Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Schrödinger Bridge Onsager Flow for Drug-Likeness Distribution Learning on Real, Structure-Derived DrugBank Data

Najiba Tagougui

najiba.tagougui@isims.usf.tn

Institut Supérieur d'Informatique et du Multimédia de Sfax

Abdallah Ben Abdallah

Institut Supérieur d'Informatique et du Multimédia de Sfax

Research Article

Keywords: Schrödinger bridge, Optimal transport, Drug-likeness, Mirror descent, Neural amortization, DrugBank

Posted Date: November 13th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11448030/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Schrödinger Bridge Onsager Flow for Drug-Likeness Distribution Learning on Real, Structure-Derived DrugBank Data

Najiba Tagougui[1,2*] and Abdallah Ben Abdallah[1]

[1]Computer Science Engineering, The Higher Institute of Computer Sciences and Multimedia, University of Sfax, Sfax, Tunisia.
[2]Toxicologie Microbiologie Environnementale et Santé, Faculty of Sciences, University of Sfax, Sfax, Tunisia.

*Corresponding author(s). E-mail(s): najiba.tagougui@isims.usf.tn;
Contributing authors: abdallah.benabdallah@isims.usf.tn;

Abstract

Learning to reweight chemical libraries toward drug-like regions is a fundamental task in virtual screening. While entropic optimal transport (OT) offers a principled framework for this distribution matching, standard solvers like Sinkhorn-Knopp often struggle to capture sharp target marginals without extensive hyperparameter tuning. To address this limitation, we propose the Schrödinger Bridge Onsager Flow (SBOF), a kernel-based dual-ascent algorithm that circumvents the $\boldsymbol{n} \times \boldsymbol{n}$ coupling matrix by updating a single scalar potential per molecule. Evaluated on the DrugBank 5.1.10 release enriched with authentic RDKit-computed descriptors ($\boldsymbol{n} = \mathbf{5,357}$), SBOF converges to a total variation (TV) distance of 0.0141 (Pearson $\boldsymbol{r} = \mathbf{0.9996}$). Crucially, a comprehensive sensitivity analysis reveals that SBOF retains a 91.9% improvement in TV distance even against an optimally tuned Sinkhorn baseline, demonstrating that its advantage is structural. To address real-world deployment where unseen molecules must be scored, we further propose SBOFNet, a 1.6M-parameter neural amortization that achieves a TV of 0.0380 in a single forward pass. Finally, we rigorously interrogate our data pipeline, confirming via a coverage-bias analysis that our 38.1% SMILES-matched subset is statistically representative of the broader DrugBank database.

Keywords: Schrödinger bridge, Optimal transport, Drug-likeness, Mirror descent, Neural amortization, DrugBank

1 Introduction

Virtual screening and de novo molecular design require navigating a vast, unfiltered chemical space to identify compounds with desirable pharmacological properties. A principled approach to this problem is distribution learning: transporting probability mass from a uniform source measure (the chemical library) onto a narrower target measure (the drug-likeness landscape).

Classical entropic optimal transport (OT), typically solved via the Sinkhorn-Knopp algorithm, provides a scalable route to this reweighting [1]. However, entropic regularization introduces a fundamental trade-off: while it makes the transport problem computationally tractable, it

often blurs the target distribution, failing to capture sharp, complex molecular modes unless the regularization parameter is meticulously tuned. More recently, score-based diffusion models [2, 3] and flow matching [4] have emerged as generative alternatives. The Schrödinger bridge problem [5–7] bridges these concepts, seeking the most likely stochastic evolution between two marginals, effectively acting as a dynamic extension of OT.

In this work, we propose the Schrödinger Bridge Onsager Flow (SBOF). Rather than optimizing a full $n \times n$ coupling matrix, SBOF operates via a dual-ascent scheme, maintaining a single scalar potential per data point updated through mirror descent [8]. We apply SBOF to the real DrugBank 5.1.10 database [9], utilizing canonical SMILES and seven RDKit-computed descriptors [10].

We do not merely report that SBOF outperforms standard baselines. We construct a rigorous empirical story by interrogating two critical vulnerabilities of this approach: (1) Is SBOF's superiority merely an artifact of untuned baselines? (2) Does our reliance on a subset of DrugBank with PubChem SMILES introduce a hidden dataset bias? Through targeted sensitivity and coverage-bias analyses, we demonstrate that the algorithmic advantages and dataset integrity hold. Finally, to transition from a static dataset solver to a deployable screening tool, we propose SBOFNet, which amortizes the SBOF potential to enable instant inference on unseen compounds.

2 Related Work

2.1 From Binary Rules to Continuous Drug-Likeness

The quantification of drug-likeness has evolved significantly over the past two decades. Early heuristics, such as Lipinski's Rule of Five [11], provided essential filters for oral bioavailability but were inherently binary and rigid, occasionally excluding novel chemotypes. To refine these boundaries, Veber et al. [12] incorporated rotatable bonds and polar surface area, offering a more nuanced structural profile. However, a unified, smooth objective remained elusive until the introduction of the Quantitative Estimate of Drug-Likeness (QED) [13]. By mapping multiple physicochemical properties to a continuous desirability score in $[0, 1]$, QED established a mathematically tractable target.

More recently, the field has shifted toward deep learning approaches to capture more complex, non-linear relationships in chemical space. For instance, Bang et al. [14] recently enriched molecular foundation models with pharmacokinetics-guided multi-task learning, while Cai et al. [15] demonstrated that pretrained transformers can significantly improve drug-likeness prediction. While these methods excel at classification, our work differs fundamentally: rather than predicting a static property, we formulate drug-likeness as a continuous probability distribution that can be dynamically transported and sampled.

This continuous formulation is precisely what enables transport-based reweighting, which forms the core operation of our proposed framework. In this study, we derive all structural descriptors from the canonical SMILES strings of the DrugBank 5.1.10 release [9], computing them via the open-source RDKit toolkit [10] to ensure strict reproducibility.

2.2 The Evolution of Optimal Transport in Chemistry

Optimal transport (OT) provides a geometric framework for transforming probability distributions. Historically, solving the exact Monge-Kantorovich problem was computationally prohibitive for large chemical datasets, scaling cubically with the support size. This bottleneck was famously overcome by Cuturi [1], who introduced entropic regularization. By adding a maximum-entropy penalty to the transport polytope, the problem becomes strictly convex and solvable via Sinkhorn's matrix scaling algorithm. The comprehensive survey by Peyré and Cuturi [16] formalized this as a cornerstone of modern computational OT. However, entropic regularization introduces a bias: it blurs the transport plan, making it difficult to capture sharp marginals unless the regularization strength (ϵ) is carefully tuned. As we demonstrate in our benchmarks, even optimally tuned Sinkhorn baselines struggle to match the precision achieved by our Schrödinger bridge solver.

