

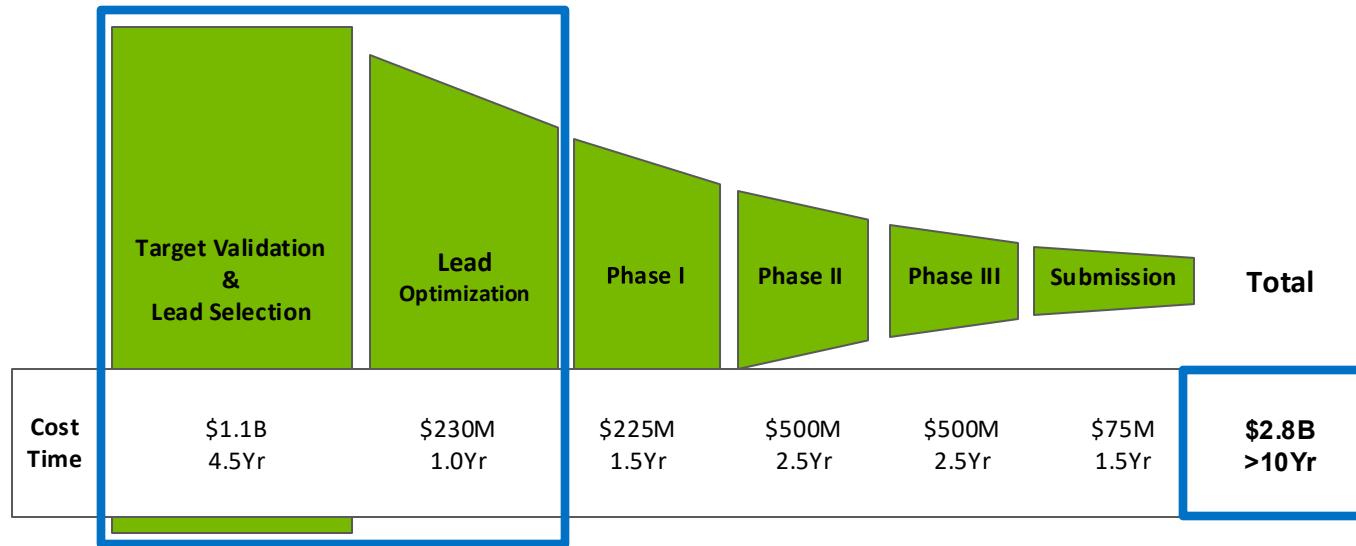


# Scientific Discovery: From the Lab Bench to the GPU

Michelle L. Gill, PhD; Tech Lead and R&D Manager, NVIDIA

PyDataNYC | 3rd November, 2023

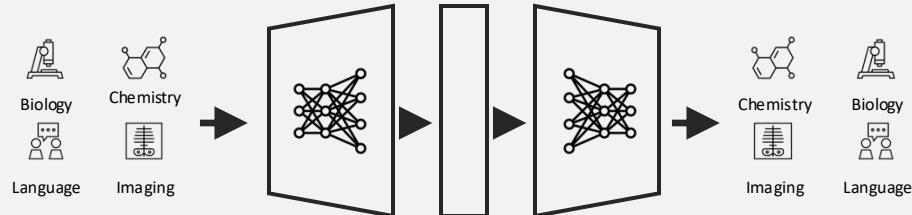
## Motivation: Drug Development is a Long and Expensive Process



**\$2.8B and >10 Years to Bring a Drug to Market**

# Language Models are Revolutionizing Discovery

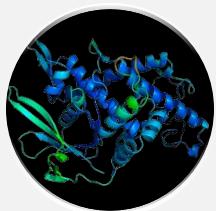
- Information from biomedical literature
- Prediction of chemical reactions
- Biomolecular property prediction
- Structure prediction and docking



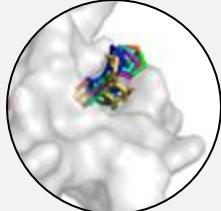
BIOMEDICAL NLP  
Learn all of PubMed



GENERATIVE CHEMISTRY  
Novel Drug Candidates



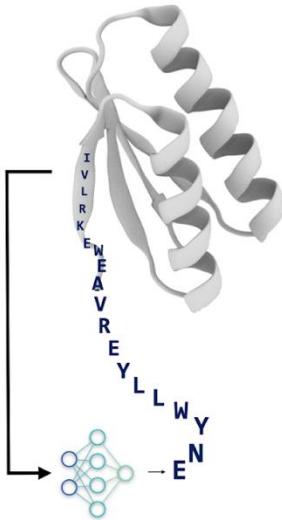
PROTEIN STRUCTURE  
Predict 3D Structures



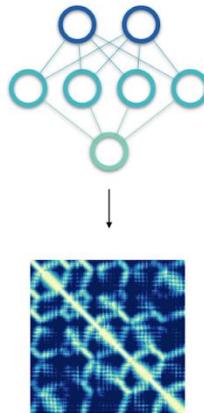
VIRTUAL SCREENING  
Docking and Pose Prediction

# From Sequence to 3D and Back Again

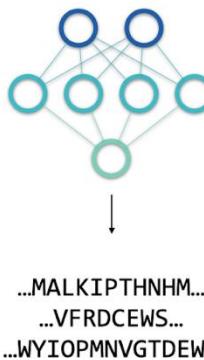
## 1 Fixed-backbone design



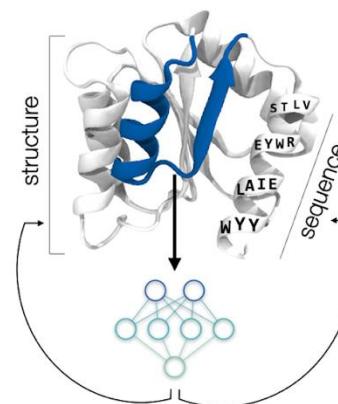
## 2 Structure Generation



## 3 Sequence generation



## 4 Sequence and structure design



# Outline

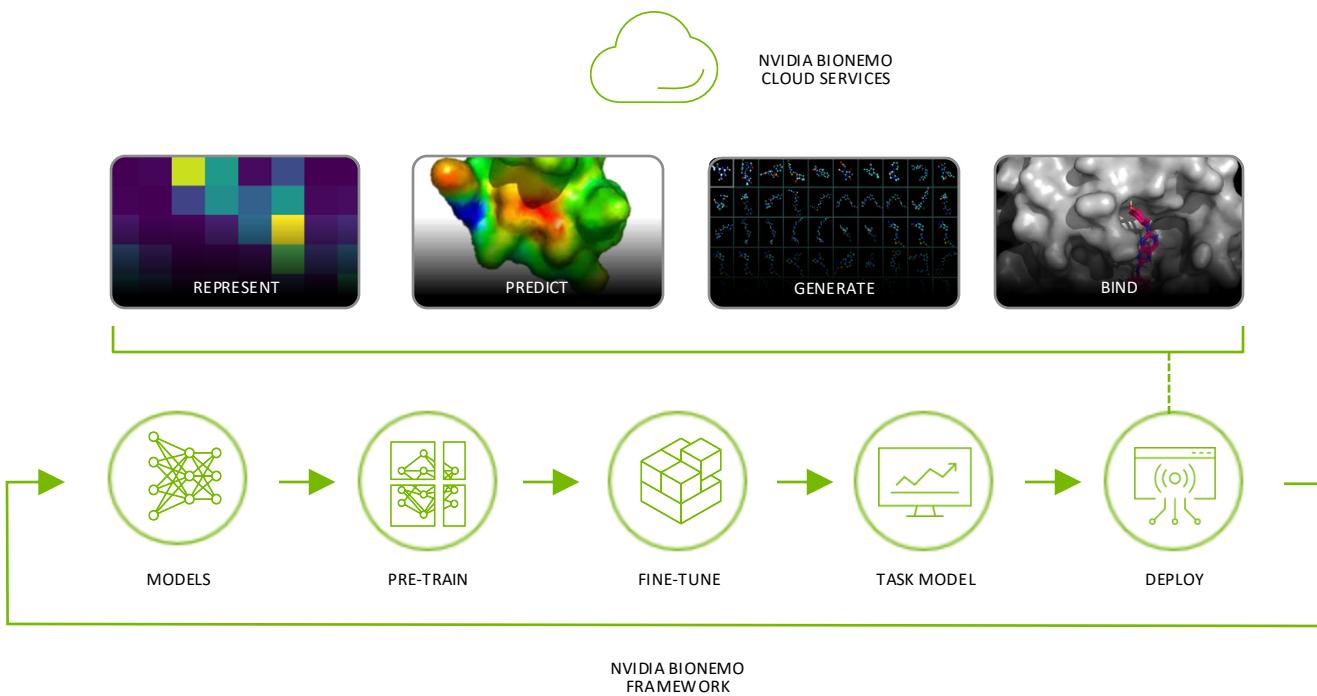
- Overview of BioNeMo: Inference Service and Training Framework
- MolMIM: Development of a Small Molecule Foundation Model for Generation
- Career Progression and Lessons from the Field

# BioNeMo Overview: Inference Service and Framework



# NVIDIA BioNeMo

AI Tools, Frameworks, and Applications for Drug Discovery



# Multiple Interfaces to a BioNeMo Model in the Inference Service

## Interactive UI and Jupyter Workflows

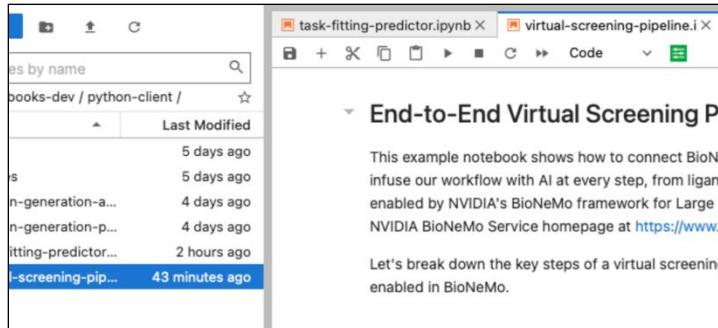


The screenshot shows the BioNeMo playground interface. At the top, there are tabs for Protein Generation, Protein Embedding, Molecule Generation, Molecule Embedding, Protein Folding (which is selected), and Docking. Below the tabs, there is a section to "Choose a model to generate sequence output. If you have a UniProt ID, input it below or you can start with one of our provided example use cases." A dropdown for "Model" is set to "AlphaFold2". A "Select an Example UniProt ID" dropdown has "Look Up ID" selected, and a "Examples" dropdown is open, showing "UniProt ID: 014763". A "Protein Sequence" text area displays a sequence of 440 characters. The main area features a 3D visualization of a protein structure with a color scale from 50 to 100 labeled "Prediction Score (0.0-1.0)".

## API and Python Client



```
1 import requests
2
3 ngc_token="<<NGC TOKEN>>"
4 headers = { "Authorization" : f"Bearer {ngc_token}" }
5
6 try:
7     response =
requests.post("https://api.stg.bionemo.ngc.nvidia.com/v1/protein-
sequence/protpt2/generate",
8                 headers=headers,
9                 json={
10                 "max_length":150,
```

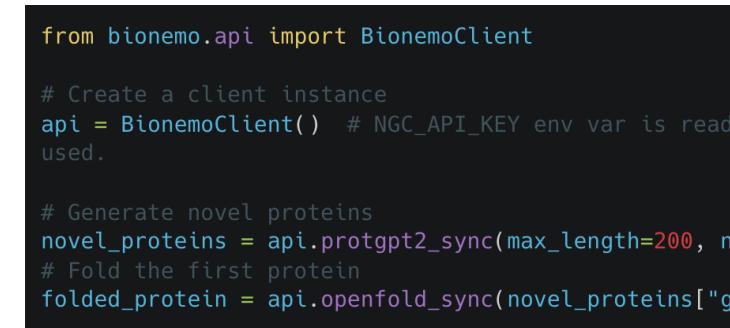


The screenshot shows a Jupyter notebook titled "task-fitting-predictor.ipynb". The notebook is part of a "books-dev / python-client" repository. The code cell contains the following text:

End-to-End Virtual Screening Pi

This example notebook shows how to connect BioNeMo to infuse our workflow with AI at every step, from ligand screening to protein folding, enabled by NVIDIA's BioNeMo framework for Large Language Models. The BioNeMo Service homepage is at <https://www.ngc.nvidia.com/catalog/publishers/nvidia/bionemo>.

Let's break down the key steps of a virtual screening pipeline enabled in BioNeMo.



```
from bionemo.api import BionemoClient

# Create a client instance
api = BionemoClient() # NGC_API_KEY env var is read
used.

# Generate novel proteins
novel_proteins = api.protpt2_sync(max_length=200, num
# Fold the first protein
folded_protein = api.openfold_sync(novel_proteins[0],
```



# Welcome to BioNemo!

Get started with a model below. Explore documentation for more information.

Secondary Action

Primary Action

## Get Started with BioNemo



### Protein Generation

These models generate proteins with a sequence distribution that mirrors the distribution of proteins on which the model was trained.

ProtGPT-2



### Protein Embedding

These models generate protein embeddings. They take an amino acid sequence and returns a learned representation.

ESM-1nv

ESM-2



### Molecule Generation

Given a seed molecule, these models can generate similar molecules

MoFlow

MegaMolBART



### Molecule Embedding

These models generate embeddings for a given molecule.

MegaMolBART



### Protein Folding

These models predict the 3D structure of a protein from only the sequence of amino acids.

ESMFold

OpenFold

AlphaFold-2



### Docking

These models take a molecule structure and a protein structure and predict the docked pose.

DiffDock



### Generate an API Key

Authenticate your identity while making queries to NeMo LLM via the REST API.

Generate API Key



### Documentation

Learn more about using NeMo LLM and dive deep with tutorials, how-to guides and examples.

Documentation



## Lab

[Protein Generation](#) [Protein Embedding](#) [Molecule Generation](#) [Molecule Embedding](#) [Protein Folding](#) [Docking](#)

Choose a model to generate sequence output. If you have a PDB ID, input it below or you can start with one of our provided example use cases.

Model ⓘ

 ▼

Output ⓘ

Enter a UniProt ID ⓘ

 Enter UniProt ID...[Look Up](#)

Or

Select an Example UniProt ID ⓘ

 ▼

Protein Sequence ⓘ

Look up a UniProt ID, choose an Example from the provided list or enter your own here...

Perform MD Refinement ⓘ

Brief description of what this does



MSA ⓘ

Upload an MSA or choose no MSA. One will be auto-generated if you take no action.

### View Code

[OpenAPI](#) [Curl](#) [Python](#)

```
1 curl -X POST "https://api.bionemo.ngc.nvidia.com/v1/protein-structure/openfold/predict" \
2   -H "Content-Type: application/json" \
3   -H "Authorization: Bearer $YOUR_NGC_API_TOKEN" \
4   -d '{
5     "sequence": "MSFSGKYQLSQENFEAFMKAIGLPEELIQGKDIKGVSEIVQNGKHFKFTITAGSKVIQNEFTVGECECLEMTGEKVKTVQLEGDNKLVTTFKNIKVTELNGDIITNTMLGDIVFKRISKRI"
6   }'
```

[Learn how to integrate the API into your application](#) [here](#)[Click here](#) to generate a new API key. [Copy Code](#)[Done](#)[Clear](#)[Generate](#)[Give Feedback](#)[View Code](#)[Download](#)



## Playground

Protein Generation Protein Embedding **Molecule Generation** Molecule Embedding Protein Folding Docking

Choose a model to generate molecules. If you have a Chemical CID, input it below or you can start with one of our provided example use cases.

