Aligning Transformers with Continuous Feedback via Energy Rank Alignment

Shriram Chennakesavalu, Frank Hu,
Sebastian Ibarraran, and Grant M. Rotskoff

NeurIPS 2025

Stanford University

Chemistry Language Models

Foundation Models for Molecular Generation

- Molecular Transformer (2019):

Input: reactants-reagents (atom-wise tokenization)

Br c 1 c c c 2 ...c(c1)c1cc3c4ccccc4c4ccccc4c3cc1n2-c1ccc2c(c1)c1ccccc1n2-c1ccccc1.CCO.
Cc1ccccc1.OB(O)c1ccc2ccc3cccnc3c2n1.c1ccc([PH](c2ccccc2)(c2ccccc2)[Pd]([PH](c2ccccc2)
(c2ccccc2)c2ccccc2)([PH](c2ccccc2)(c2ccccc2)c2ccccc2)[PH](c2ccccc2)(c2ccccc2)c2ccccc2)cc1

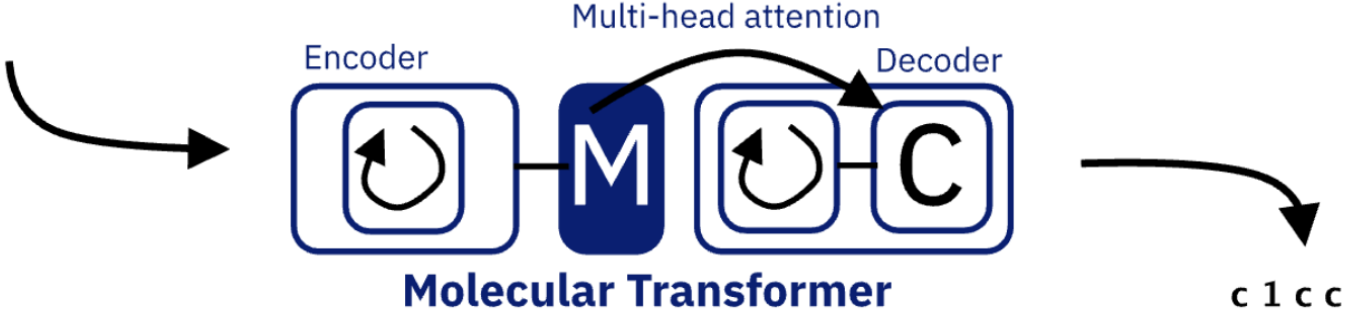

Molecular Transformer

c 1 c c c (– n
...2c3ccccc3c3cc(–n4c5ccc(–
c6ccc7ccc8cccnc8c7n6)cc5c5cc6c7ccccc7c7ccccc7c6cc54)ccc32)cc1

Target: most likely products

Schwaller et al., *ACS Cent Sci*, **2019**

Energy Rank Alignment (ERA)

- Simple gradient-based objective
- Explicitly incorporates a reward/energy function
- Avoids greedy policies
- Controllable regularization
- State-of-the-art performance on molecular generation benchmarks

Theoretical Framework

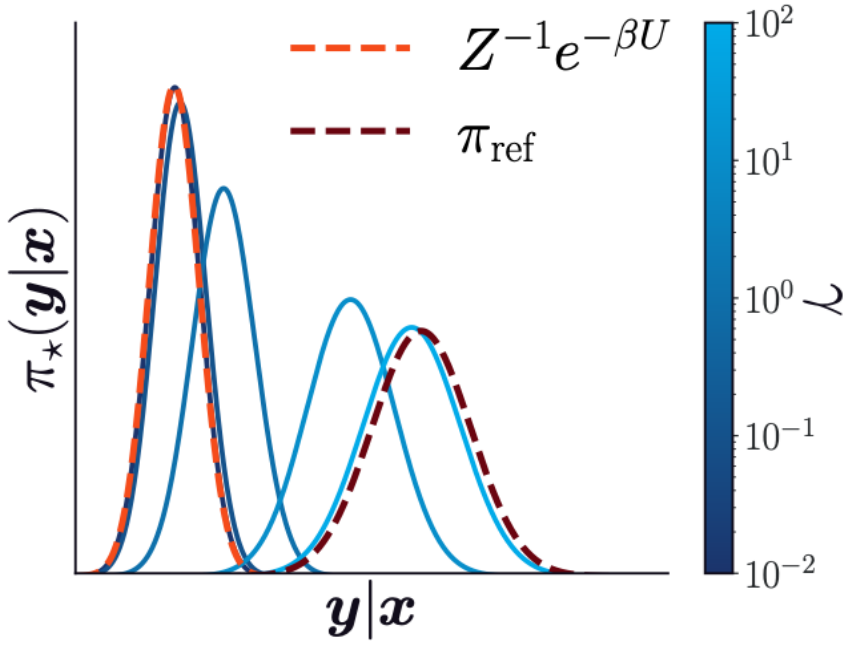
$Z^{-1}e^{-\beta U}$
π_{ref}
$\pi_\star(\boldsymbol{y}|\boldsymbol{x})$
$\boldsymbol{y}|\boldsymbol{x}$
10^2
10^1
10^0
10^{-1}
10^{-2}
λ