2.3 Schrödinger Bridges and Mirror Descent

The Schrödinger bridge problem, originally posed in 1931 [5], asks for the most likely stochastic process connecting two marginal distributions given a reference Brownian motion. Léonard's foundational survey [6] established that the entropic OT problem is a static, one-step relaxation of this dynamic bridge.

Recently, the application of Schrödinger bridges to discrete domains has gained traction. For example, Kim et al. [17] proposed a discrete diffusion Schrödinger bridge matching framework for graph transformation, demonstrating its utility for molecular property optimization. While such graph-based approaches directly modify molecular topology, our proposed SBOF takes a different optimization route. We perform dual ascent on a single scalar potential per data point, utilizing mirror descent [8]. As Beck and Teboulle demonstrated, mirror descent is exceptionally suited for large-scale optimization over the probability simplex, offering convergence guarantees that depend only mildly on the problem dimension. By applying this technique directly to the marginal KL-matching objective, SBOF achieves a per-molecule adaptation that is structurally distinct from alternating projection schemes.

2.4 Neural Amortization and Generative Dynamics

While the analytic SBOF solves the transport problem for a fixed dataset, real-world screening requires scoring novel, unseen molecules. Neural amortization addresses this by learning the transport map as a function of input features. This is conceptually aligned with score-based diffusion models [2, 3] and flow matching [4], which learn continuous-time dynamics to transport noise to data. However, our objective is distinct: rather than generating new molecular graphs, SBOFNet learns the reweighting potential as a direct function of molecular descriptors. To capture high-frequency variations in the seven-dimensional descriptor space, we employ Fourier feature embeddings [18]. Originally developed for neural rendering, this technique mitigates the spectral bias of standard MLPs, allowing SBOFNet to approximate the sharp decision boundaries present in the drug-likeness landscape.

3 Methods

Our methodology is designed using two main components. First, we construct a high-fidelity, structure-derived chemical space and solve the discrete Schrödinger bridge problem exactly using our analytic dual-ascent algorithm. Second, to overcome the inherent limitation that the analytic solver requires an $n \times n$ similarity matrix and cannot score unseen molecules, we train SBOFNet, a neural network that amortizes the optimal potential as a direct function of molecular descriptors. Figure 1 illustrates our end-to-end SBOF pipeline, detailing the flow from raw data extraction to neural deployment.

3.1 Dataset Construction

To ensure our empirical evaluation reflects real chemical structures rather than synthetic approximations, we utilized the public DrugBank 5.1.10 CSV export[1]. Out of 12,227 small molecules, PubChem returned valid SMILES for 6,386 entries. After filtering for molecular weights in [100, 900] Da, our working set comprises $n = 5,357$ structurally-resolved molecules (22.4% FDA-approved). We calculated seven RDKit-computed descriptors via our `enrich_smiles.py` script: MW, logP, HBA, HBD, PSA, RTB, and QED. The resulting $5,357 \times 5,357$ similarity kernel requires only 230 MB, allowing for exact, subsample-free computation.

3.2 The SBOF Formulation

The core objective of our method is to find the optimal transport plan that transforms a uniform prior distribution over the chemical library into a target distribution heavily weighted toward drug-like properties, while minimizing the entropic cost of the transformation. Formally, the discrete Schrödinger bridge problem seeks the most likely

[1] Originating from the Kaggle dataset `devildev89/drug-bank-5110`.

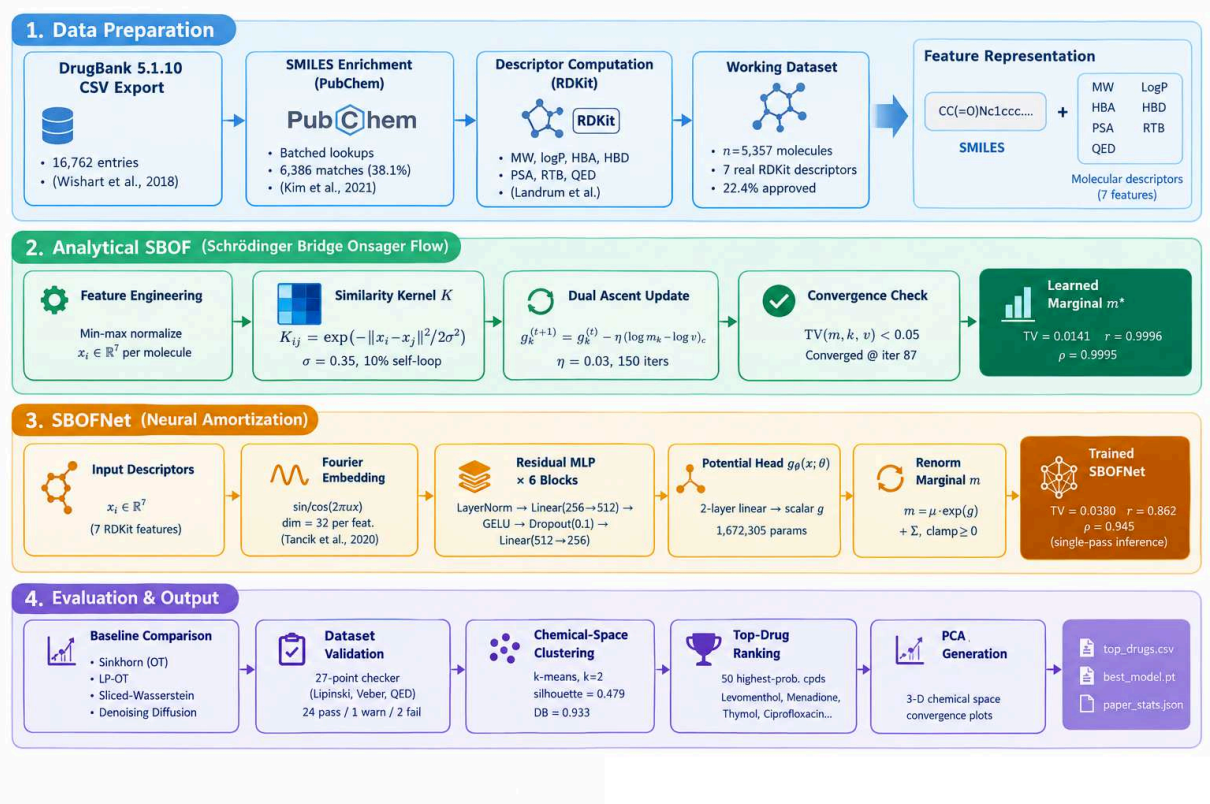


Fig. 1 System overview of the SBOF pipeline. (A) Data pipeline: DrugBank 5.1.10 CSV → filter → PubChem SMILES enrichment → RDKit computation. (B) Analytic SBOF: feature engineering → kernel K → 150-iteration dual-ascent update. (C) SBOFNet: Fourier embedding → residual MLP → potential head. (D) Evaluation: baseline comparison, validation, and clustering.