[Learn More](#)

Model ①

MoFlow

Select an Example CID ①

Look Up ID

Examples

Dicloxacillin

SMILES ①

73 of 510 chars

```
Cc1onc(-c2c(Cl)cccc2Cl)c1C(=O)N[C@@H]1C(=O)N2C[C@H]1SC(C)C[C@H]2C(=O)O
```

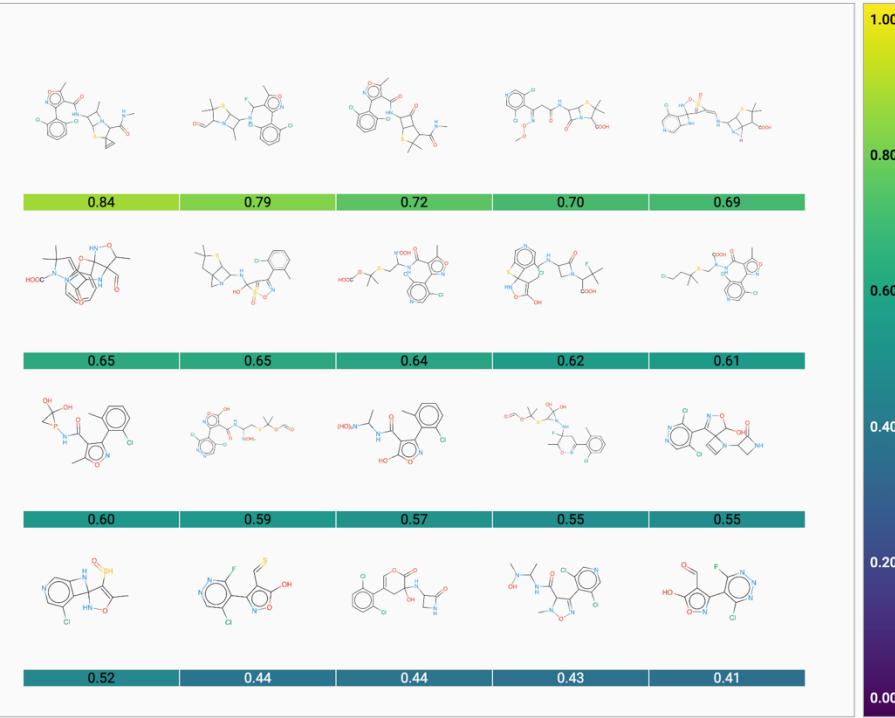
Number of Molecules ①

20

Sample Temperature ①

0.20  0.35

Output ①

[Clear](#)[Generate](#)[Give Feedback](#)[View Code](#)[Download](#)[Collapse](#)

## Playground

Protein Generation Protein Embedding Molecule Generation Molecule Embedding Protein Folding **Docking**

Choose a model to generate docking poses. Provide a molecule and a target protein file.

[Learn More](#)

Model ①

DiffDock

Molecule ①

Ensitrelvir\_analog

Target Protein ①

SARS\_CoV\_2\_MPro

Generated Poses ①

20

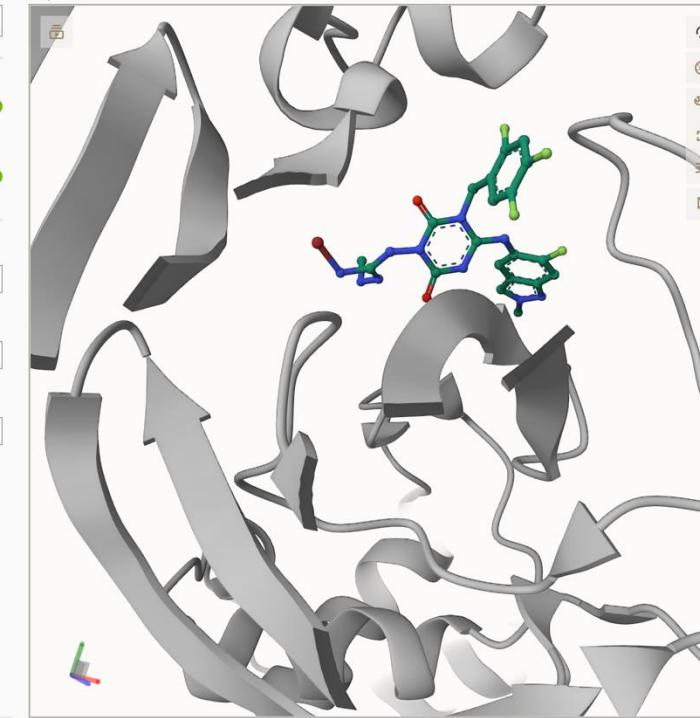
Diffusion Steps ①

18

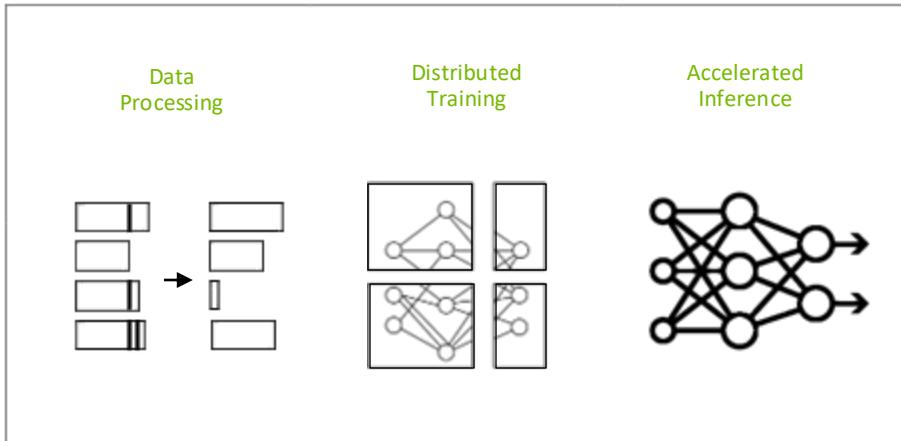
Diffusion Time Divisions ①

20

Output ①

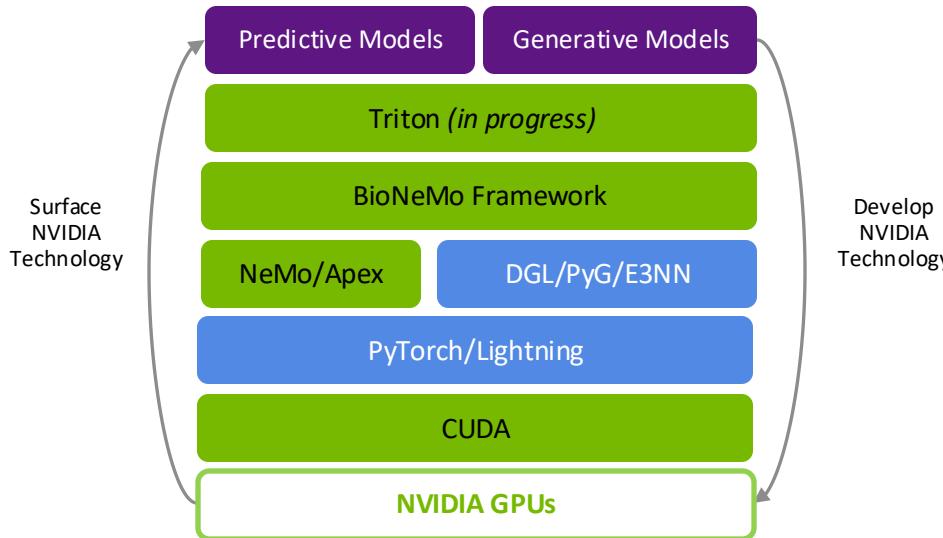
[Center Pose](#) [Reset View](#) [View All Poses](#) < >Rank: 1 Score:  
-0.567Rank: 2 Score:  
-0.769Rank: 3 Score:  
-0.789Rank: 4 Score:  
-1.155Rank: 5 Score:  
-1.254Rank: 6 Score:  
-1.621Rank: 7 Score:  
-1.655Rank: 8 Score:  
-2.039Rank: 9 Score:  
-2.144Rank: 10 Score:  
-2.184Rank: 11 Score:  
-2.372Rank: 12 Score:  
-2.576Rank: 13 Score:  
0.100[Clear](#)[Generate](#)[Give Feedback](#)[View Code](#)[Download](#)

# BioNeMo Framework Overview



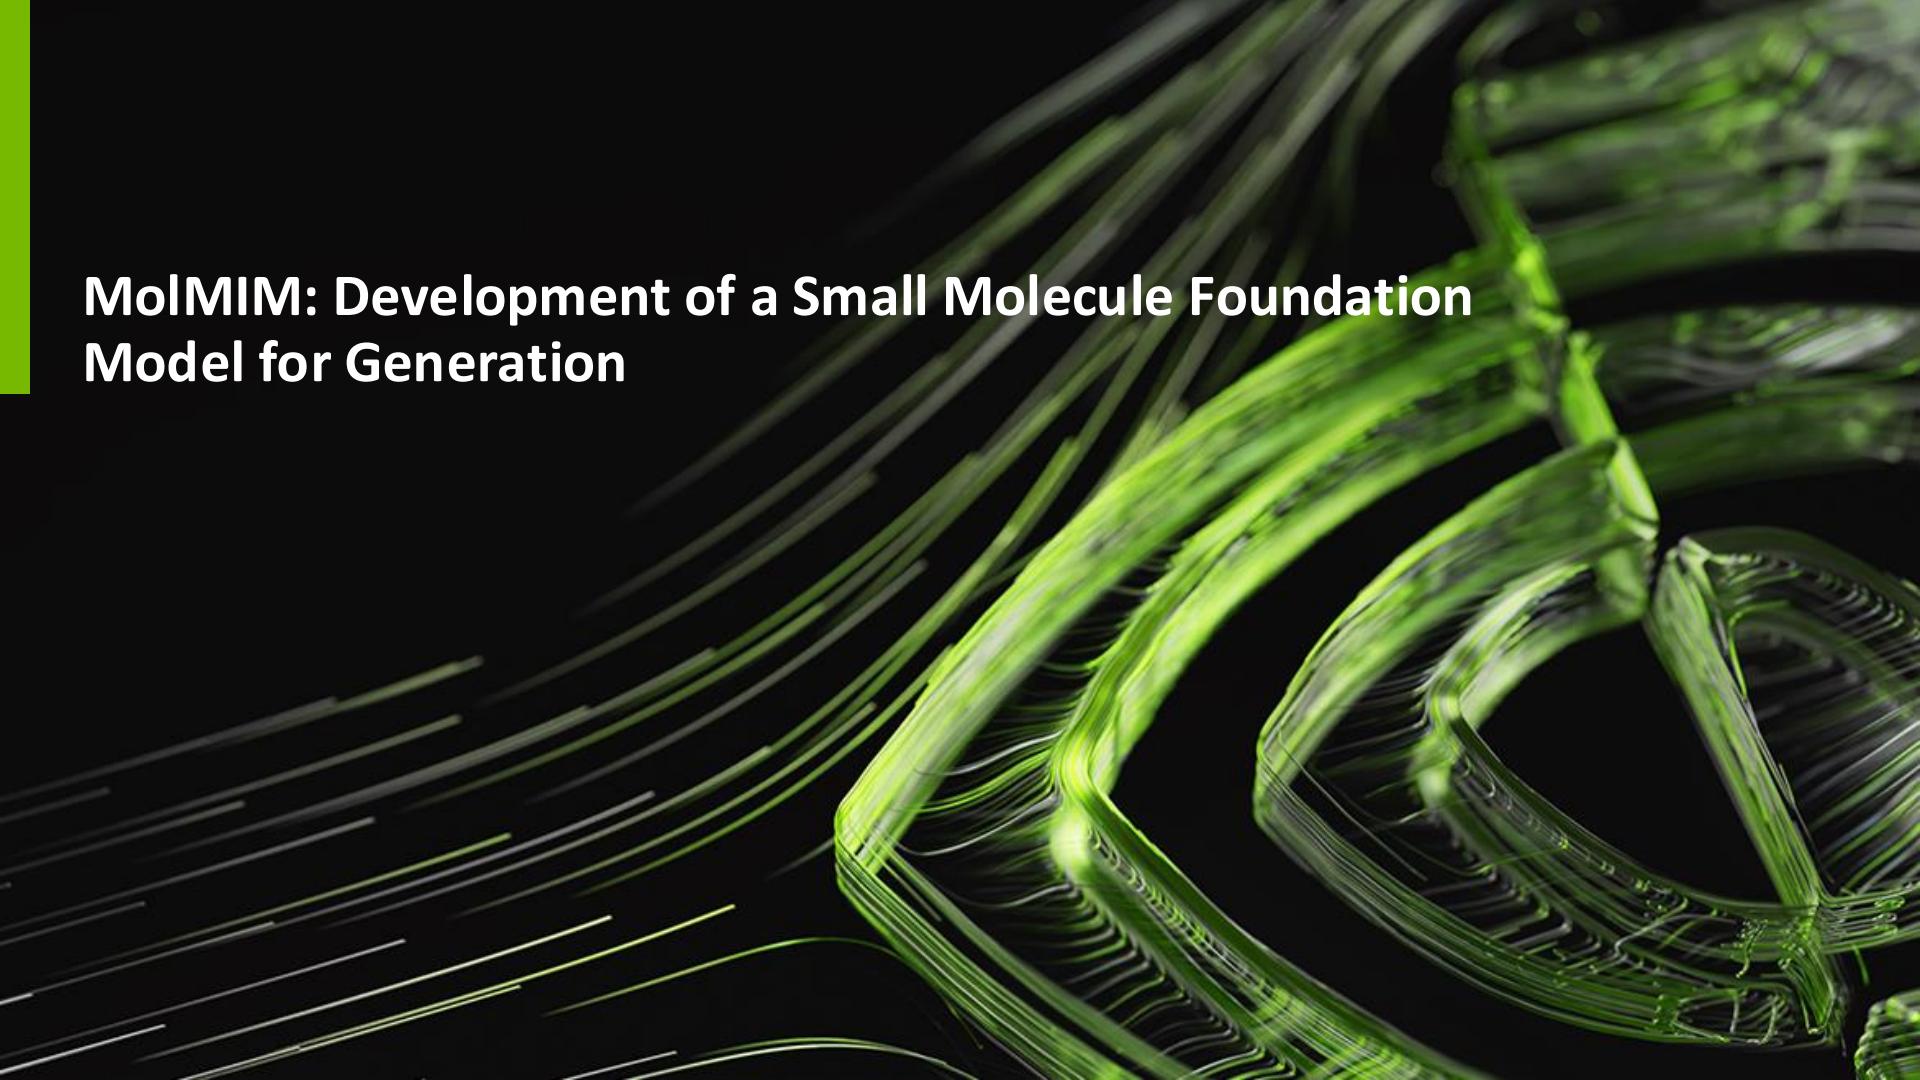
- Includes dataset processing, training, fine tuning, and example downstream tasks
- Support for multi-GPU and multi-node training
- Data parallelism, and three types of model parallelism
- Currently three LLM models for cheminformatics and protein applications – more models and model types coming soon