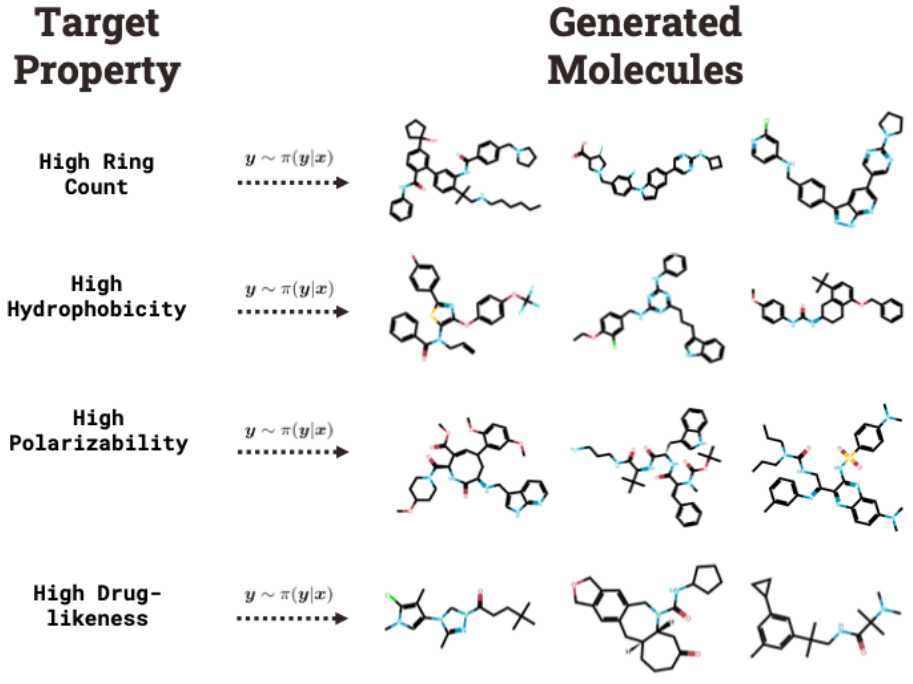
Target Property
Generated Molecules
High Ring Count
High Hydrophobicity
High Polarizability
High Drug-likeness
$y \sim \pi(y|x)$

Experiments

Molecular Generation

- Foundation model: transformer trained on 2.4M ChEMBL SMILES
- Focused on easily computed *in silico* oracles
  - RDKit: QED, hydrophobicity, ring count, Tanimoto similarity
  - Docking: random forest predictors of docking scores against two kinase targets – GSK3β and JNK3
- Tested both unprompted (BOS token only) and prompted (conditioned on a SMILES) alignment

