final bridge remains sensitive to the requested target condition, supporting solvent-specific generation rather than generic conformational expansion.

H.4 CHEMICAL FEASIBILITY AND FAILURE ANALYSIS

As shown in Table 8, source-carried geometry provides the dominant contribution to chemical validity, while clash-aware supervision further improves generation quality. Failure analysis indicates that residual invalid proposals are mainly associated with steric clashes rather than cyclic closure.

Table 8 decomposes the contribution of reconstruction geometry, self-consistent training, and clash-aware supervision. The results show that validity gains primarily originate from chemically consistent reconstruction, with feasibility supervision providing an additional refinement.

Table 8: Factors controlling chemical validity on the generation development panel.

Method	Valid W→H	Valid H→W	Δu_τ W→H	Δu_τ H→W
Frozen bridge + SC decode	0.7221	0.7356	–	–
SC training, no clash	0.7108	0.7352	–	–
SC + clash, scratch	0.7692	0.7817	–	–
SC + clash, final	**0.7604**	**0.7868**	+0.0095	+0.0006

H.5 SELECTION EFFECTS AND MODE-RECOVERY CONTROLS

Observable-guided selection provides additional refinement after target-aware generation. Table 9 summarizes this downstream effect, while Table 10 reports proposal-budget and target-condition controls for target-emergent mode recovery.

Table 9: Effect of observable-guided selection on target-aware generated ensembles.

Method	Direction	Torsion MMD
Source copy	W→H	1.217
Target-aware after selection	W→H	0.944
Source copy	H→W	1.223
Target-aware after selection	H→W	0.841

Table 10: Target-emergent mode recovery under proposal-budget and target-condition controls.

Condition	W→H	H→W
Source ensemble	0.079	0.057
SolvMorph (20 proposals)	0.169	0.237
SolvMorph (full proposals)	0.336	0.410

H.6 EMPIRICAL MODE RECOVERY AND POPULATION LIMITS

As shown in Table 10, SolvMorph recovers target-emergent modes beyond the observed source ensemble while retaining the advantage under proposal-budget matching. The reduction under wrong-target conditioning supports solvent-specific mode recovery rather than generic conformational expansion.

Population recovery remains limited by finite trajectory references. We therefore interpret mode discovery as the primary generation result rather than claiming accurate recovery of target populations.

I SUPPLEMENTARY GENERATION ABLATION RESULTS

Additional generation ablations are summarized in Figure 8. Adding source conformers improves candidate coverage relative to a source-independent prior. Target conditioning changes proposal placement, while wrong-target controls show that the generated support depends on the requested

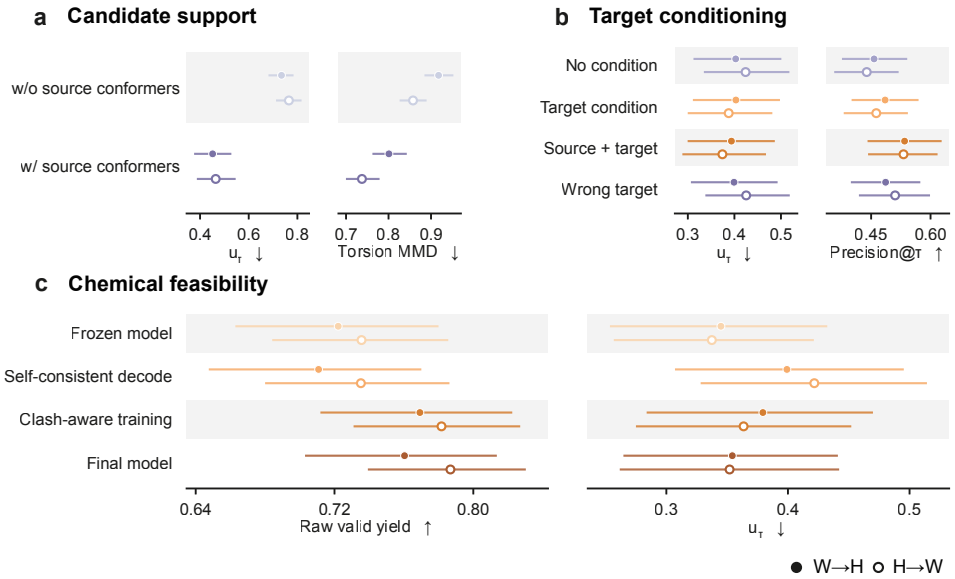


Figure 8: **Supplementary generation ablations on PILOT96.** Additional ablations for candidate support, target conditioning, and chemical feasibility.

solvent condition. Chemical-feasibility ablations further show that clash-aware training improves validity while maintaining competitive coverage.