# BioNeMo Framework Technology Stack

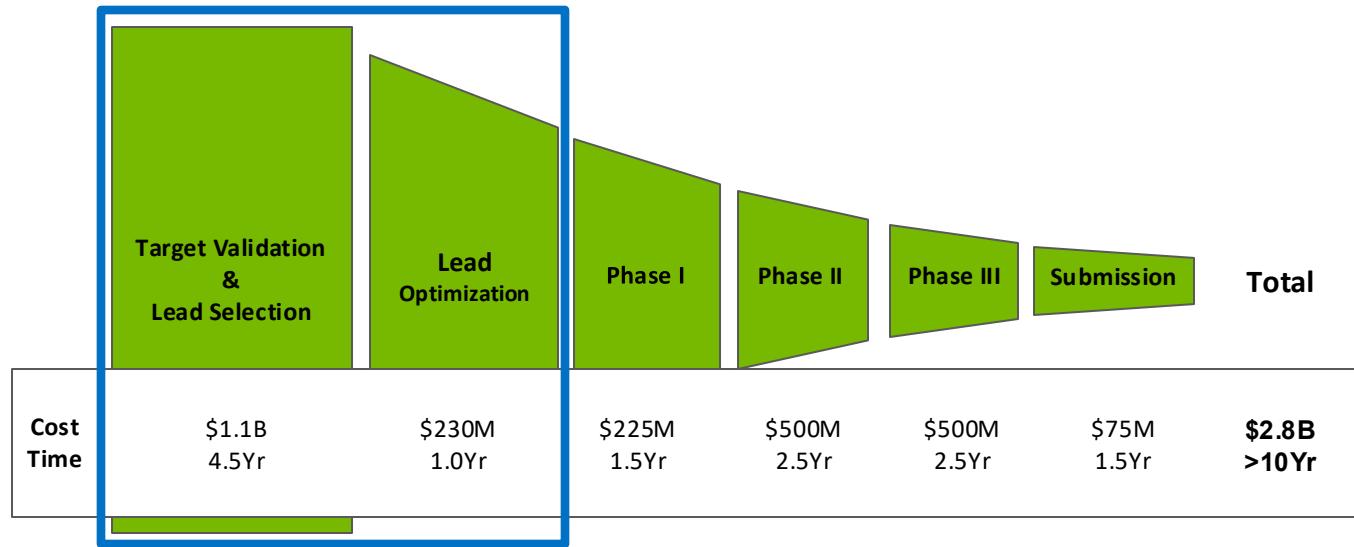


- Based on NVIDIA NeMo, which is a library for development and training of LLMs
- Automated deployment with Triton is in progress
- Surface and develop new software and hardware technology

# MolMIM: Development of a Small Molecule Foundation Model for Generation

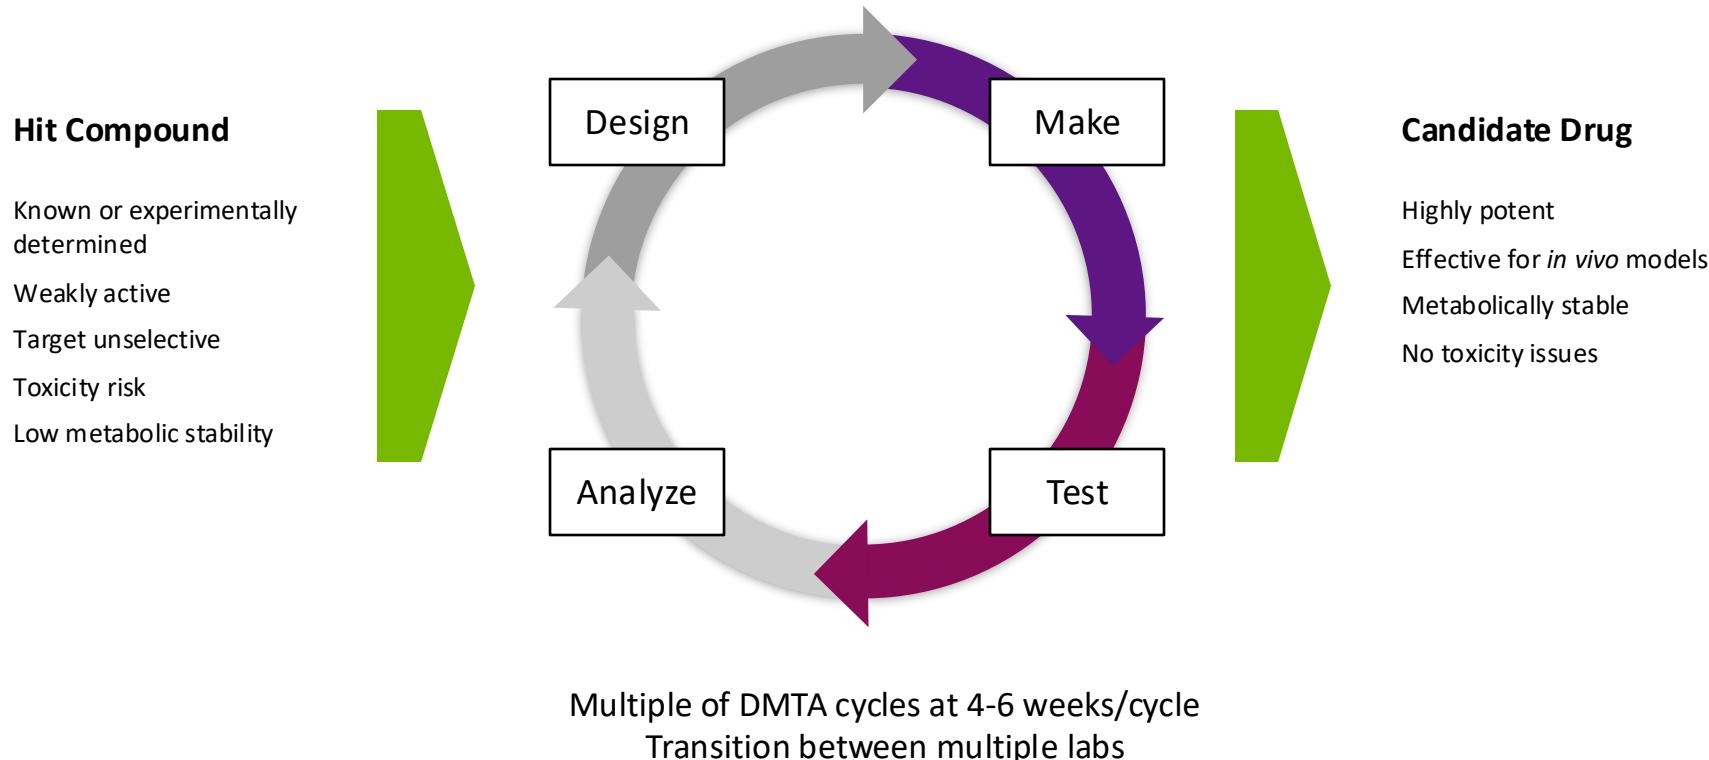


## Motivation: Drug Development is a Long and Expensive Process



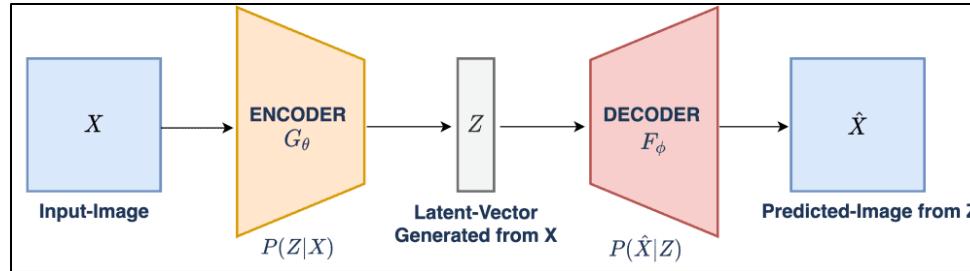
**\$2.8B and >10 Years to Bring a Drug to Market**

# Lead Discovery: Three Years for Design-Make-Test-Analyze Cycle

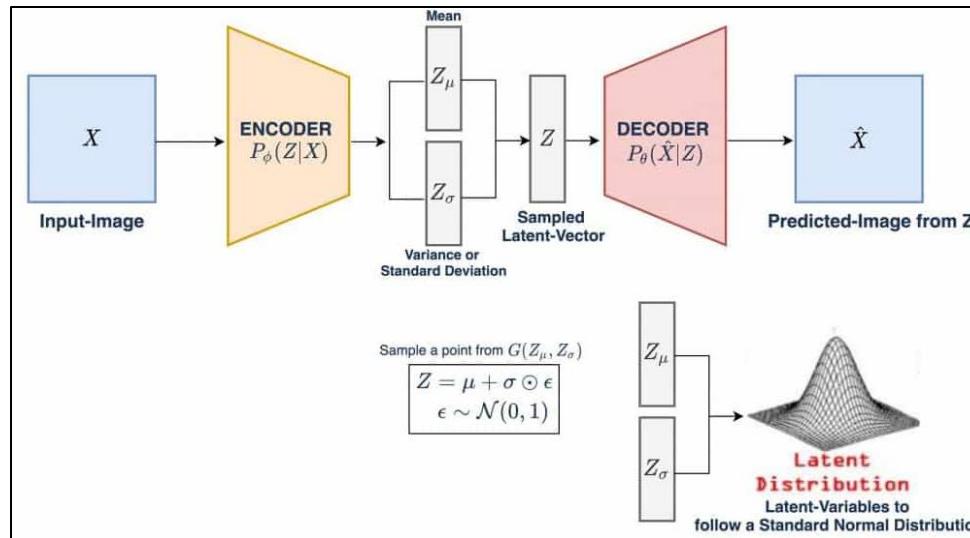


# Autoencoder Models in a Nutshell

Autoencoder



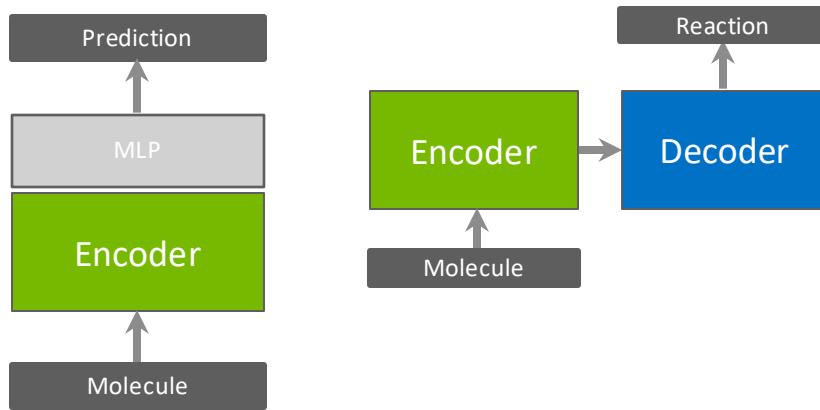
Variational  
Autoencoder (VAE)



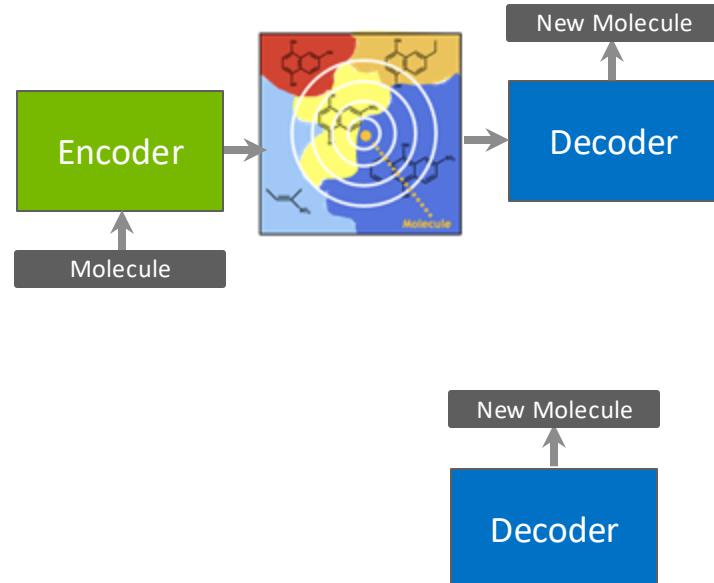
Also works  
with  
sequences --  
seq2seq  
models

# Cheminformatics Foundation Model Objectives

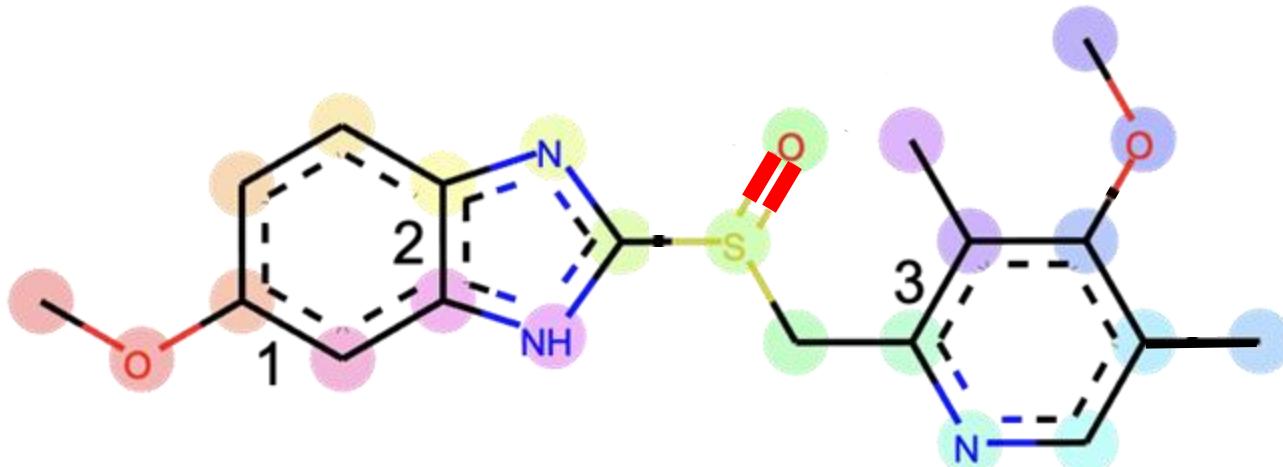
## Representation and Translation



## Generation



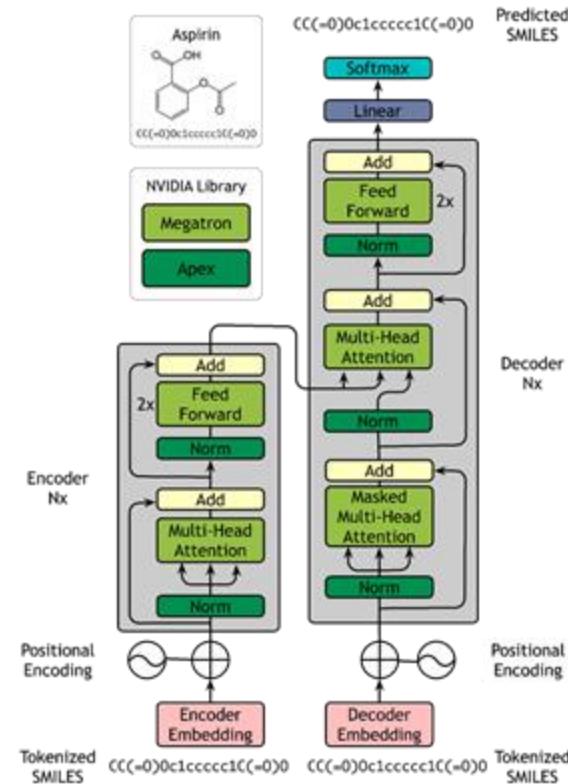
## SMILES: a Natural Language Representation of Small Molecules