Single-Property Alignment

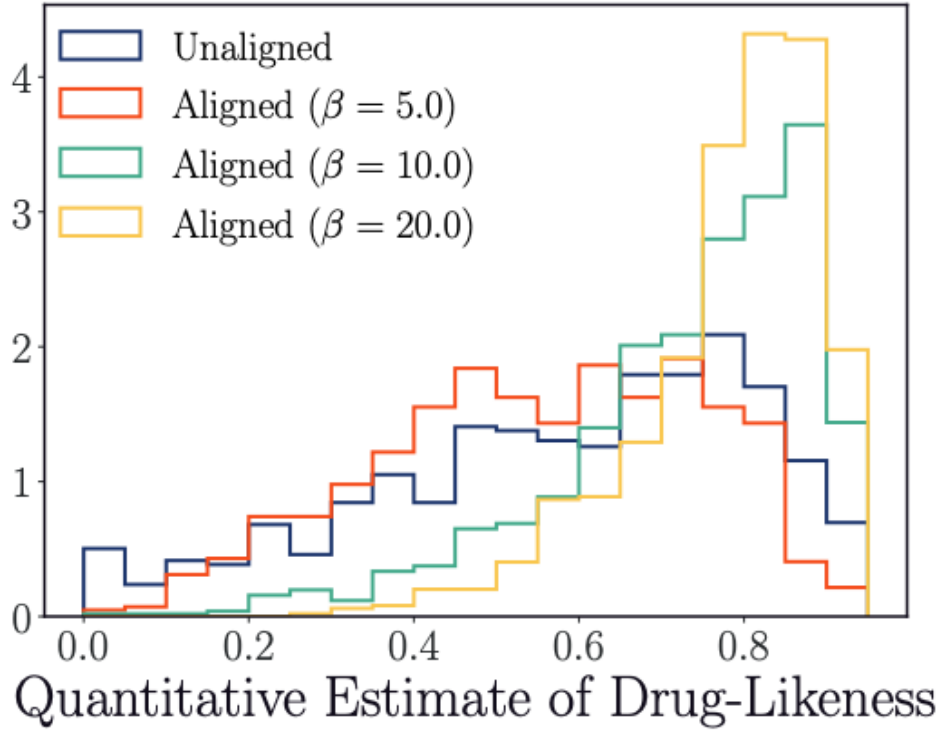

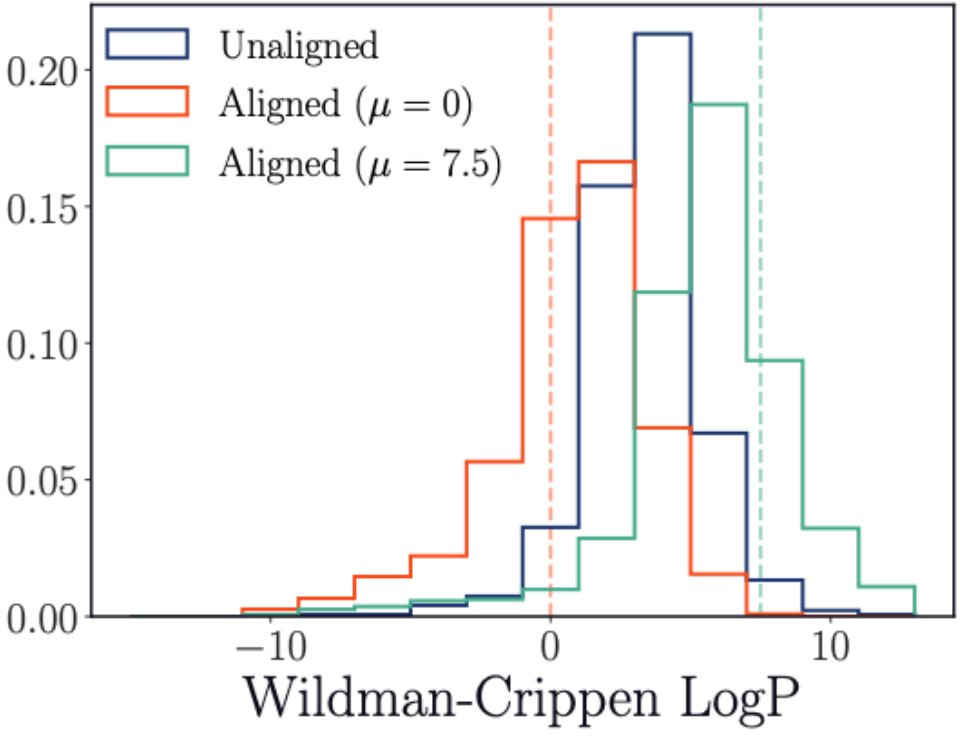

Multi-Property Alignment

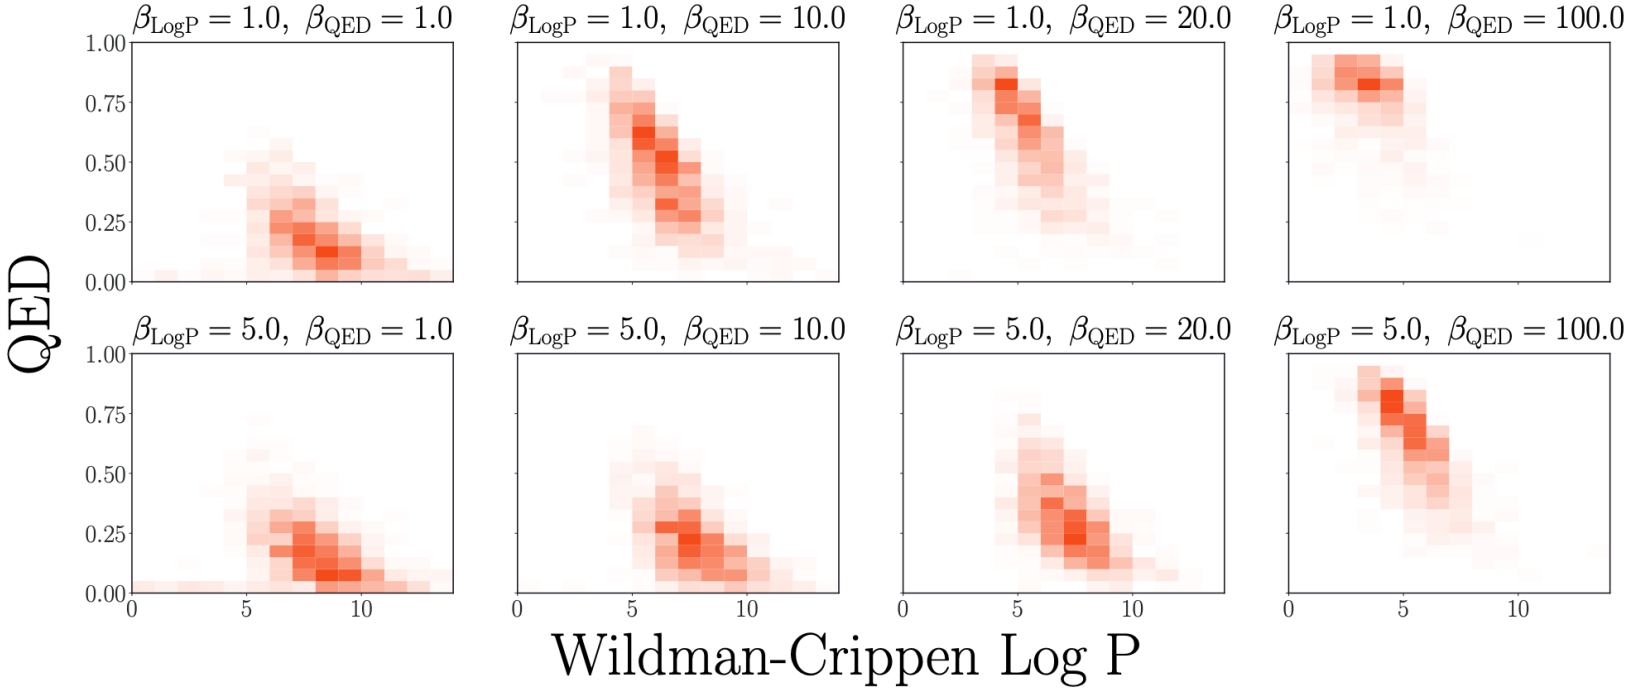

Prompted Multi-Property Alignment

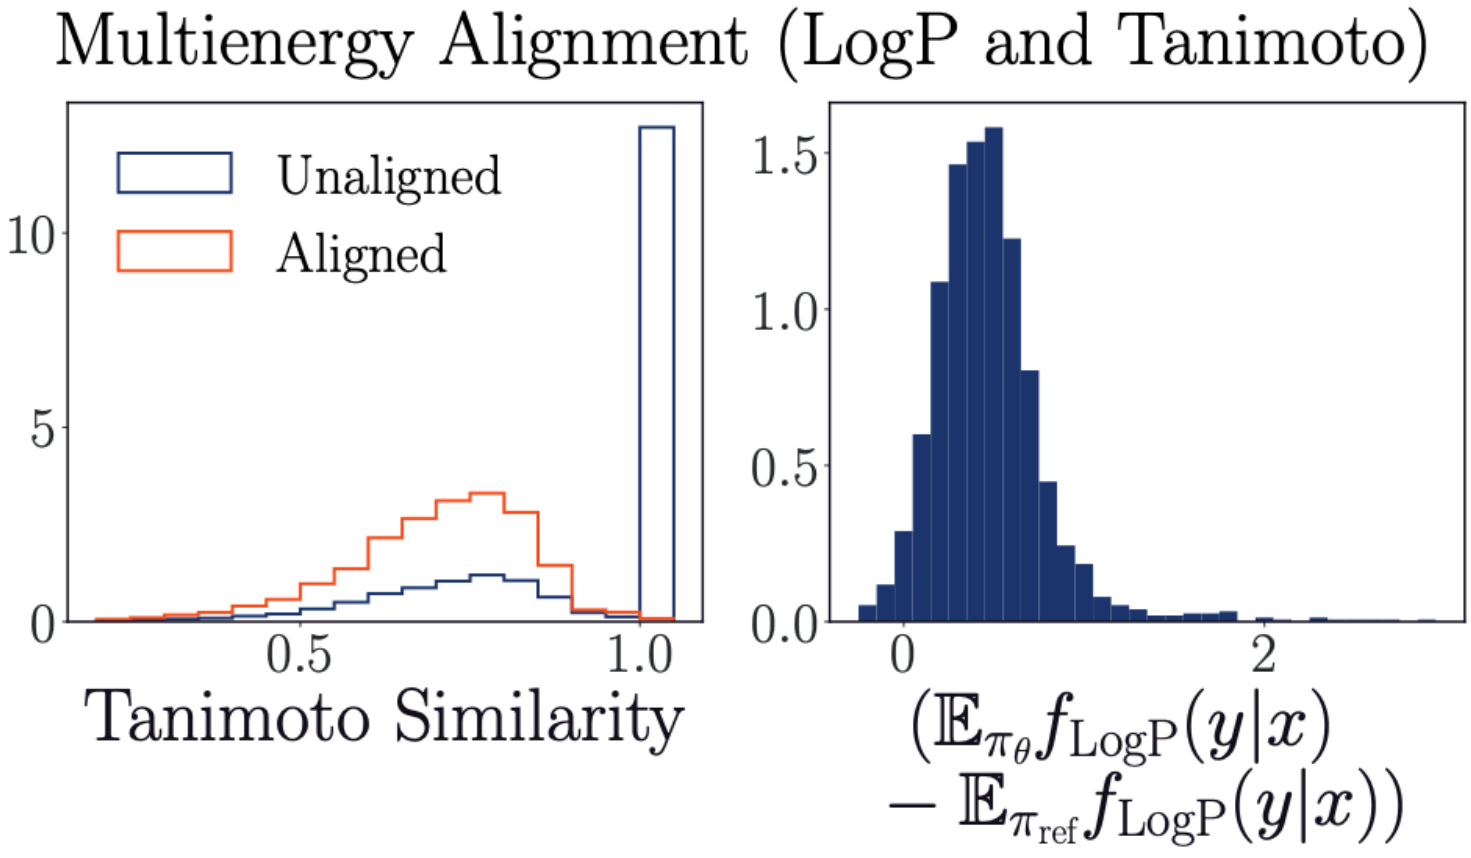

Docking Benchmarks

- Generate more diverse molecules (lower IntDiv) with higher bioactivities

	GSK3β top-100		JNK3 top-100	
	mean score	IntDiv	mean score	IntDiv