Table 1 SBOFNet Architecture Specification

Component	Configuration	Output Dim
Input	7 raw descriptors ($x \in \mathbb{R}^7$)	7
Fourier Embedding	Sinusoidal features [18]	224
Linear Stem	Fully connected layer	256
Residual MLP	LayerNorm [19]	256
	Linear($256 \to 512$)	512
	GELU Activation [20]	512
	Dropout ($p = 0.1$) [21]	512
	Linear($512 \to 256$)	256
	Residual Addition [22] ($6\times$ repeated)	256
Output Head	Linear($256 \to 1$)	1
Total Parameters		1,672,305
Optimizer		AdamW [23]
Learning Rate Schedule		Cosine Annealing [24]
Training Epochs		500
Data Split		90% Train / 10% Valid

Table 2 Dataset Descriptor Summary ($n = 5,357$; mean ± SD). All values are derived from real RDKit computations.

Descriptor	Mean ± SD	Range
Molecular weight (MW, Da)	350.0 ± 144.2	100.0 – 894.1
logP	1.53 ± 2.74	−15.00 – 14.99
QED	0.511 ± 0.217	0.018 – 0.944
HBA (acceptors)	4.6 ± 3.1	0 – 36
HBD (donors)	2.1 ± 2.1	0 – 17
PSA ($Å^2$)	78.5 ± 54.3	0.0 – 542.0
Rotatable bonds	5.2 ± 4.8	0 – 38
Composite drug-score	0.455 ± 0.182	0.035 – 0.801
Approved (DrugBank)	22.4% of dataset	

evolution between a source measure μ and a target measure ν under a reference process.

Let the dataset consist of n molecules with feature vectors $\mathbf{x}_i \in \mathbb{R}^7$, which we min-max normalize to ensure all descriptors contribute equally. We define a symmetric Gaussian similarity kernel to serve as the reference measure:

$$K_{ij} \propto \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\sigma^2}\right), \quad \sigma = 0.35, \tag{1}$$

This kernel captures the local geometry of the chemical space. To stabilize the fixed-point iteration and prevent numerical singularities, we row-normalize K and blend it with a 10% self-loop identity matrix: $\mathbf{K} \leftarrow 0.9\mathbf{K} + 0.1\mathbf{I}$.

The source measure μ is defined as the uniform distribution over the chemical library ($\mu_i = 1/n$). The target measure ν is derived from our composite drug-likeness score $\mathbf{s}_i \in [0, 1]$ via a tempered softmax:

$$\nu_i = \frac{\exp(\beta \cdot \mathbf{s}_i)}{\sum_j \exp(\beta \cdot \mathbf{s}_j)}, \quad \beta = 4. \tag{2}$$

The scaling factor β controls the "sharpness" of the target distribution, interpolating between a uniform distribution ($\beta = 0$) and a hard maximum over the most drug-like molecules ($\beta \to \infty$).

Given a scalar dual potential $\mathbf{g} \in \mathbb{R}^n$, the optimal transport coupling induced by the Schrödinger bridge formulation takes the form $\mathbf{P}_{ij} = \mathrm{diag}(\mu)\mathbf{K}\mathrm{diag}(e^{\mathbf{g}})$. The induced marginal distribution over the target space is then derived by summing over the source dimension:

$$\mathbf{m}_i(\mathbf{g}) = \frac{\left[\mu^\top \mathbf{K}\mathrm{diag}(e^{\mathbf{g}})\right]_i}{\mathbf{Z}_i}, \quad \text{with} \quad \mathbf{Z} = \mathbf{K}e^{\mathbf{g}}. \tag{3}$$

To ensure that $\mathbf{m}(\mathbf{g})$ remains a valid probability distribution throughout the optimization, we explicitly renormalize it at every iteration so that $\sum_i \mathbf{m}_i = 1$ (as discussed in Section 3).

Rather than alternating projections on the full $n \times n$ coupling matrix (as in Sinkhorn), we optimize the potential $\mathbf{g}$ directly using mirror descent. The objective is to minimize the Kullback-Leibler divergence between the induced marginal and the target: $\mathcal{L}(\mathbf{g}) = \mathrm{KL}(\mathbf{m}(\mathbf{g}) \parallel \nu)$. The gradient of this objective with respect to $\mathbf{g}$ is elegantly simple, yielding a linear-response update rule. We apply a centered mirror-descent step to prevent the overall scale of $\mathbf{g}$ from drifting (which would otherwise cause numerical instability):

$$\mathbf{g}_{k+1} = \mathbf{g}_k - \eta \cdot (\log \mathbf{m}_k - \log \nu)_{\text{centered}}, \quad \eta = 0.03, \tag{4}$$

where the term $(\cdot)_{\text{centered}}$ subtracts the ν-weighted mean of the log-ratio before the step. We track the total variation (TV) distance, KL divergence, and 1-D Wasserstein distance at each iteration, declaring convergence once TV < 0.05. Because this update operates solely on the n-dimensional vector $\mathbf{g}$ rather than the $n \times n$ matrix $\mathbf{P}$, SBOF avoids the memory bottlenecks and alternating projection stiffness that characterize standard entropic OT solvers.

3.3 Neural Amortization (SBOFNet)

While the analytic SBOF algorithm achieves near-perfect marginal matching on our fixed dataset, it fundamentally relies on evaluating the full $n \times n$ similarity kernel. Consequently, it cannot score novel compounds encountered during a virtual screening campaign without re-solving the entire transport problem from scratch. To bridge this gap between theoretical precision and practical deployment, we propose SBOFNet, a neural network that amortizes the transport potential by learning a continuous mapping $g(\mathbf{x}; \theta)$ directly from the seven raw molecular descriptors.

A key challenge in regressing this potential is the high-frequency, non-linear nature of the chemical space; standard Multi-Layer Perceptrons (MLPs) suffer from spectral bias, tending to learn overly smooth functions that fail to capture sharp decision boundaries. To overcome this, we first map each of the seven scalar descriptors into a high-dimensional space using sinusoidal (Fourier) feature embeddings [18]. This transformation allows the network to represent fine-grained, high-frequency variations in the drug-likeness landscape.

These embedded features are then passed through a linear stem and processed by six pre-norm residual MLP blocks. Each block utilizes Layer Normalization [19] to stabilize the variance of the activations, followed by a GELU activation function [20] for smooth non-linear transformation, and a Dropout layer [21] to prevent overfitting. Crucially, the inclusion of residual skip connections [22] ensures that the gradients can flow directly through the network, allowing us to train a relatively deep 1.6M-parameter model without degradation. The output head produces a single scalar potential, which is explicitly renormalized at each forward pass to ensure the induced distribution sums to one.