COc1ccc2n c(S(=O)Cc3ncc(C)c(OC)c3C)[nH]c2c

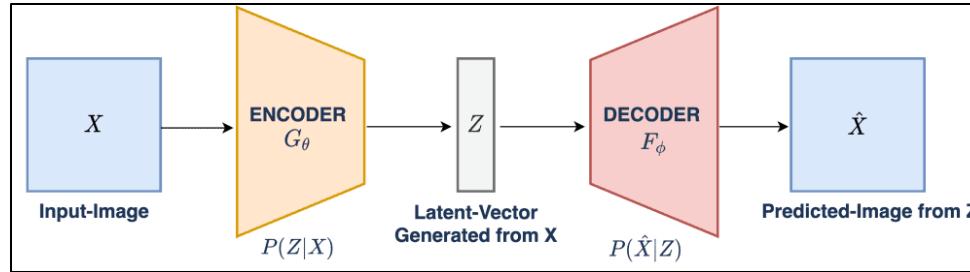
# MegaMolBART Molecule Representations

- MegaMolBART is a sequence-to-sequence developed in collaboration with AstraZeneca
- Based on BART NLP model
- Trained on 1.5B small molecules in SMILES format
- Useful for representation and sequence translation tasks
- Not well suited for generation tasks -- lacks an organized and uniformly shaped latent space

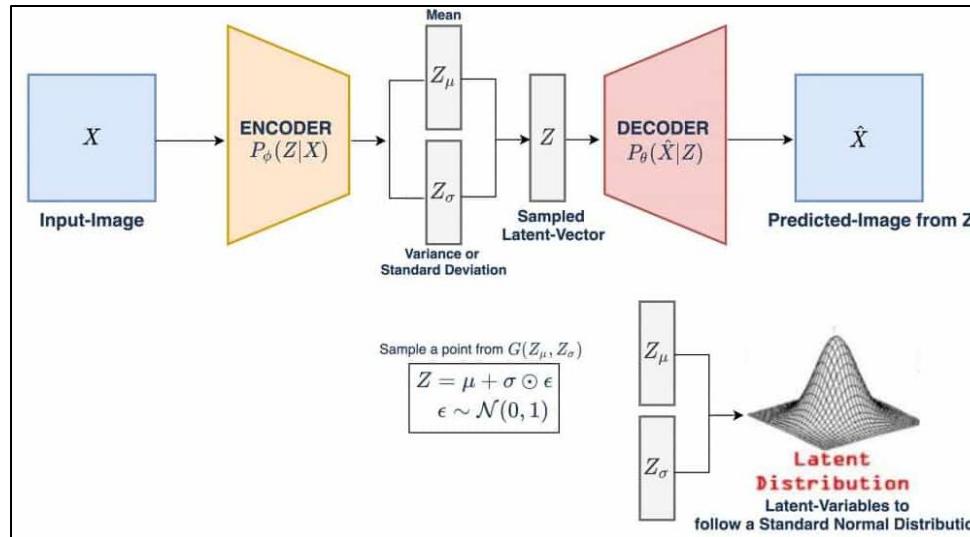


# Autoencoder Models in a Nutshell

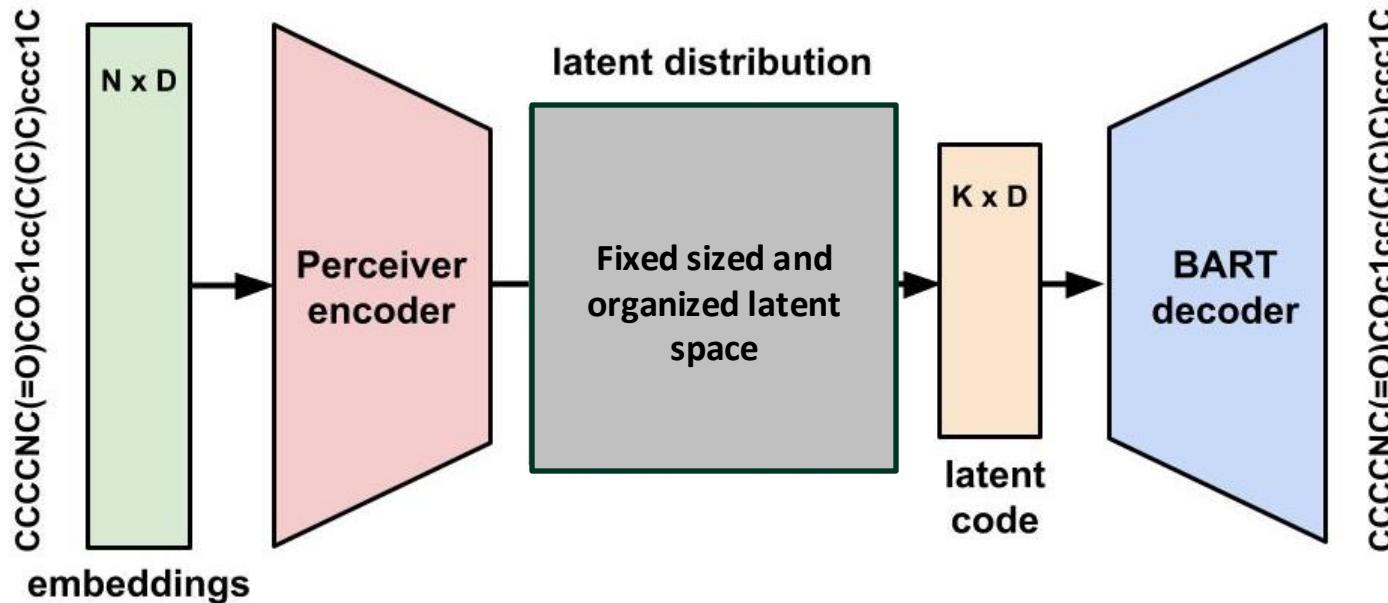
Autoencoder



Variational  
Autoencoder (VAE)

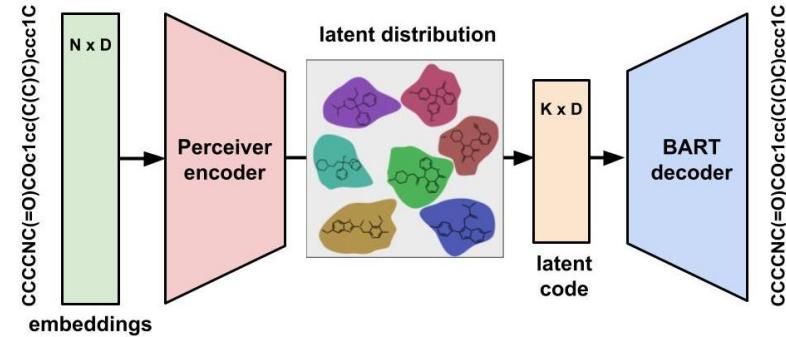
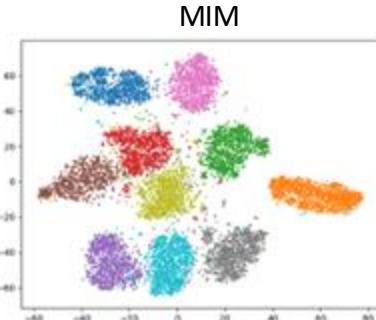
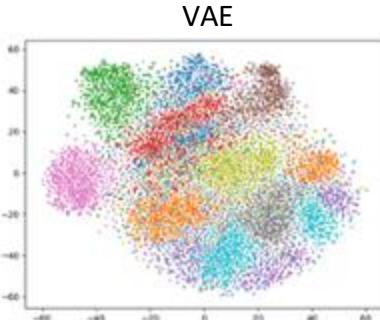


# Development of MolMIM for Molecule Generation



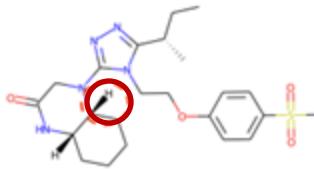
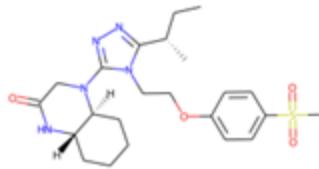
# A Clustered Latent Space with Mutual Information Machine

- Mutual information machine (MIM) has a loss function that maximizes mutual information and minimizes marginal entropy
- MIM loss results in a clustered space while variational autoencoder (VAE) loss smooths the latent space resulting in blurring

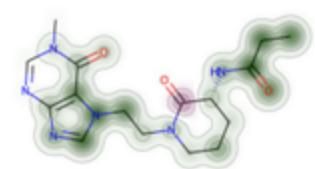
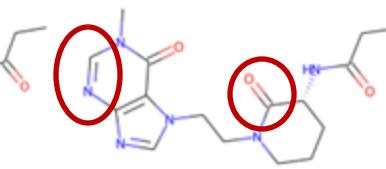
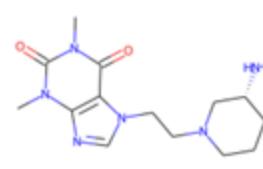


# MolMIM – Sampling Distance Can Be Tuned for Similarity

**Small Perturbations**

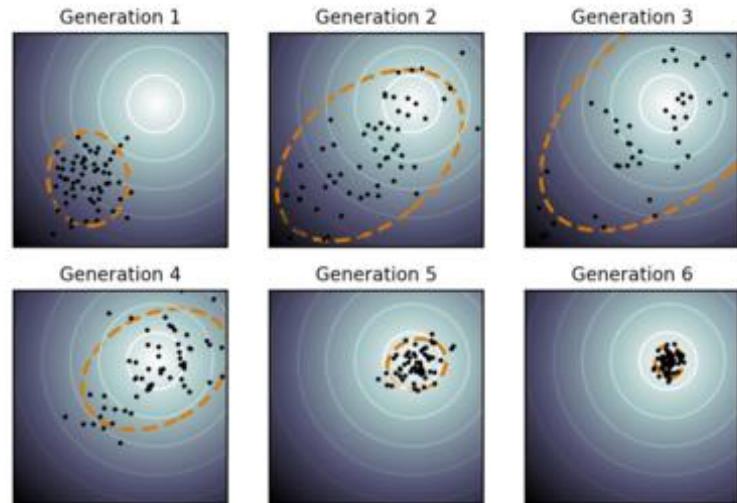


**Larger Perturbations**



# Measuring the Controllability of MolMIM Generation

- **Hypothesis:** having a structured latent space will improve performance of property guided optimization
- Chose covariance matrix adaptation (CMA-ES), which is a zeroth order optimization method
- CMA-ES is non-parametric and uses only a single scoring function per sample



# Multi-Objective Property Optimization

- Performed multi-objective optimization to jointly optimize two molecule properties (QED, SA) and binding to two protein targets (JNK3, GSK4 $\beta$ )
- Novelty is proportion of molecules with similarity metric (0.0 – 1.0) less than  $\leq 0.4$  relative to any other molecule
- Diversity is average similarity across all compounds
- MolMIM is competitive for success and diversity, but novelty has room for improvement

Model	QED + SA + JNK3 + GSK4 $\beta$		
	Success (%)	Novelty (%)	Diversity
RationaleRL	74.8	56.1	0.621
MARS	92.3	82.4	0.719
JANUS	<b>100</b>	32.6	<b>0.821</b>
FaST	<b>100</b>	<b>100</b>	0.716
MolMIM (R)	97.5	71.1	0.791
MolMIM (A)	96.6	63.3	0.807
MolMIM (E)	98.3	55.1	0.767
MolMIM (E) <sup>+</sup>	99.2	54.8	0.772

# MolMIM: Research to Productization



**arXiv > cs > arXiv:2208.09016**

**Computer Science > Machine Learning**

[Submitted on 18 Aug 2022 (v1), last revised 29 Mar 2023 (this version, v2)]

**Improving Small Molecule Generation using Mutual Information Machine**

Danny Reidenbach, Micha Livne, Rajesh K. Ilango, Michelle Gill, Johnny Israeli

We address the task of controlled generation of small molecules, which entails finding novel molecules with desired properties under certain constraints (e.g., similarity to a reference molecule). Here we introduce MolMIM, a probabilistic auto-encoder for small molecule drug discovery that learns an informative and clustered latent space. MolMIM is trained with Mutual Information Machine (MIM) learning, and provides a fixed length representation of variable length SMILES strings. Since encoder-decoder models can learn representations with "holes" of invalid samples, here we propose a novel extension to the training procedure which promotes a dense latent space, and allows the model to sample valid molecules from random perturbations of latent codes. We provide a thorough comparison of MolMIM to several variable-size and fixed-size encoder-decoder models, demonstrating MolMIM's superior generation as measured in terms of validity, uniqueness, and novelty. We then utilize CMA-ES, a naive black-box and gradient free search algorithm, over MolMIM's latent space for the task of property guided molecule optimization. We achieve state-of-the-art results in several constrained tasks.

**ICLR**

Poster  
in  
Workshop: [Machine Learning for Drug Discovery \(MLDD\)](#)

**Improving Small Molecule Generation using Mutual Information Machine**

Danny Reidenbach · Micha Livne · Rajesh Ilango · Michelle Gill · Johnny Israeli

[ Abstract ] [ Project Page ]

[  Poster ] [  OpenReview ]

Fri 5 May 10 a.m. PDT – 10:55 a.m. PDT

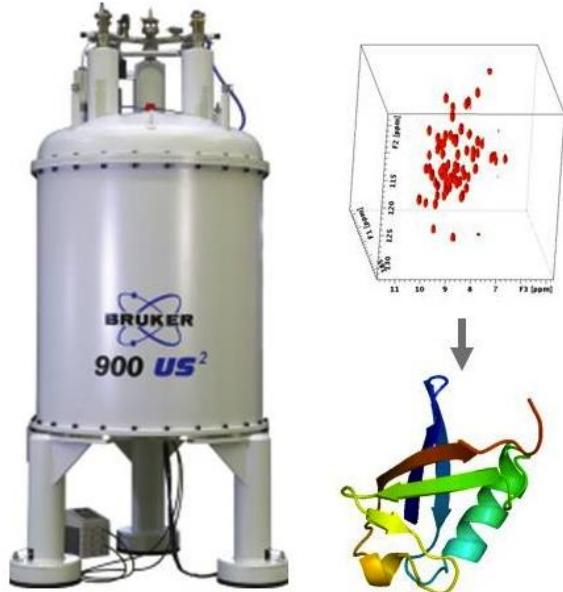
- Integration of MolMIM model into BioNeMo inference service
- Productionize model architecture and training framework
- Accelerated inference
- Improving encoder representations

The background of the slide features a dark, black space filled with glowing green lines. These lines are thin and wispy on the left side, creating a sense of motion and depth. As they move towards the right, they converge and thicken into a dense, glowing green mass that forms a large, stylized letter 'A'. The lines have a slight glow and some internal structure, resembling fiber optics or light trails.