ERA	0.996 ± 0.000	$\mathbf{0.219 \pm 0.002}$	$\mathbf{0.987 \pm 0.001}$	$\mathbf{0.264 \pm 0.005}$
MolRL-MGPT	$\mathbf{1.000 \pm 0.000}$	0.362 ± 0.015	0.961 ± 0.010	0.372 ± 0.025
GFlowNet	0.649 ± 0.072	0.715 ± 0.104	0.437 ± 0.219	0.716 ± 0.145
GraphGA	0.919 ± 0.016	0.365 ± 0.024	0.875 ± 0.025	0.380 ± 0.015
JT-VAE	0.235 ± 0.083	0.770 ± 0.067	0.159 ± 0.040	0.781 ± 0.127
REINVENT	0.965 ± 0.011	0.308 ± 0.035	0.942 ± 0.019	0.368 ± 0.021

Acknowledgements

Yinuo Ren
Ph.D. (Institute for Computational and Mathmatical Engineering)

Andy Mitchell
Ph.D. (Chemistry)

Grant Rotskoff
Assistant Professor of Chemistry

Jérémie Klinger
Postdoc

Abigail Park
Ph.D. (Chemistry)

Sherry Li
Ph.D. (Chemistry)

Frank Hu
Ph.D. (Chemistry, Markland Lab)

Emmit Pert
Ph.D. (Chemistry)

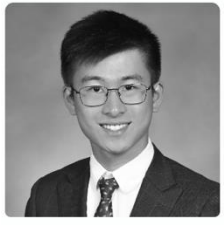

Josh Liu
Ph.D. (Biophysics)

Nick Juntunen
Ph.D. (Chemistry)

Steven Dunne
Ph.D. (Biophysics)

Wenhao Gao
Postdoc