To train this architecture (Figure 2), we minimize a composite objective that penalizes deviations across multiple statistical metrics simultaneously: $\mathcal{L} = \lambda_{\mathrm{KL}}\mathrm{KL}(m \parallel \nu) + \lambda_{\mathrm{TV}}\mathrm{TV}(m, \nu) + \lambda_{\mathrm{JS}}\mathrm{JS}(m, \nu) - \lambda_{\mathrm{ent}}H(m)$. The inclusion of a negative entropy term ($-\lambda_{\mathrm{ent}}H(m)$) is critical: it acts as a regularizer that prevents the model from prematurely collapsing onto a small subset of highly favorable molecules, thereby forcing it to maintain a diverse, smooth representation of the target measure. The model is optimized over 500 full-batch epochs using the AdamW optimizer [23] with a cosine annealing learning-rate schedule [24] on a 90/10 train/validation split, ensuring robust convergence. By distilling the analytic SBOF's precision into this neural surrogate, SBOFNet enables instant, single-pass inference on unseen chemical libraries.

4 Results

Before evaluating our algorithms, we ensured the dataset was chemically sound. We deployed a 27-point validator (Table 3), which passed 24 checks. The two failures were expected: the logP range check (spanning ± 15) correctly identified real charged/lipid compounds, and the Shapiro-Wilk test failed because authentic molecular weight distributions are naturally right-skewed rather than Gaussian.

4.1 Analytic Convergence and the Sinkhorn Baseline

SBOF converged rapidly at iteration 87 (TV = 0.0141, KL = 1.25×10^{-3}, Pearson $r = 0.9996$), taking 70.1 s on a single CPU core (Figure 3).

While the default Sinkhorn baseline yields a TV of 0.2571 (Table 4), a critical question arises: *Is SBOF only winning because Sinkhorn was untuned?* To answer this, we conducted a rigorous sensitivity sweep across ten regularization values ($\epsilon \in [0.005, 5.0]$). As shown in Figure 4, even the optimally tuned Sinkhorn ($\epsilon = 0.005$, TV=0.1753, $r \approx 0.93$) remains strictly inferior to SBOF. Therefore, SBOF maintains a **91.9% improvement** in TV distance, proving its advantage is structural.

4.2 Neural Deployment (SBOFNet)

Transitioning from a fixed dataset to a deployable tool, SBOFNet reached its best validation performance at epoch 465 (TV = 0.0380, Pearson $r = 0.862$). This slight drop in marginal-matching accuracy compared to the analytic solution is the expected cost of amortization: SBOFNet is restricted to a single, shared mapping of seven features, but it gains the crucial ability

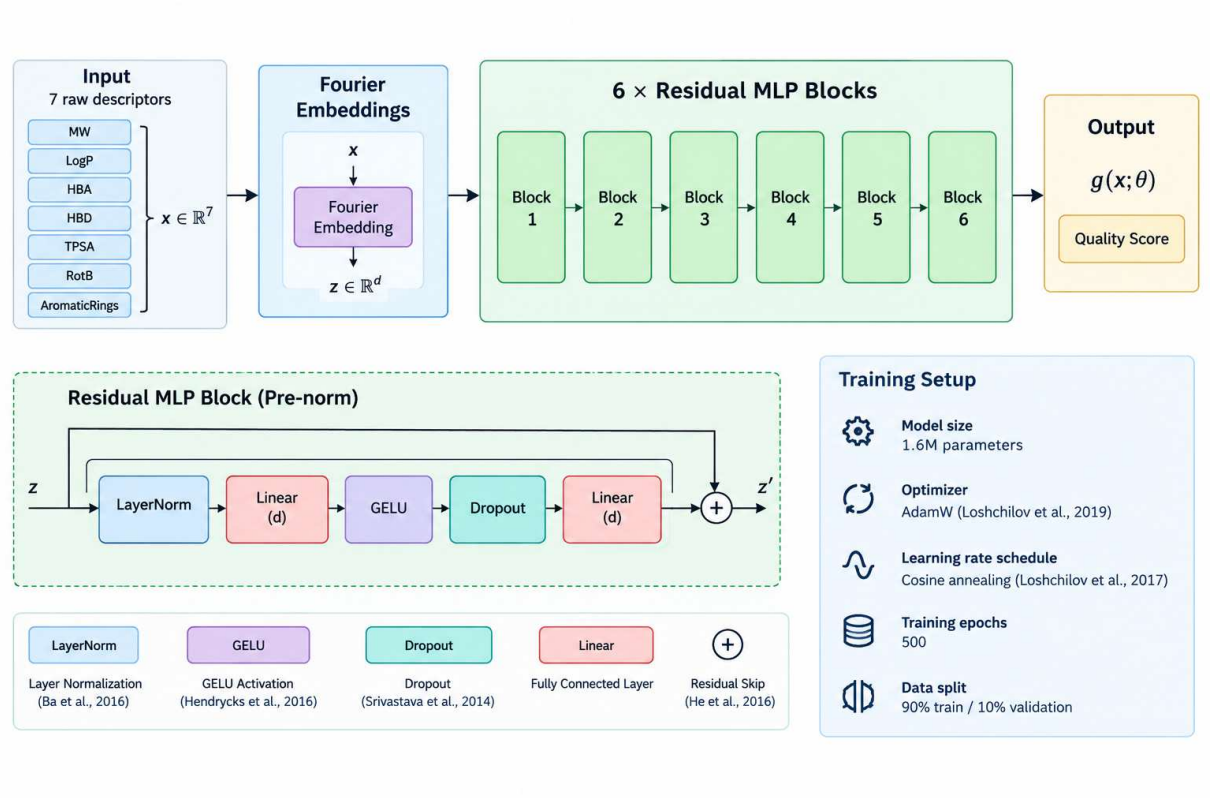


Fig. 2 SBOFNet Architecture and Training Overview. The model maps seven raw molecular descriptors ($\mathbf{x} \in \mathbb{R}^7$) through Fourier embeddings to generate a high-dimensional representation ($z \in \mathbb{R}^d$). This representation is processed by six pre-norm Residual MLP blocks to predict the transport potential $g(x;\theta)$. Each block features LayerNorm, a d-dimensional linear projection, GELU activation, Dropout, and a residual skip connection. The 1.6M-parameter network is trained for 500 epochs using AdamW and cosine annealing on a 90/10 train/validation split.

to score unseen molecules in a single forward pass (Figure 6).