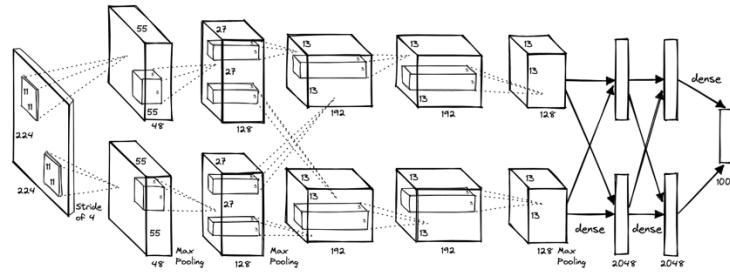
**“How I Got Here” and Lessons Learned Along  
the Way**

# From Structural Biologist to Data Scientist

Postdoctoral Research: Enzyme Dynamics by  
NMR Spectroscopy



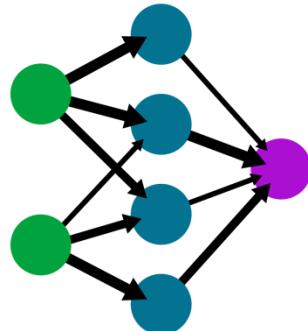
AlexNet Won ImageNet Challenge in 2012



AlexNet didn't just win; it dominated. AlexNet was unlike the other competitors. This new model demonstrated unparalleled performance on the largest image dataset of the time, ImageNet. This event made AlexNet the first widely acknowledged, successful application of deep learning.

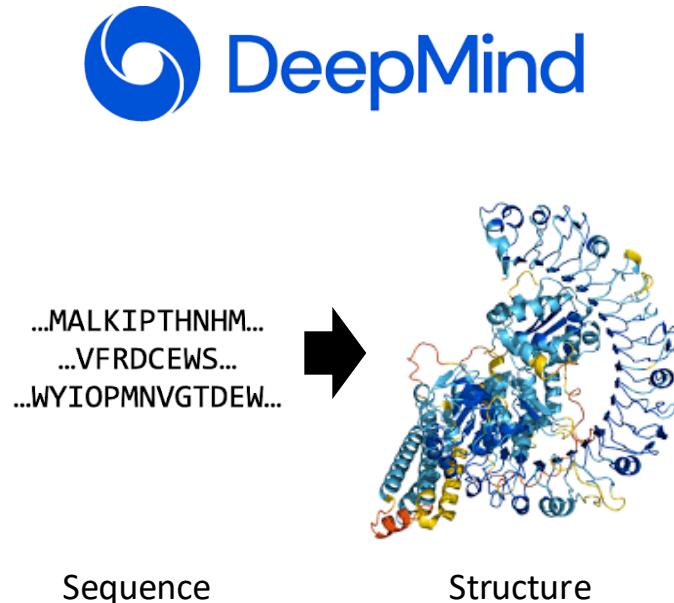
Don't miss the bigger picture: Machine learning will have an impact on every industry.

## From Structural Biologist to Data Scientist



**nVIDIA.**<sup>®</sup>

# A Deep Learning Model Became the World's Best Protein Structure Predictor



C  
A  
S  
P  
13

Google DeepMind About Technologies Impact Discover

Overview Blog The Podcast Visualising AI

RESEARCH

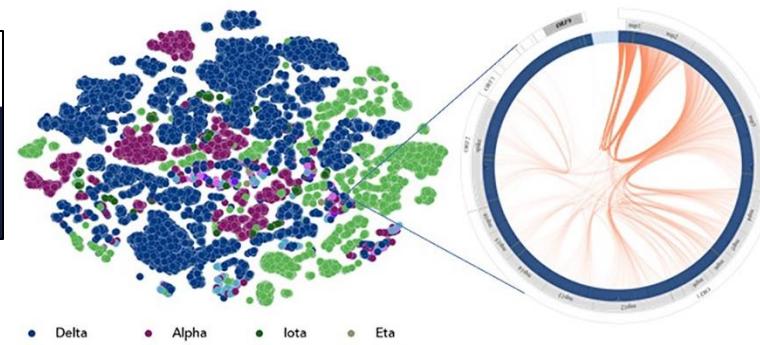
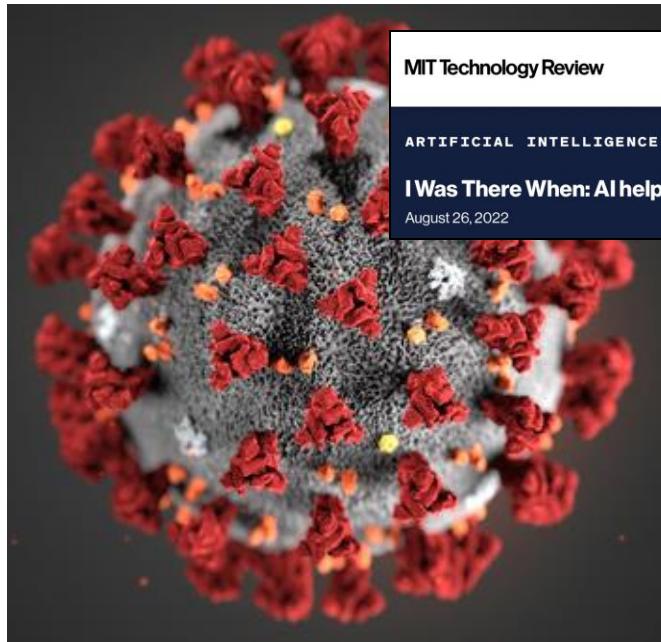
AlphaFold: a solution to a 50-year-old grand challenge in biology

CASP15: AlphaFold's success spurs new challenges in ...

Dec 14, 2022 — Two years later, **AlphaFold** still dominates the competition. Deepmind itself did not participate in this round, but **AlphaFold** has been open ...

AlphaFold won the Critical Assessment of Protein Structure Prediction (CASP13) Competition in 2018 ... and has done so every year since

# AI and the Race for a COVID-19 Vaccine



Genome-scale language models (GenSLMs) discover distinct evolutionary patterns in SARS-CoV-2

Argonne NATIONAL LABORATORY

## Argonne researchers win Gordon Bell Special Prize for adapting language models to track virus variants

BY KEVIN JACKSON | NOVEMBER 29, 2022

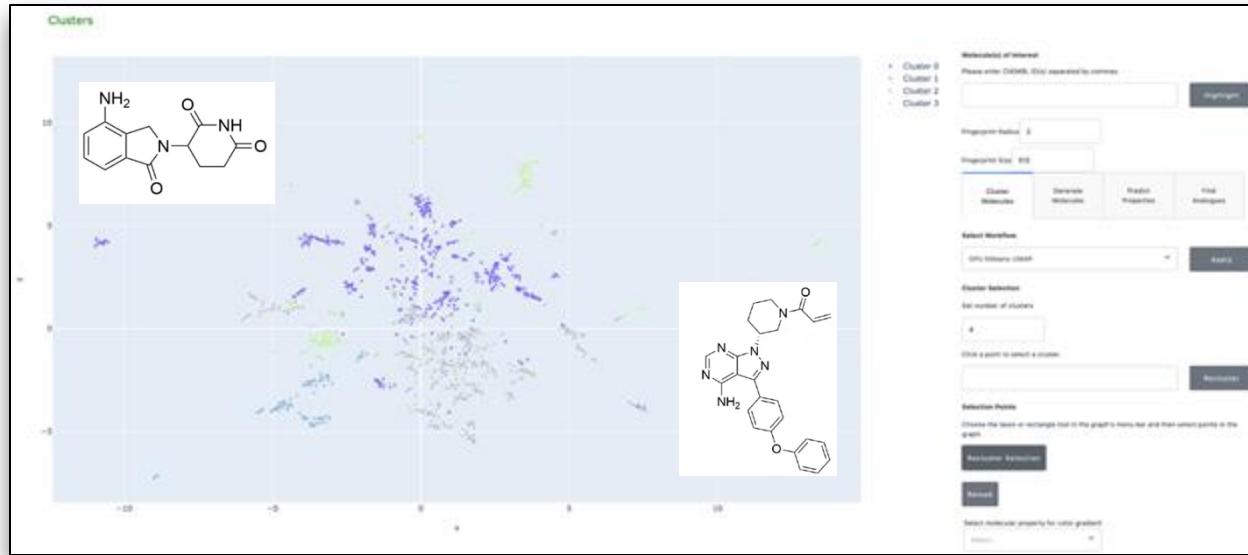
Groundbreaking research focuses on understanding genomic sequences to catch more deadly variants of COVID-19.

News

Media Contacts

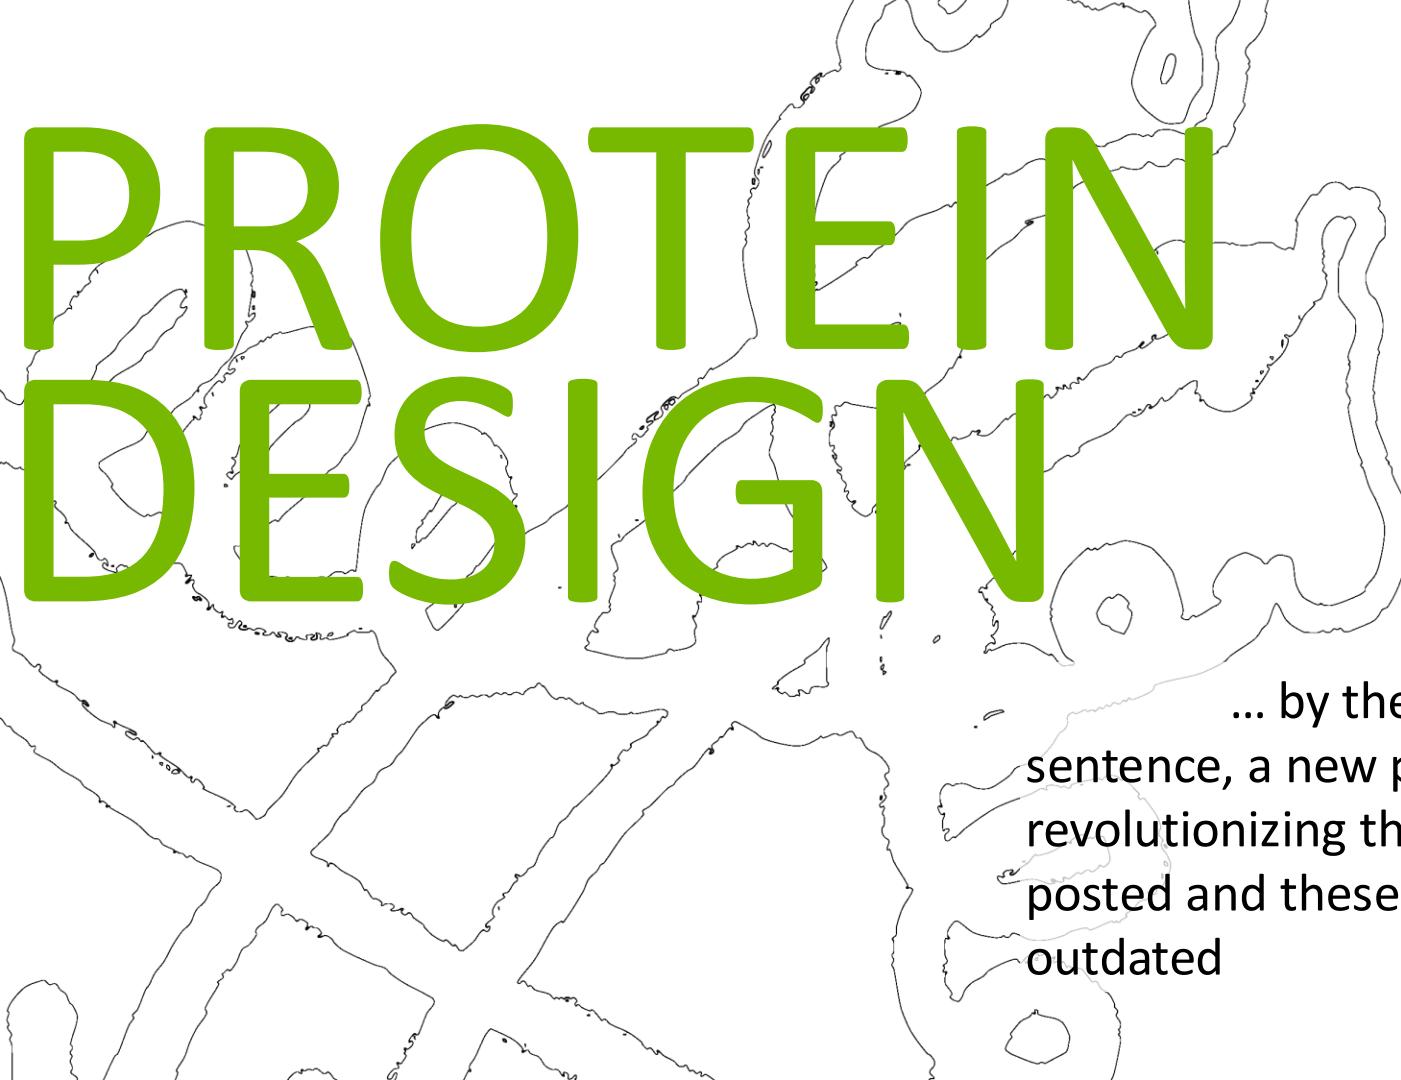
Experts Guide

# First Effort: Interface for Clustering and Visualization of Small Molecules

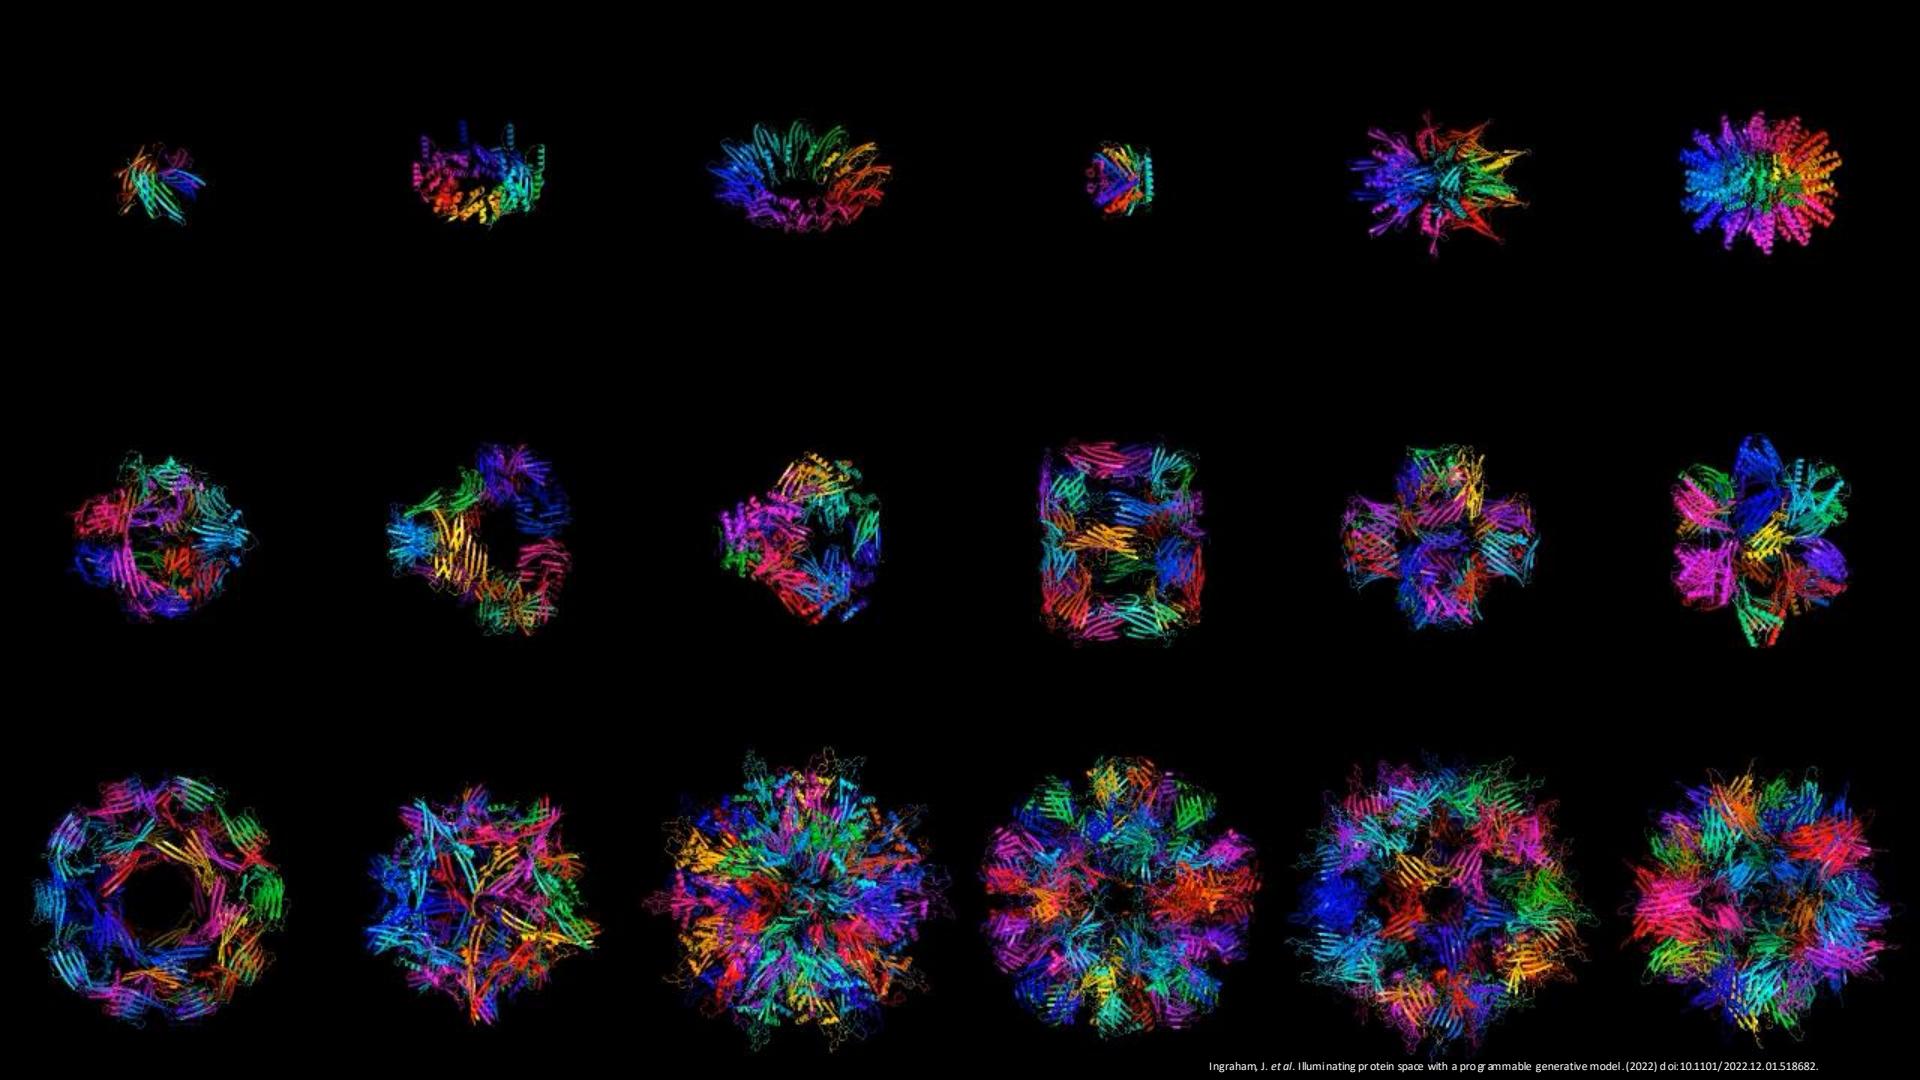


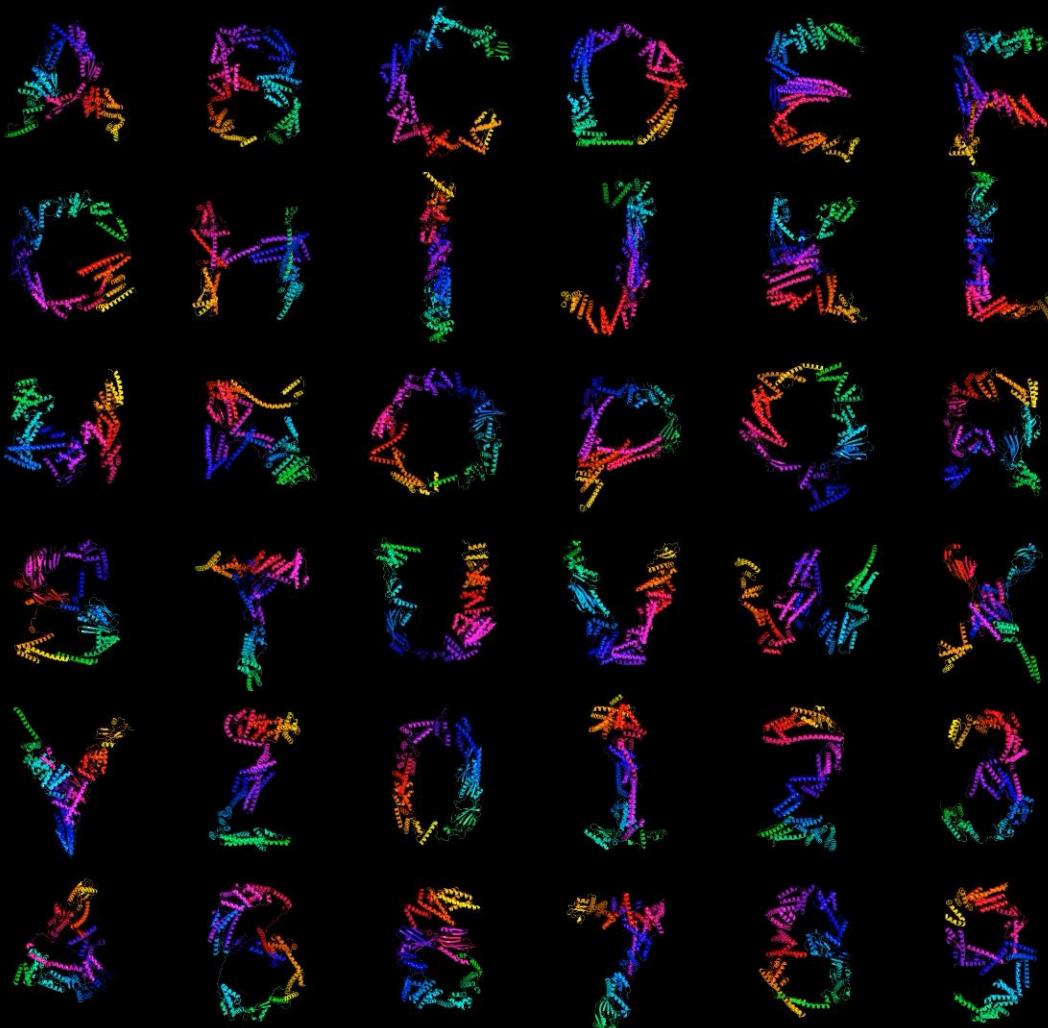
Deep learning is high risk. Ensure the project will succeed if deep learning fails.

# PROTEIN DESIGN

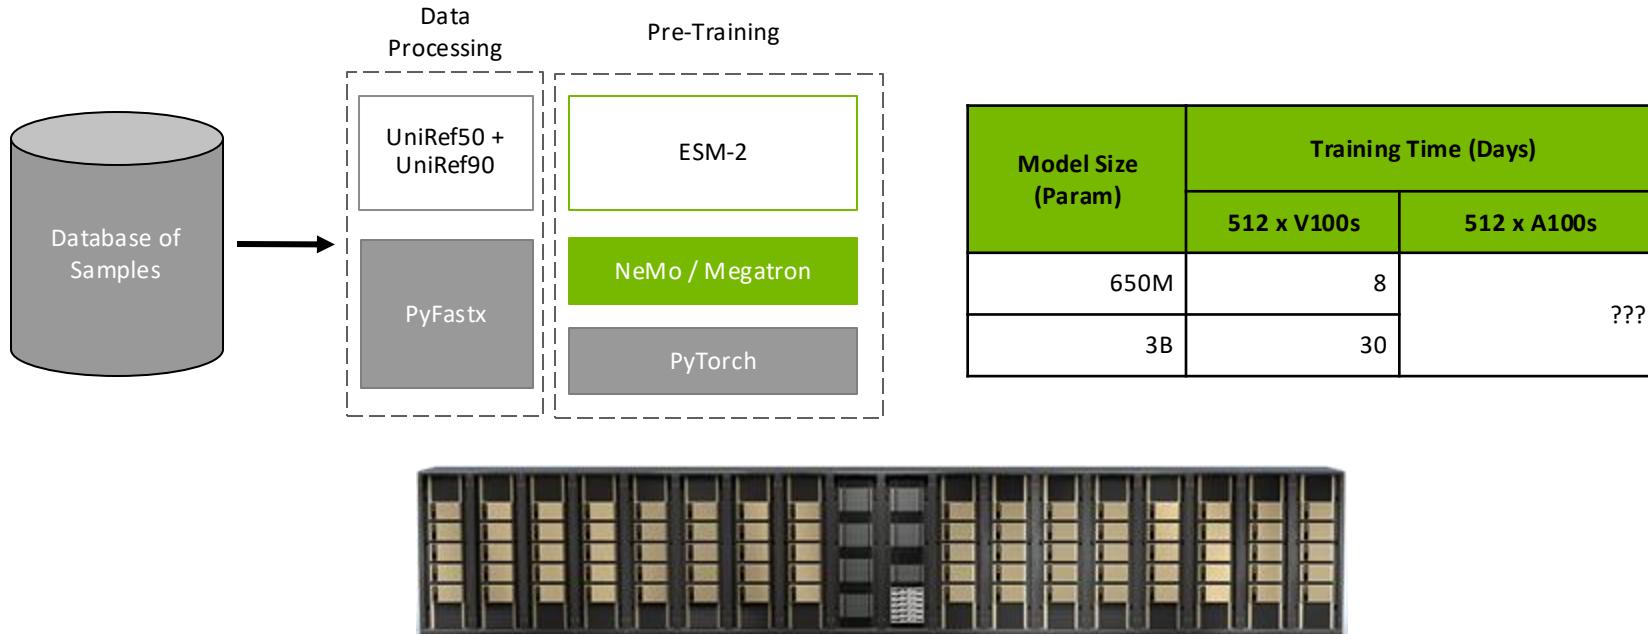


... by the time you've read this sentence, a new pre-print revolutionizing the field has been posted and these slides are totally outdated



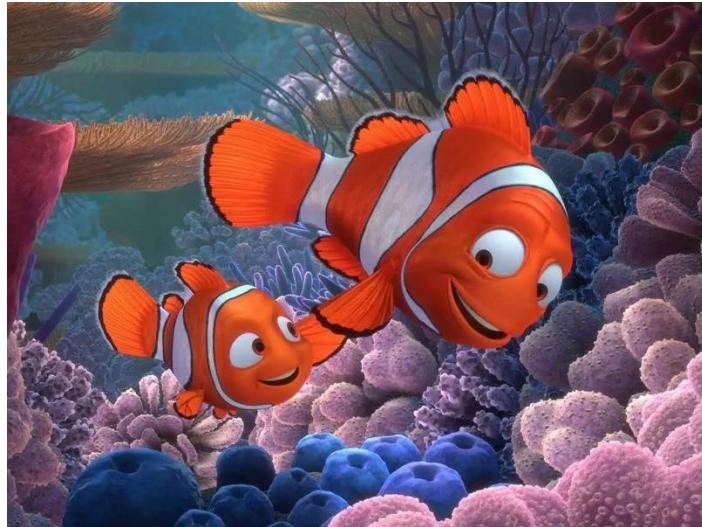


# Developing Deep Learning Models at Scale

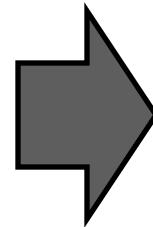


Successes from calculated risks provide justification for growing a team.

## Rapid Team Growth and Adventures in Management



Two Engineers



< Two Years



Over Thirty Engineers

Deep learning is hard, but growing and managing a team is the most challenging problem.

## Conclusions

- BioNeMo is a framework and inference service for developing, training, deploying, and using deep learning models and tools for drug discovery
- MolMIM is a cheminformatics language model trained on SMILES with a structured latent space for molecule design
- Careers are long compared to the pace of machine learning advancement
- Capitalize on new opportunities and enjoy the ride!

BioNeMo Inference Service early access : <https://www.nvidia.com/bionemo>  
BioNeMo Framework general access coming next week!