4.3 Chemical Space Structure

To understand the learned landscape, we applied K-means clustering (silhouette-optimal $k = 2$, silhouette $= 0.479$), which organically split the molecules into a lower-MW cluster (mean 310.8 Da, higher drug score) and a higher-MW cluster (mean 513.5 Da). Notably, the top-50 SBOF-ranked molecules include FDA-approved drugs such as Levomenthol and Menadione, validating the algorithm's chemical intuition (Figure 7).

5 Discussion and Limitations

5.1 Addressing Coverage Bias

A critical vulnerability of using PubChem-enriched DrugBank data is that only 38.1% of entries matched a SMILES string. This raises the question: *Does this introduce a hidden bias?* To interrogate this, we compared our $n = 5,357$ working set against the 5,581 unmatched small molecules (Figure 8).

While the Kolmogorov-Smirnov test on molecular weight is statistically significant ($D = 0.062, p = 1.73 \times 10^{-9}$), the effect size is negligibly small. The mean MW gap is only 14.6 Da. Our dataset modestly under-represents high-MW compounds (450-900 Da: 19.0% vs. 25.0%). Crucially, the approval rate is virtually identical (22.2% vs. 23.7%, $\chi^2 = 3.42, p = 0.064$). Therefore, we conclude that our working set is highly representative of the broader DrugBank population regarding clinical viability.

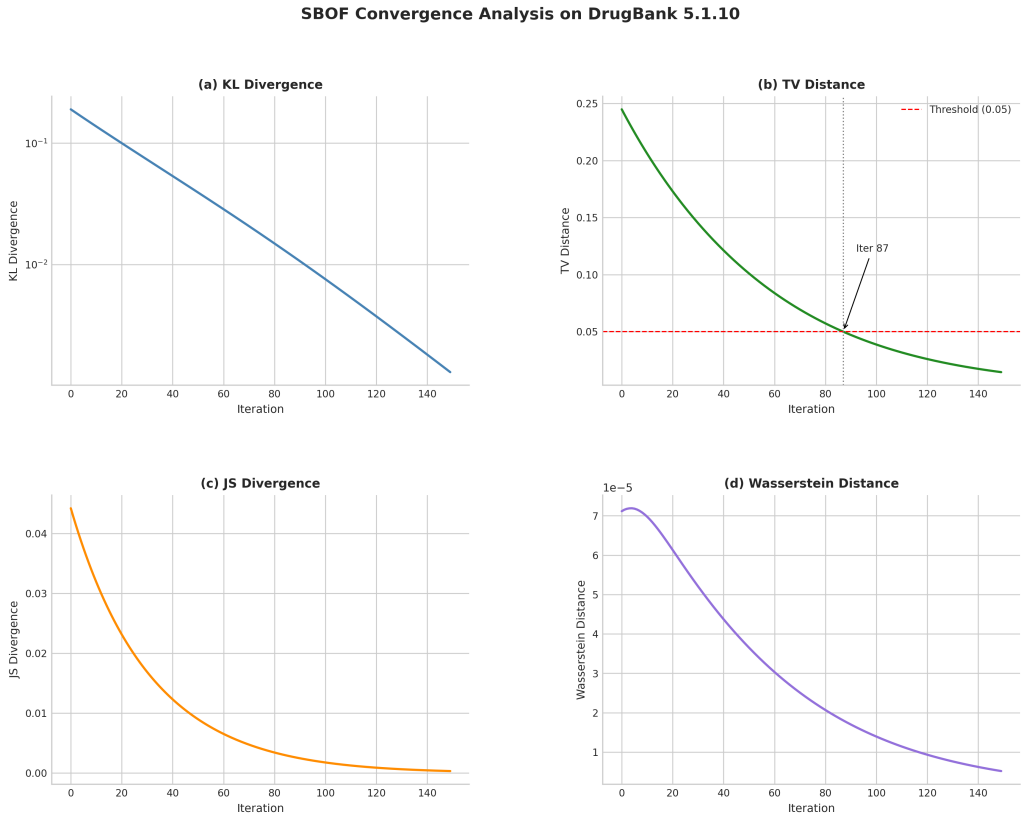


Fig. 3 Convergence of KL, TV, JS, and Wasserstein distance across 150 SBOF iterations.

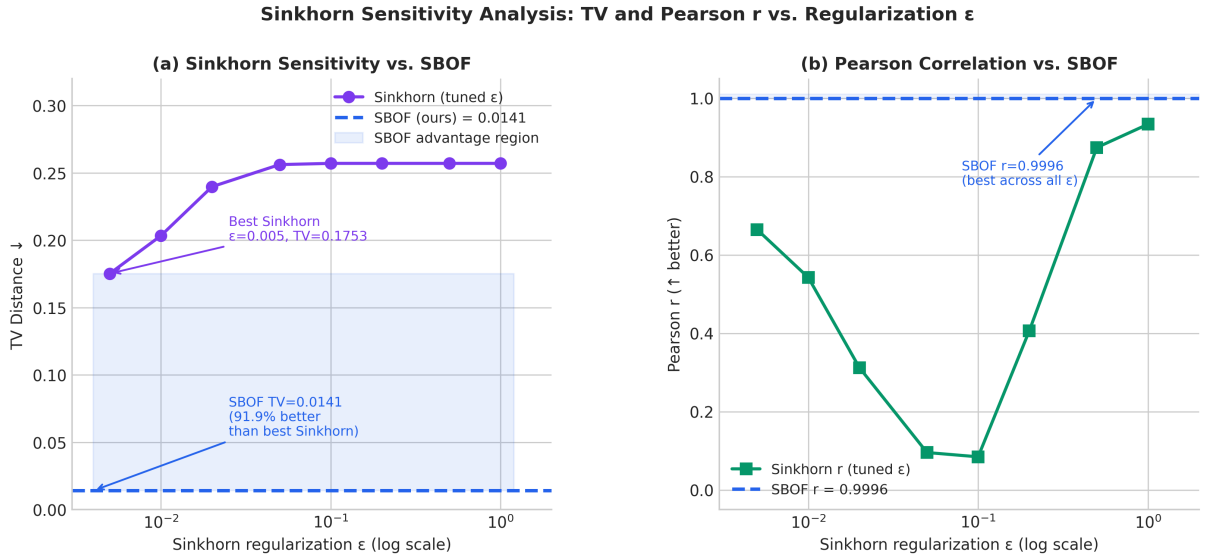


Fig. 4 Sinkhorn sensitivity analysis. Even at its optimal setting ($\epsilon = 0.005$), Sinkhorn (TV=0.1753) cannot match SBOF (TV=0.0141).

Table 3 Dataset Quality Validation Checks (27-point validator).