# The BioNeMo Team

Johnny Israeli	Gagan Kaushik	Ohad Mosafi
	George Armstrong	Pablo Ribalta
Alireza Moradzadeh	Guoqing Zhou	Rajesh Ilango
Arkadiusz Nowaczynski	Han-Yi Chou	Sara Rabhi
Camir Ricketts	Jasleen Grewal	Simon Chu
Danny Reidenbach	Kevin Boyd	Srimukh Veccham
Dejun Lin	Maria Korshunova	Steven Kohen-Hill
Dorota Toczydlowska	Mario Geiger	Tomasz Grzegorzek
Emine Kucukbenli	Marta Stepniewska-Dziubinska	Timur Rvachov
Eric Dawson	Micha Livne	Yuxing Peng
Farhad Ramezanghorbani	Neha Tadimeti	Zachary McClure

# Thank You!

## Questions:

Fireside Chat

10:15 – 10:55am

Central Park East

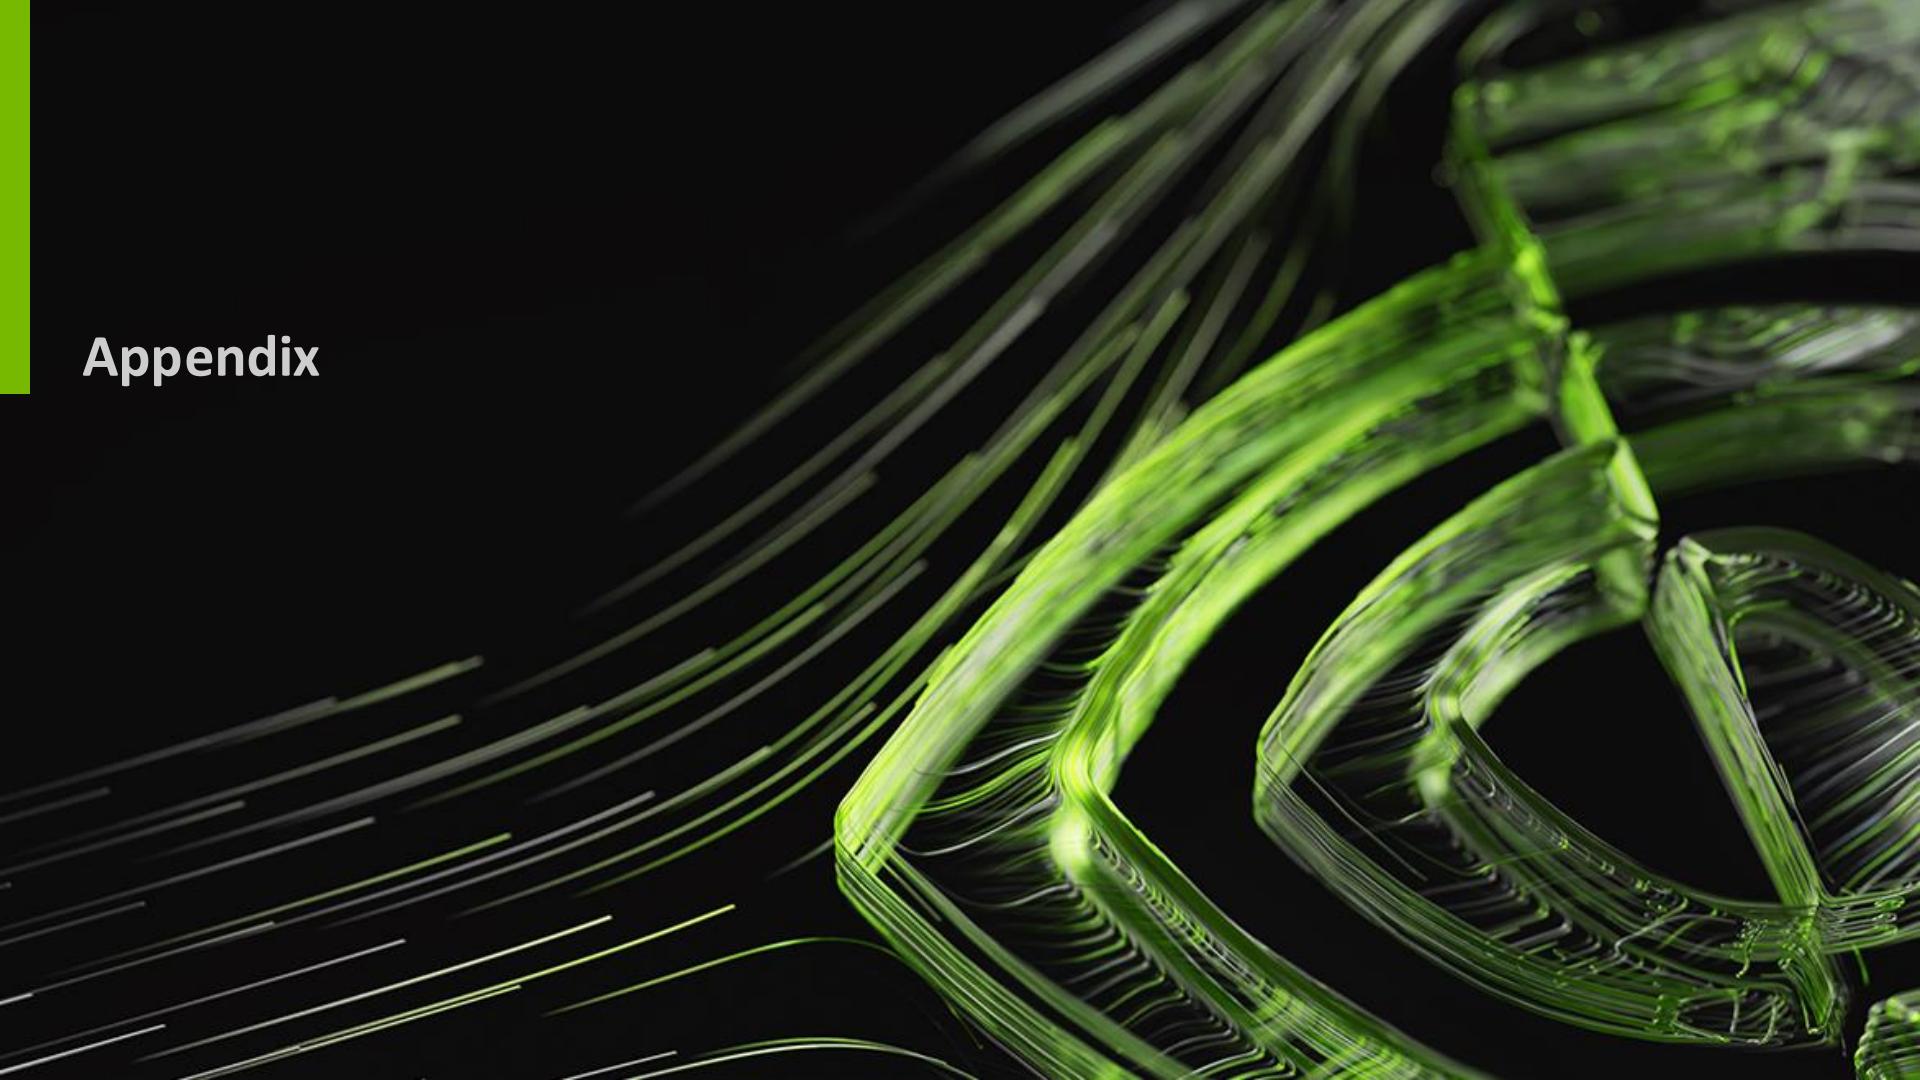


[mgill@nvidia.com](mailto:mgill@nvidia.com)

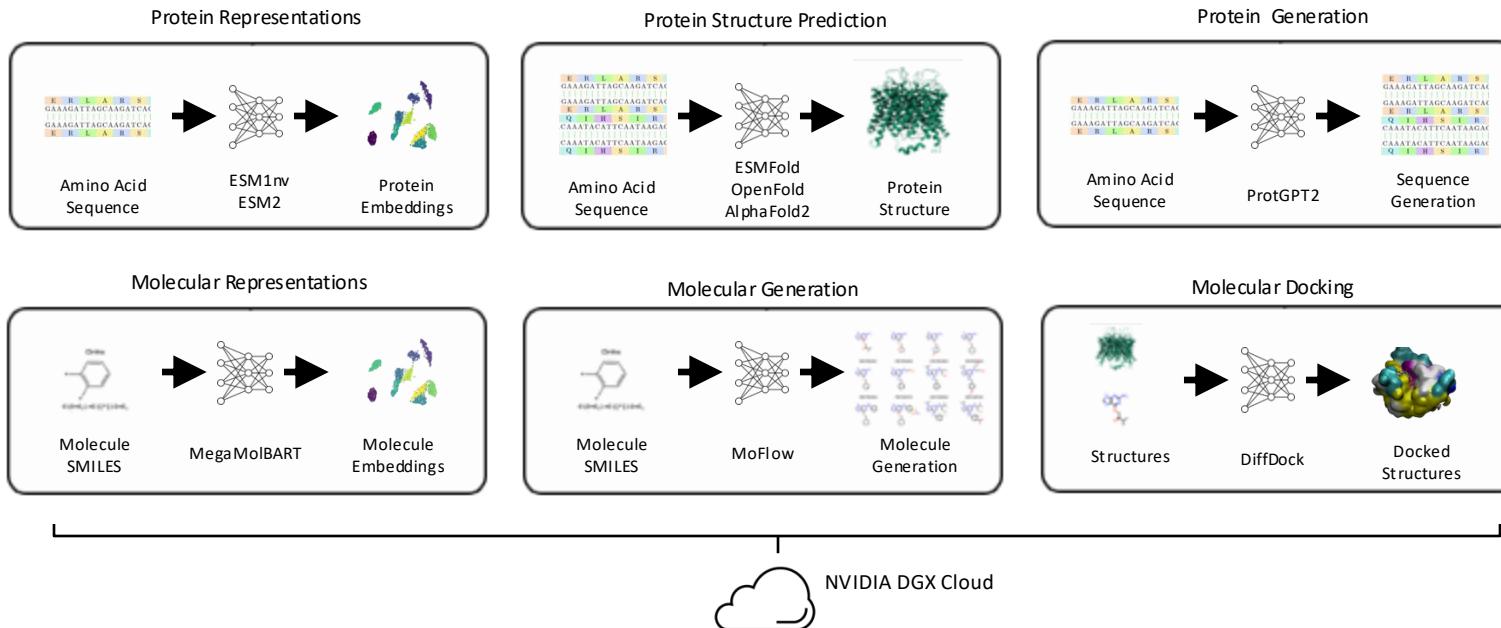


[michellelynngill.com](http://michellelynngill.com)

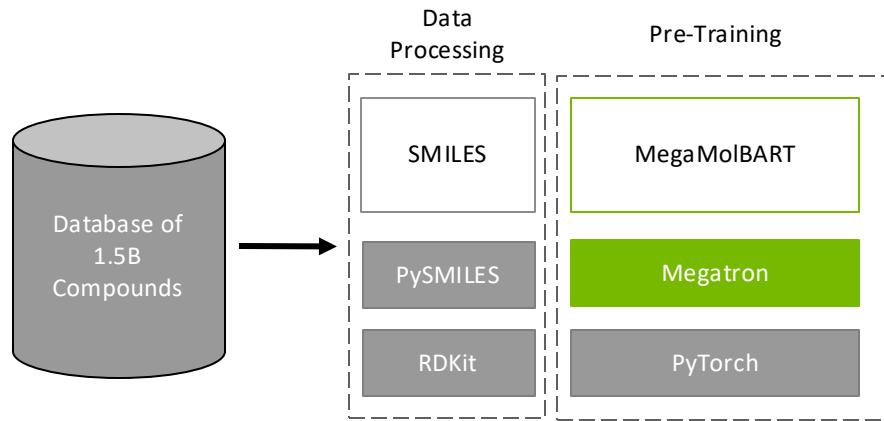
# Appendix



# Nine Models in Inference Service for Drug Discovery Applications



# Deep Learning at Scale

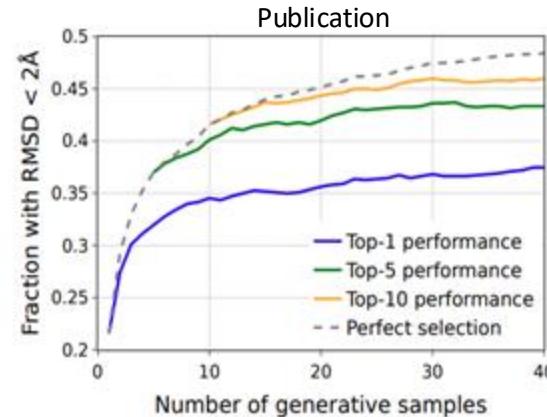
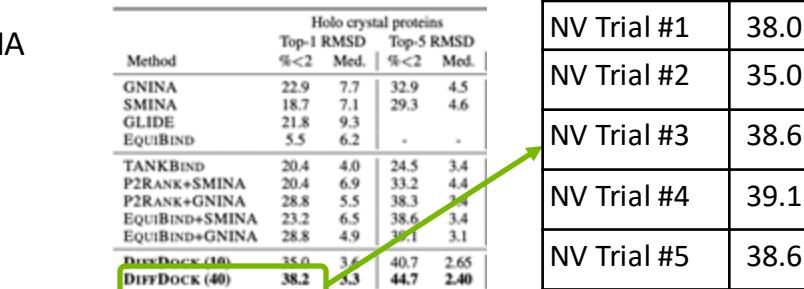
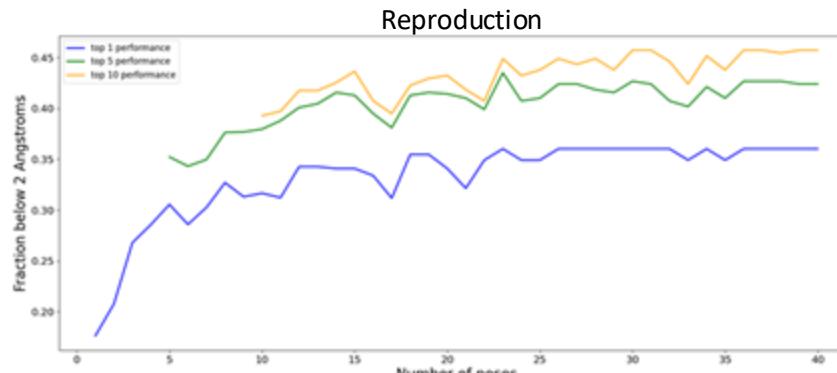


Attention Heads	Layers	Hidden Size	Feed Forward	Parameters
8	4	256	1024	10M
8	6	512	2048	45M
16	8	1024	4096	230M



# Life Cycle of a BioNeMo Model in the Inference Service

- Model checkpoints are accelerated using a variety of NVIDIA tools – standard tricks to custom CUDA kernels
- All quantitative and qualitative results are reproduced
- For DiffDock, the RMSD metrics were reproduced under a variety of different conditions



# Proteins Generated from Evozyne's ProT-VAE Models

ProT-VAE: Protein Transformer Variational AutoEncoder for Functional Protein Design