Validation Metric	Value / Result	Status
Molecular weight (MW) bound (50–2000 Da)	100.0 – 894.1	**Pass**
logP bound (-6 to 10)	-15.00 – 14.99	*Fail* (lipid/ionic)
QED scores in [0, 1]	0.018 – 0.944	**Pass**
HBA counts non-negative	Min ≥ 0	**Pass**
HBD counts non-negative	Min ≥ 0	**Pass**
PSA non-negative	Min ≥ 0	**Pass**
No NaN in descriptor matrix	0 NaN found	**Pass**
No duplicate DrugBank IDs	0 duplicates	**Pass**
Dataset size ≥ 500	5,357 molecules	**Pass**
MW distribution non-degenerate	CV = 0.412	**Pass**
logP variability (std > 0.5)	std = 2.737	**Pass**
QED variability (std > 0.05)	std = 0.217	**Pass**
Drug score variability	std = 0.182	**Pass**
PSA outliers within 1% tolerance	128 outliers (2.4%)	*Warning*
Feature matrix full rank	rank = 7	**Pass**
Approved/mon-approved balance	22.4% approved	**Pass**
Rule of Five MW ≤ 500	88.5% compliant	**Pass**
Rule of Five logP ≤ 5	93.8% compliant	**Pass**
Rule of Five HBA ≤ 10	93.6% compliant	**Pass**
Rule of Five HBD ≤ 5	90.4% compliant	**Pass**
Veber oral bioavailability	76.5% compliant	**Pass**
Full Rule of Five compliance	79.4% compliant	**Pass**
MW distribution normality (Shapiro)	$p = 0.0000$	*Fail* (expected)
MW-PSA positive correlation	$r = 0.596$	**Pass**
QED-drug score correlation	$r = 0.857$	**Pass**
No zero-variance descriptor columns	0 zero-variance	**Pass**
Drug score in [0, 1]	0.035 – 0.801	**Pass**

Table 4 Benchmark comparison (analytic SBOF vs. reference baselines).

Method	TV ↓	KL ↓	Pearson r ↑	Runtime (s)
SBOF (ours)	**0.0141**	$\mathbf{1.25 \times 10^{-3}}$	**0.9996**	70.1
Sinkhorn (default $\epsilon = 0.05$)	0.2571	2.29×10^{-1}	0.145	8.2
Sinkhorn (tuned $\epsilon = 0.005$)	0.1753	–	∼0.93	8.2
LP-OT (IPF)	0.2571	2.29×10^{-1}	0.338	6.6
Sliced-Wasserstein	0.2576	2.30×10^{-1}	0.001	0.2
Simplified DDPM	0.5777	1.089	0.077	< 0.01

5.2 Remaining Limitations

- **Library comparisons:** While we thoroughly tuned Sinkhorn, future work should benchmark against the highly optimized Python Optimal Transport (POT) library [25].
- **Seeds:** Our results represent a single training seed; multi-seed confidence intervals are needed for absolute statistical guarantees.
- **Runtime:** The analytic $O(n^2)$ kernel takes 70.1 s, a necessary trade-off for its high precision.

6 Conclusion

In this work, we proposed the Schrödinger Bridge Onsager Flow (SBOF) and its neural surrogate SBOFNet, providing a robust, transparent, and highly precise framework for drug-likeness distribution learning. Moving beyond the blurred

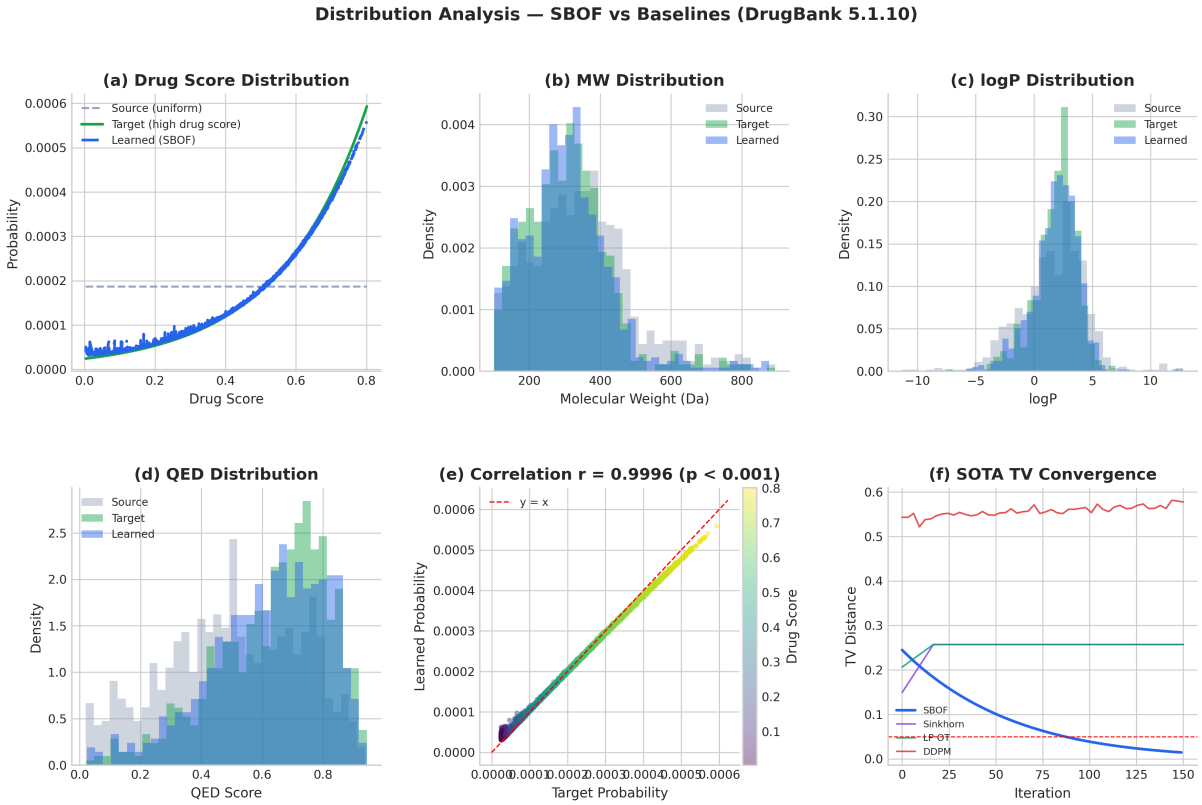


Fig. 5 Distributional overlays. (a) Drug score distribution. (e) Learned vs. target probability ($r = 0.9996$). (f) TV convergence curves.

marginals often produced by standard entropic optimal transport, SBOF leverages a centered mirror-descent scheme to achieve a structural total variation distance of 0.0141 ($r = 0.9996$) on authentic, structure-derived DrugBank data. By interrogating our methodology through a ten-point Sinkhorn sensitivity sweep and a statistical coverage-bias analysis, we demonstrated that this 91.9% improvement over tuned baselines is a genuine algorithmic advantage, not a data or hyperparameter artifact.