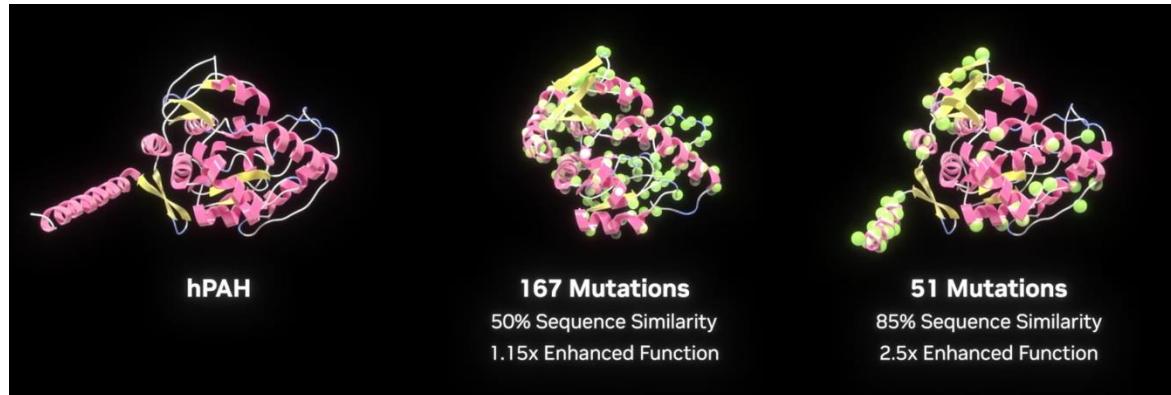
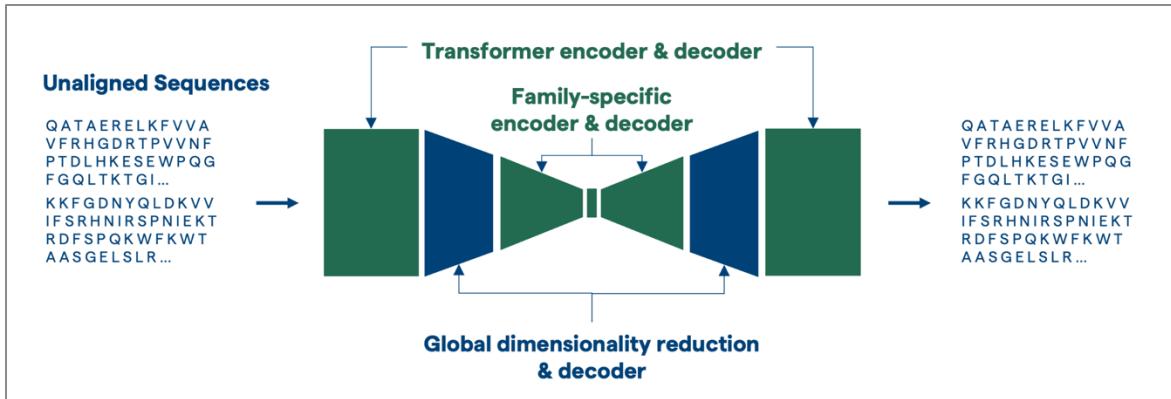
Emre Sevgen<sup>1†</sup>, Joshua Moller<sup>1†</sup>, Adrian Lange<sup>1</sup>, John Parker<sup>1</sup>, Sean Quigley<sup>1</sup>, Jeff Mayer<sup>1</sup>, Poonam Srivastava<sup>1</sup>, Sitaram Gayatri<sup>1</sup>, David Hosfield<sup>1</sup>, Maria Korshunova<sup>2</sup>, Micha Livne<sup>2</sup>, Michelle Gill<sup>2</sup>, Rama Ranganathan<sup>1</sup>, Anthony B. Costa<sup>2\*</sup> and Andrew L. Ferguson<sup>1\*</sup>

<sup>1</sup>Evozyne, Inc., 2430 N Halsted Street, Chicago, 60614, IL, USA.

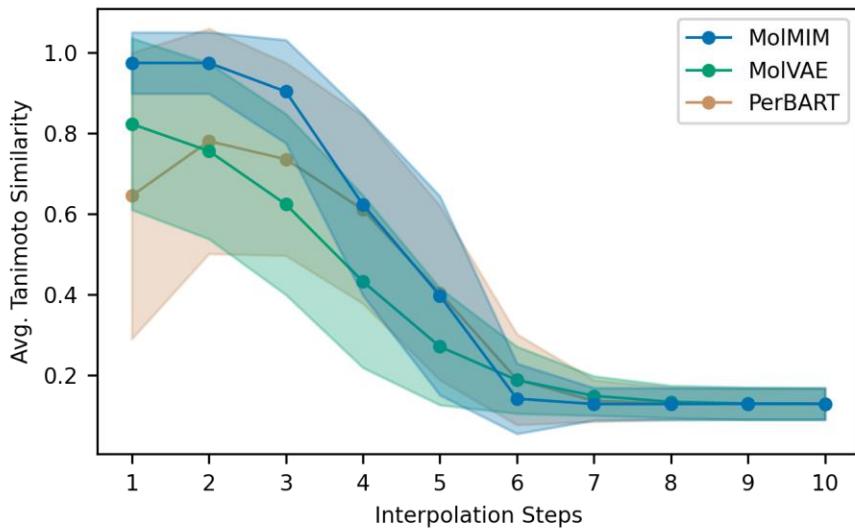
<sup>2</sup>NVIDIA, 2788 San Tomas Expressway, Santa Clara, 95051, CA, USA.

\*Corresponding author(s). E-mail(s): [acosta@nvidia.com](mailto:acosta@nvidia.com);  
[andrew.ferguson@evozyne.com](mailto:andrew.ferguson@evozyne.com);

†These authors contributed equally to this work.



# Probing Latent Structure by Molecule Interpolation



- Pairwise interpolations performed at ten evenly spaced steps for 1,000 ZINC15 molecules
- Average Tanimoto similarity to first molecule was calculated at each step
- Molecules sampled from Perceiver BART and MolVAE have reduced similarity to start and a large degree of variability at early interpolation steps
- Molecules sampled from MolMIM are similar and have the smallest variance at early steps

# MolMIM – Performance on Seed Based Molecule Sampling

- Randomly sampled ten molecules for each of 20k molecules from test split
- Effective novelty is percentage of molecules that are valid, unique, not identical to seed, and novel
- Sampling radius empirically determined to maximize effective novelty
- CDDD used as baseline model – trained with molecular property loss
- Perceiver BART sampling speed improved relative to MegaMolBART
- MolVAE and MolMIM show significant improvements in validity and effective novelty

Model	Latent Dim	Validity (%)	Uniqueness (%)	Novelty (%)	Effective Novelty (%)	Test Runtime
MegaMolBART	Variable	75.0	84.8	94.2	51.1	8.7 hours
Perceiver BART	2048	71.8	94.9	94.6	59.1	38 min
MolVAE	2048	95.7	<b>100.0</b>	98.1	93.9	64 min
<b>MolMIM</b>	<b>512</b>	<b>98.7</b>	<b>100.0</b>	95.5	<b>94.2</b>	30 min
CDDD	512	84.5	98.9	<b>99.5</b>	82.2	12 hours <sup>†</sup>

<sup>†</sup>CDDD decoding speed limited by batch size.

# Single Property Optimization with CMA-ES

Model	QED (%)		Penalized logP $\delta \geq 0.6$
	$\delta \geq 0.4$	$\delta \geq 0.4$	
AtomG2G	73.6	-	-
HeirG2G	76.9	-	-
DESMILES	77.8	-	-
QMO	92.8	$7.71 \pm 5.65$	$3.73 \pm 2.85$
MolGrow	-	$8.34 \pm 6.85$	$4.06 \pm 5.61$
GraphAF	-	$8.21 \pm 6.51$	$4.98 \pm 6.49$
GraphDF	-	$9.19 \pm 6.43$	$4.51 \pm 5.80$
CDGS	-	$9.56 \pm 6.33$	$5.10 \pm 5.80$
FaST	-	$18.09 \pm 8.72$	$8.98 \pm 6.31$
MolMIM	<b>94.6</b>	<b><math>28.45 \pm 54.67</math></b>	<b><math>7.60 \pm 23.62</math></b>
MolMIM		$9.44 \pm 4.12^{\dagger}$	$4.57 \pm 3.87^{\dagger}$

- Performed optimization of QED or penalized logP with query budget of 50,000 oracle calls per input molecule
- Success is % of molecules with  $\text{QED} \geq 0.9$  or penalized logP improvement while maintaining Tanimoto similarity  $\delta \geq \{0.4, 0.6\}$
- MolMIM achieves the highest QED and logP success rates
- Penalized logP results impacted by known exploit where identical functional groups are repeatedly added

Results above solid bar as in B. Chen, X. Fu, R. Barzilay, T. Jaakkola, ArXiv (2021) and S. C. Hoffman, *et al*, Nat Mach Intell. 4, 21–31 (2022)  
 QED and logP oracles from Therapeutic Data Commons.  
<sup>†</sup>logP improvement limited to  $\leq 20$

# Single Property Optimization with CMA-ES

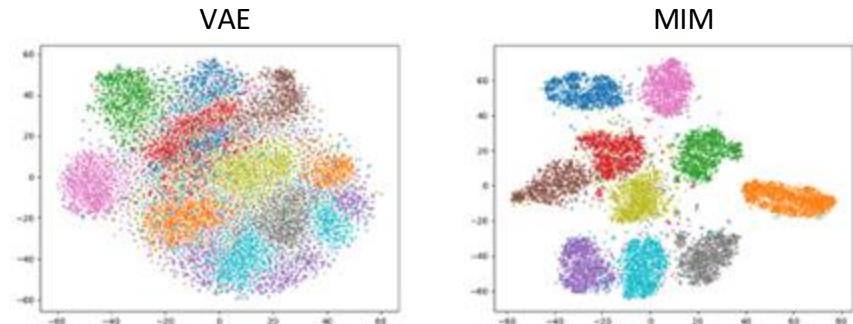
Model	QED (%)		Penalized logP $\delta \geq 0.6$
	$\delta \geq 0.4$	$\delta \geq 0.4$	
AtomG2G	73.6	-	-
HeirG2G	76.9	-	-
DESMILES	77.8	-	-
<b>QMO</b>	<b>92.8</b>	$7.71 \pm 5.65$	$3.73 \pm 2.85$
MolGrow	-	$8.34 \pm 6.85$	$4.06 \pm 5.61$
GraphAF	-	$8.21 \pm 6.51$	$4.98 \pm 6.49$
GraphDF	-	$9.19 \pm 6.43$	$4.51 \pm 5.80$
CDGS	-	$9.56 \pm 6.33$	$5.10 \pm 5.80$
<b>FaST</b>	-	$18.09 \pm 8.72$	$8.98 \pm 6.31$
<b>MolMIM</b>	<b>94.6</b>	<b><math>28.45 \pm 54.67</math></b>	<b><math>7.60 \pm 23.62</math></b>
MolMIM		$9.44 \pm 4.12^{\dagger}$	$4.57 \pm 3.87^{\dagger}$

- Performed optimization of QED or penalized logP with query budget of 50,000 oracle calls per input molecule
- Success is % of molecules with  $\text{QED} \geq 0.9$  or penalized logP improvement while maintaining Tanimoto similarity  $\delta \geq \{0.4, 0.6\}$
- MolMIM achieves the highest QED and logP success rates
- Penalized logP results impacted by known exploit where identical functional groups are repeatedly added
- Recall: MolMIM trained without chemical properties, activity, or fragment knowledge

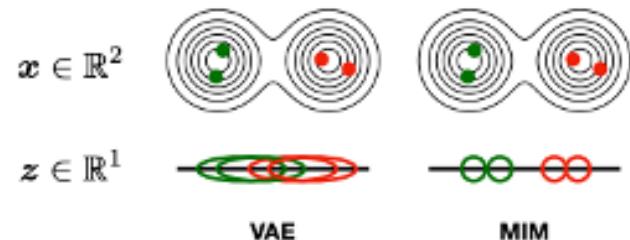
Results above solid bar as in B. Chen, X. Fu, R. Barzilay, T. Jaakkola, ArXiv (2021) and S. C. Hoffman, *et al*, Nat Mach Intell. 4, 21–31 (2022)  
 QED and logP oracles from Therapeutic Data Commons.  
<sup>†</sup>logP improvement limited to  $\leq 20$

# A Clustered Latent Space with Mutual Information Machine

- Same architecture as VAE, but loss maximizes mutual information and minimizes marginal entropy
- MIM results in an informative and clustered latent space



$$\begin{aligned}\mathcal{L}_{\text{A-MIM}}(\theta) &= \frac{1}{2} \left( CE(\mathcal{M}_S^q(\mathbf{x}, \mathbf{z}), q_{\theta}(\mathbf{x}, \mathbf{z})) \right. \\ &\quad \left. + CE(\mathcal{M}_S^q(\mathbf{x}, \mathbf{z}), p_{\theta}(\mathbf{x}, \mathbf{z})) \right) \\ &\geq H_{\mathcal{M}_S^q}(\mathbf{x}) + H_{\mathcal{M}_S^q}(\mathbf{z}) - I_{\mathcal{M}_S^q}(\mathbf{x}; \mathbf{z})\end{aligned}$$



Model	QED (%)		Penalized logP	
	$\delta \geq 0.4$	$\delta \geq 0.4$	$\delta \geq 0.6$	$\delta \geq 0.6$
JT-VAE	8.8	$1.03 \pm 1.39$	$0.28 \pm 0.79$	
GCPN	9.4	$2.49 \pm 1.30$	$0.79 \pm 0.63$	
MoIDQN	-	$3.37 \pm 1.62$	$1.86 \pm 1.21$	
MMPA	32.9	-	-	-
VSeq2Seq	58.5	$3.37 \pm 1.75$	$2.33 \pm 1.17$	
VJTNN+GAN	60.6	-	-	-
VJTNN	-	$3.55 \pm 1.67$	$2.33 \pm 1.24$	
MoFlow	-	$4.71 \pm 4.55$	$2.10 \pm 2.86$	
GA	-	$5.93 \pm 1.41$	$3.44 \pm 1.09$	
AtomG2G	73.6	-	-	-
HeirG2G	76.9	-	-	-
DESMILES	77.8	-	-	-
QMO	92.8	$7.71 \pm 5.65$	$3.73 \pm 2.85$	
MolGrow	-	$8.34 \pm 6.85$	$4.06 \pm 5.61$	
GraphAF	-	$8.21 \pm 6.51$	$4.98 \pm 6.49$	
GraphDF	-	$9.19 \pm 6.43$	$4.51 \pm 5.80$	
CDGS	-	$9.56 \pm 6.33$	$5.10 \pm 5.80$	
FaST	-	$18.09 \pm 8.72$	$8.98 \pm 6.31$	
MolMIM	<b>94.6</b>	<b><math>28.45 \pm 54.67</math></b>	<b><math>7.60 \pm 23.62</math></b>	
MolMIM		$9.44 \pm 4.12^{\dagger}$	$4.57 \pm 3.87^{\dagger}$	

## Single Property Optimization

- Repeated QED and penalized logP optimization with query budget of 50,000 oracle calls per input molecule
- Success is % of molecules with  $\text{QED} \geq 0.9$  or penalized logP improvement while maintaining Tanimoto similarity  $\delta \geq \{0.4, 0.6\}$
- MolMIM achieves the highest QED and logP success rates
- Penalized logP results impacted by known exploit where identical functional groups are repeatedly added
- MolMIM results were repeated with logP improvement limited

Results above solid bar as in B. Chen, X. Fu, R. Barzilay, T. Jaakkola, ArXiv (2021) and S. C.

Hoffman, *et al*, Nat Mach Intell. 4, 21–31 (2022).

<sup>†</sup>logP improvement limited to  $\leq 20$

## Perspective on BioNeMo

- Models have a finite lifespan, the value is in the learnings
- Developing and productizing internal research is useful for driving improvements to the platform
- Scalability and acceleration are differentiating factors
- Surface NVIDIA technologies, and use bottlenecks to drive the development software and hardware improvements