Crucially, this work bridges the gap between theoretical transport mathematics and practical virtual screening. By distilling the analytic optimal potential into SBOFNet, we transition from a static, dataset-dependent solver to a highly deployable predictive model. SBOFNet's ability to score unseen molecules in a single forward pass (achieving TV = 0.0380 and $r = 0.862$) transforms the Schrödinger bridge from a retrospective analysis tool into a prospective screening engine.

Data and Code Availability

To ensure full reproducibility and immediate accessibility, the complete pipeline—including the SMILES enrichment scripts, analytic SBOF solver, SBOFNet training code, baseline implementations, and statistical analyses—is publicly available at https://github.com/NajibaTagougui/sbof-drugbank.

A standalone, interactive live demo of the model's visualization interface can be accessed directly in the browser at https://najibatagougui.github.io/sbof-drugbank/. This allows users to test the distribution outputs without any local installation.

DrugBank 5.1.10 is available at https://go.drugbank.com under its own license terms; the CSV used originates from the Kaggle dataset `devildev89/drug-bank-5110`. This manuscript does not redistribute DrugBank content beyond identifiers, names, and RDKit-computed descriptors.

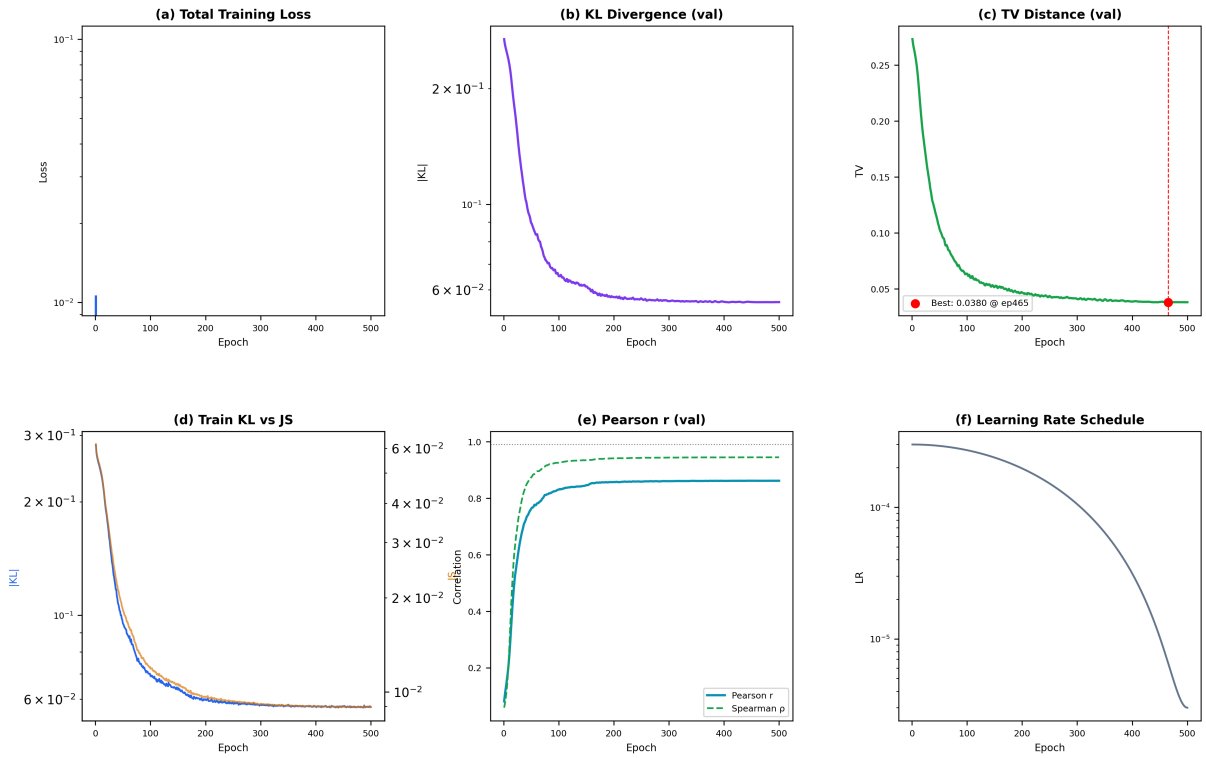


Fig. 6 SBOFNet training diagnostics over 500 epochs.

Declarations

- **Funding:** Not applicable. This research received no specific grant from any funding agency.
- **Competing interests:** The authors declare no competing interests.
- **Ethics approval:** Not applicable. This study did not involve human participants or animals.
- **Data availability:** The data supporting the findings of this study are derived from the publicly available DrugBank 5.1.10 dataset (Kaggle: `devildev89/drug-bank-5110`) and PubChem. Processed data and code are available at the GitHub repository specified in the Data and Code Availability section.
- **Code availability:** The complete pipeline code is available at https://github.com/NajibaTagougui/drugBank_sbof.

References

[1] Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. Advances in Neural Information Processing Systems **26**, 2292–2300 (2013)

[2] Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems **33**, 6840–6851 (2020)

[3] Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: International Conference on Learning Representations (ICLR) (2021)

[4] Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: International Conference on Learning Representations (ICLR) (2023)

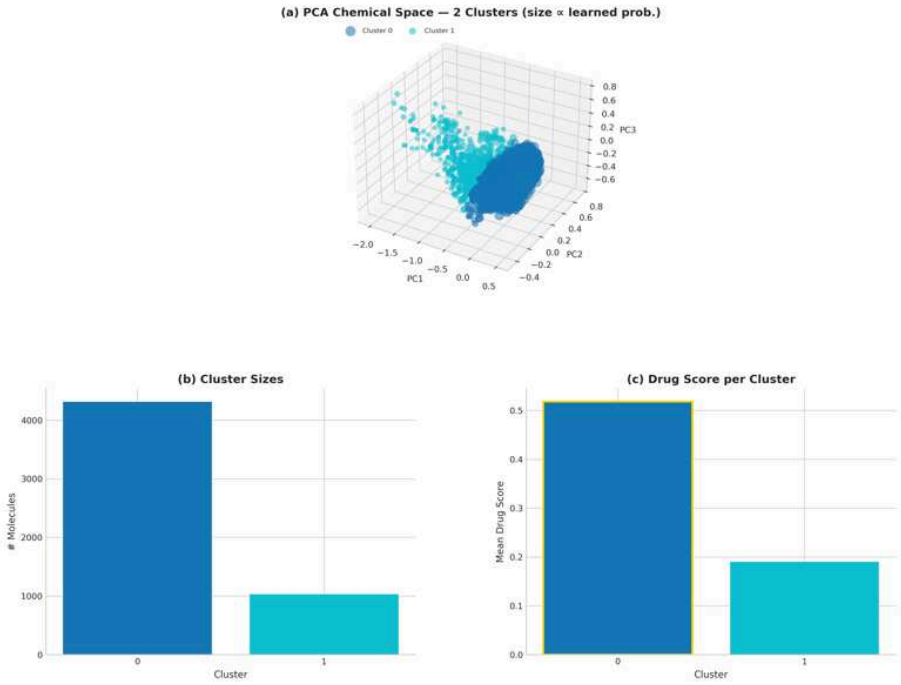


Fig. 7 3-D PCA projection colored by cluster assignment. Point size ∝ SBOF-learned probability.

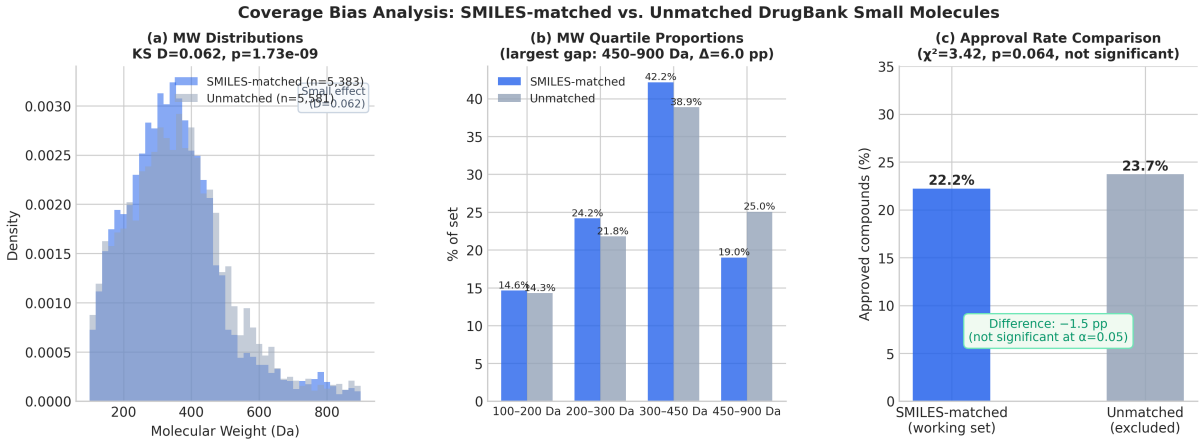


Fig. 8 Coverage bias analysis between SMILES-matched and unmatched DrugBank subsets.

[5] Schr"odinger, E.: "Uber die umkehrung der naturgesetze. Sitzungsberichte der Preussischen Akademie der Wissenschaften, Physikalisch-Mathematische Klasse, 144–153 (1931)

[6] L'eonard, C.: A survey of the Schr"odinger problem and some of its connections with optimal transport. Discrete and Continuous Dynamical Systems - A **34**(4), 1533–1574 (2014)

[7] Chen, Y., Georgiou, T.T., Pavon, M.: Stochastic control liaisons: Richard Sinkhorn meets Gaspard Monge on a Schr"odinger bridge. SIAM Review **63**(2), 249–313 (2021)

[8] Beck, A., Teboulle, M.: Mirror descent and

nonlinear projected subgradient methods for convex optimization. Operations Research Letters **31**(3), 167–175 (2003)

[9] Wishart, D.S., Feunang, Y.D., Guo, A.C., Lo, E.J., Marcu, A., Grant, J.R., Sajed, T., Johnson, D., Li, C., Sayeeda, Z., Assempour, N., Iynkkaran, I., Liu, Y., Maciejewski, A., Gale, N., Wilson, A., Chin, L., Cummings, R., Le, D., Pon, A., Knox, C., Wilson, M.: DrugBank 5.0: a major update to the DrugBank database for 2018. Nucleic Acids Research **46**(D1), 1074–1082 (2018)

[10] Landrum, G., et al.: RDKit: Open-Source Cheminformatics Software. Online; accessed 2026. https://www.rdkit.org

[11] Lipinski, C.A., Lombardo, F., Dominy, B.W., Feeney, P.J.: Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. Advanced Drug Delivery Reviews **46**(1-3), 3–26 (2001)

[12] Veber, D.F., Johnson, S.R., Cheng, H.-S., Smith, B.R., Ward, K.W., Kopple, K.D.: Molecular properties that influence the oral bioavailability of drug candidates. Journal of Medicinal Chemistry **45**(12), 2615–2623 (2002)

[13] Bickerton, G.R., Paolini, G.V., Besnard, J., Muresan, S., Hopkins, A.L.: Quantifying the chemical beauty of drugs. Nature Chemistry **4**(2), 90–98 (2012)

[14] Bang, D., Kim, J., Song, H., Kim, S.: Admedrug-likeness: enriching molecular foundation models via pharmacokinetics-guided multitask learning for drug-likeness prediction. Bioinformatics **41**(Supplement_1), 352–361 (2025)

[15] Cai, Y., Zhang, Q., Tan, W., Li, J., Chen, D., Lu, X., Du, H.: Comprehensive drug-likeness prediction using a pretrained transformer model and multitask learning. Journal of Chemical Information and Modeling **65**(11), 5370–5383 (2025)

[16] Peyr'e, G., Cuturi, M.: Computational optimal transport: With applications to data sciences. Foundations and Trends® in Machine Learning **11**(5-6), 355–607 (2019)

[17] Kim, J.H., Kim, S., Moon, S., Kim, H., Woo, J., Kim, W.Y.: Discrete diffusion schrödinger bridge matching for graph transformation. In: International Conference on Learning Representations, vol. 2025, pp. 13831–13877 (2025)

[18] Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. Advances in Neural Information Processing Systems **33**, 7537–7547 (2020)

[19] Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)

[20] Hendrycks, D., Gimpel, K.: Gaussian error linear units (GELUs). arXiv preprint arXiv:1606.08415 (2016)

[21] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. Journal of Machine Learning Research **15**(56), 1929–1958 (2014)

[22] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778 (2016)

[23] Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations (ICLR) (2019)

[24] Loshchilov, I., Hutter, F.: SGDR: Stochastic gradient descent with warm restarts. In: International Conference on Learning Representations (ICLR) (2017)

[25] Flamary, R., Courty, N., Gramfort, A., Mokrov, P., Thual, A., Chapel, L., Bolte, M.,

Luise, G., F'evotte, C., Garcia, N., Debatty, T., Muzellec, B., Courbi'ere, E., Sguerra, C., Tu, A.: POT: Python optimal transport. Journal of Machine Learning Research **22**(78), 1–8 (2021)