
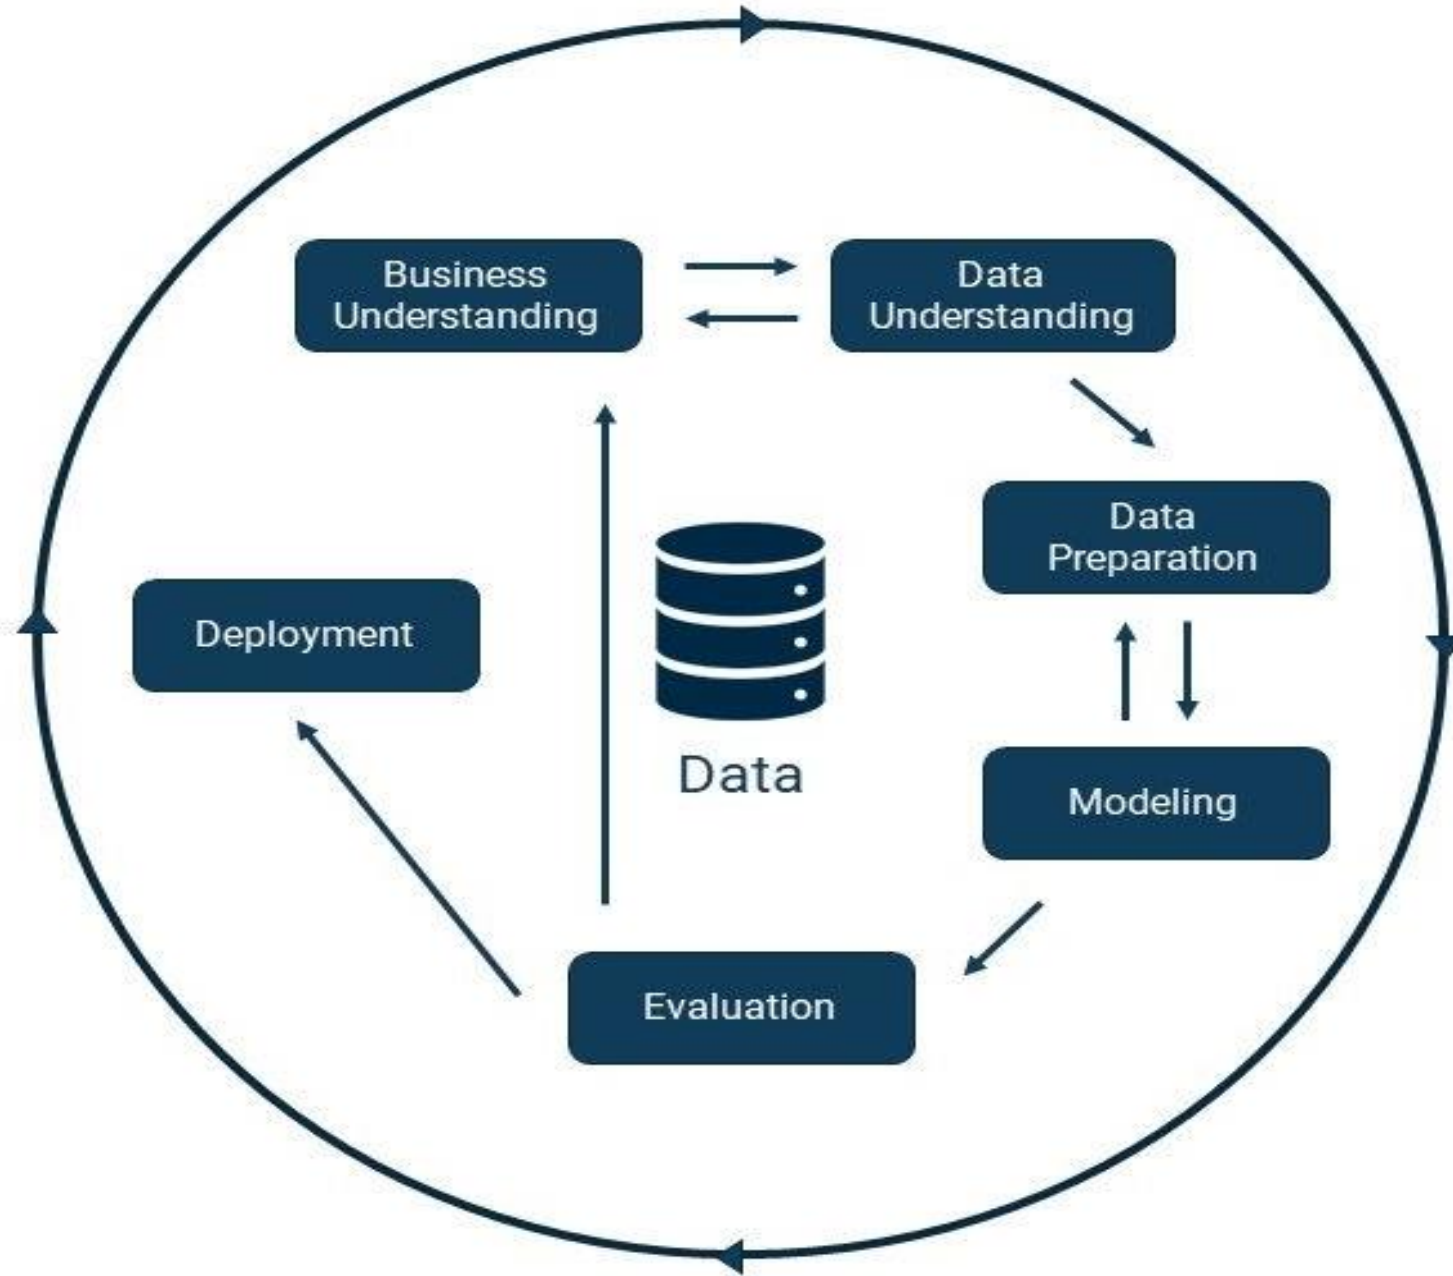


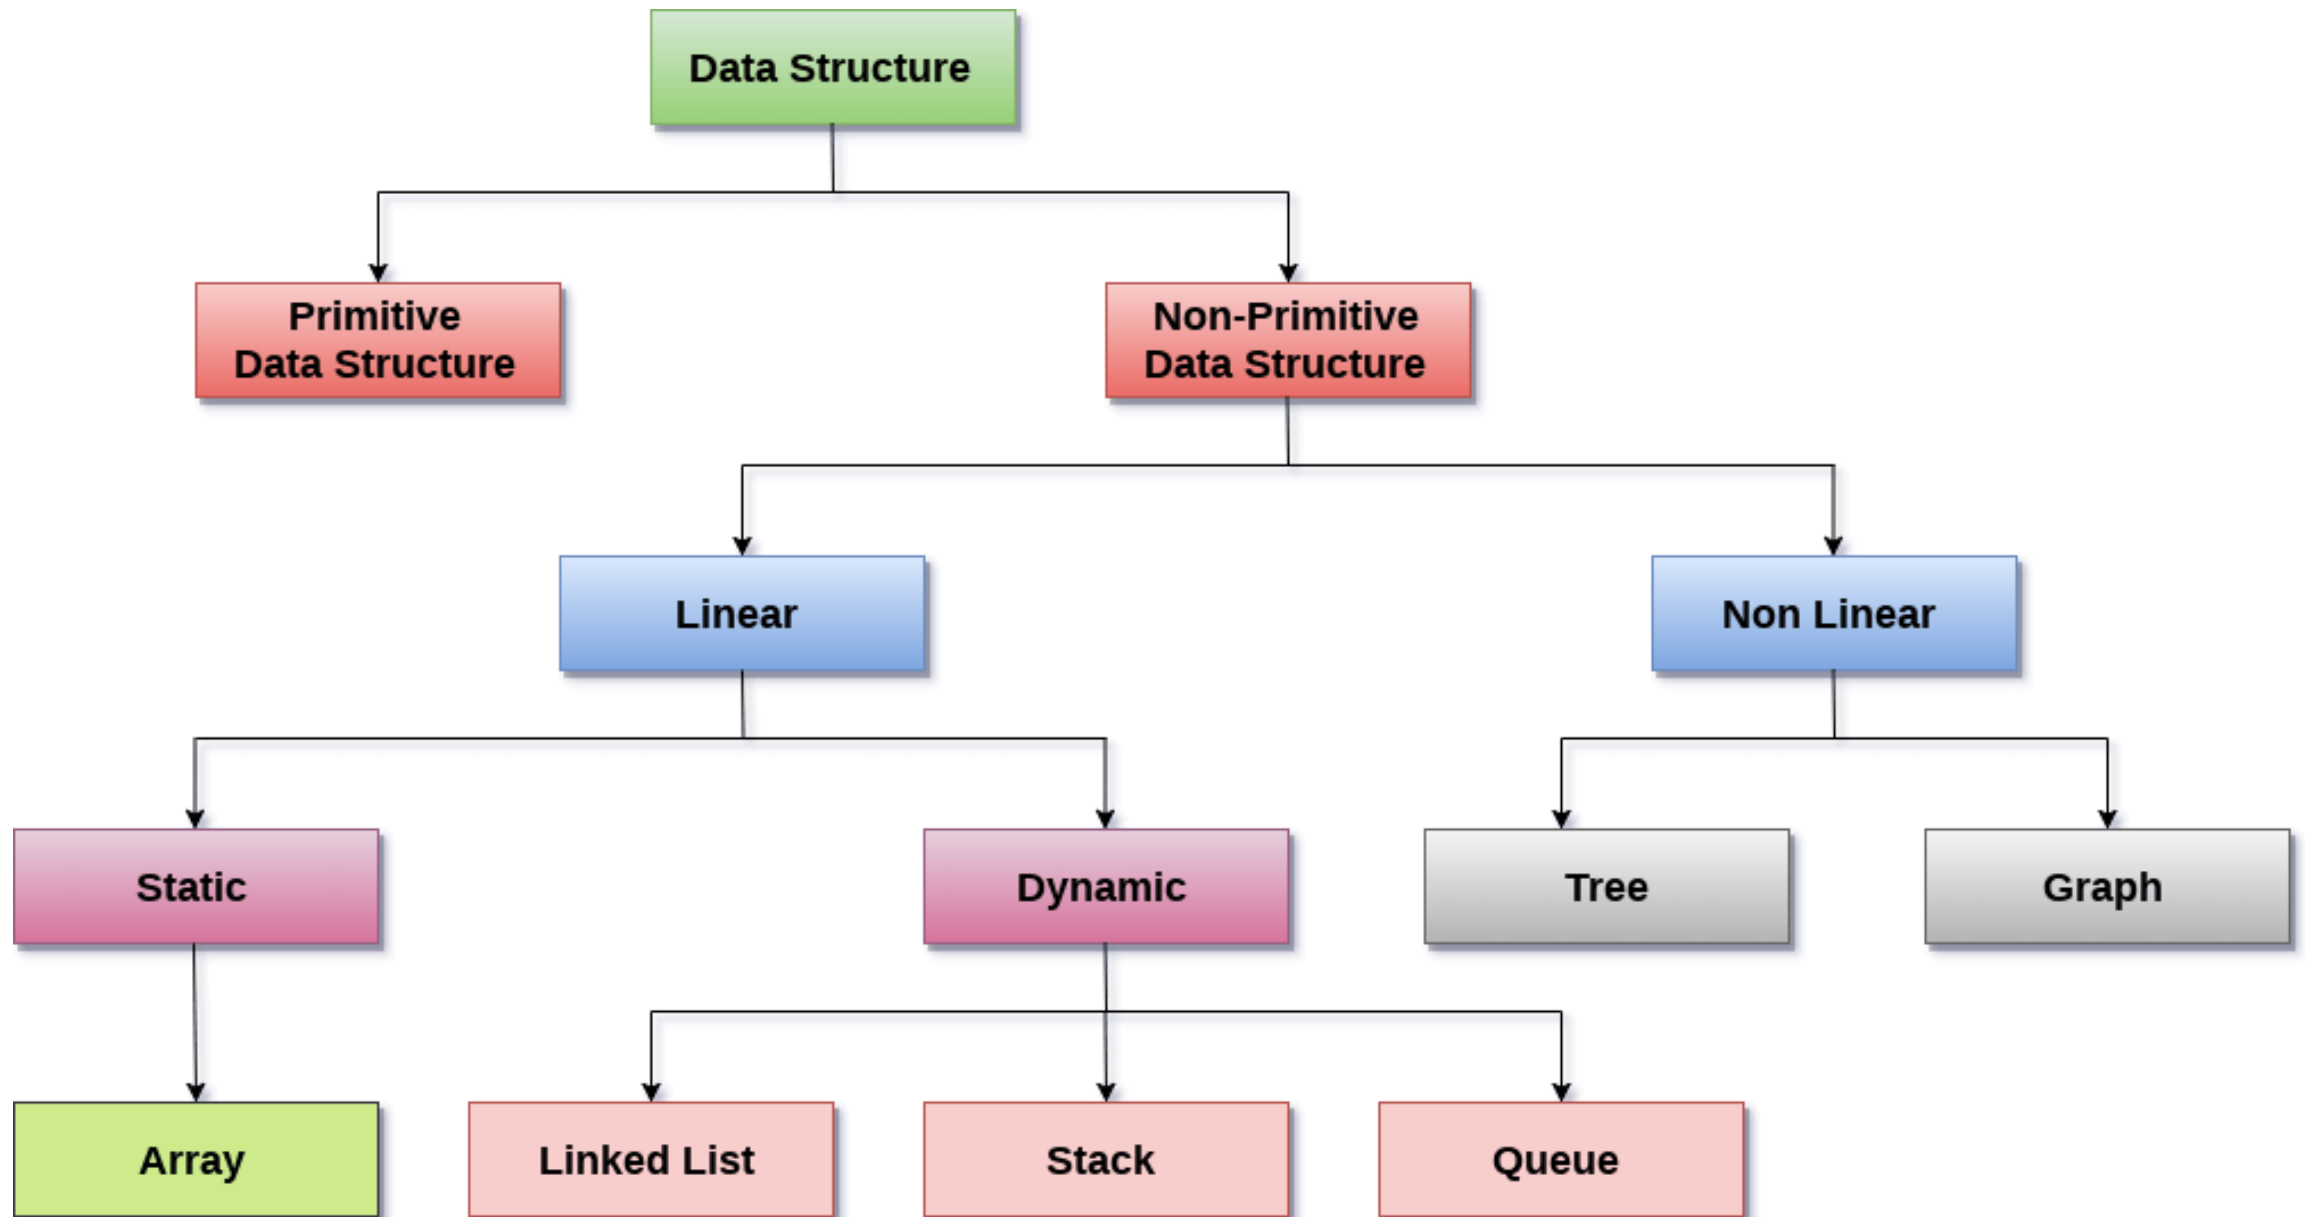
# *Next-Generation Drug Discovery: The Role of AI in Global Pharma Innovation*

*Houman Kazemzadeh, PharmD,  
PhD Student in Medicinal Chemistry,  
Tehran University of Medical Sciences*



# CRISP-DM





# *Computer Aided Drug Design (CADD)*

➤ *Molecular Modeling + Cheminformatics*

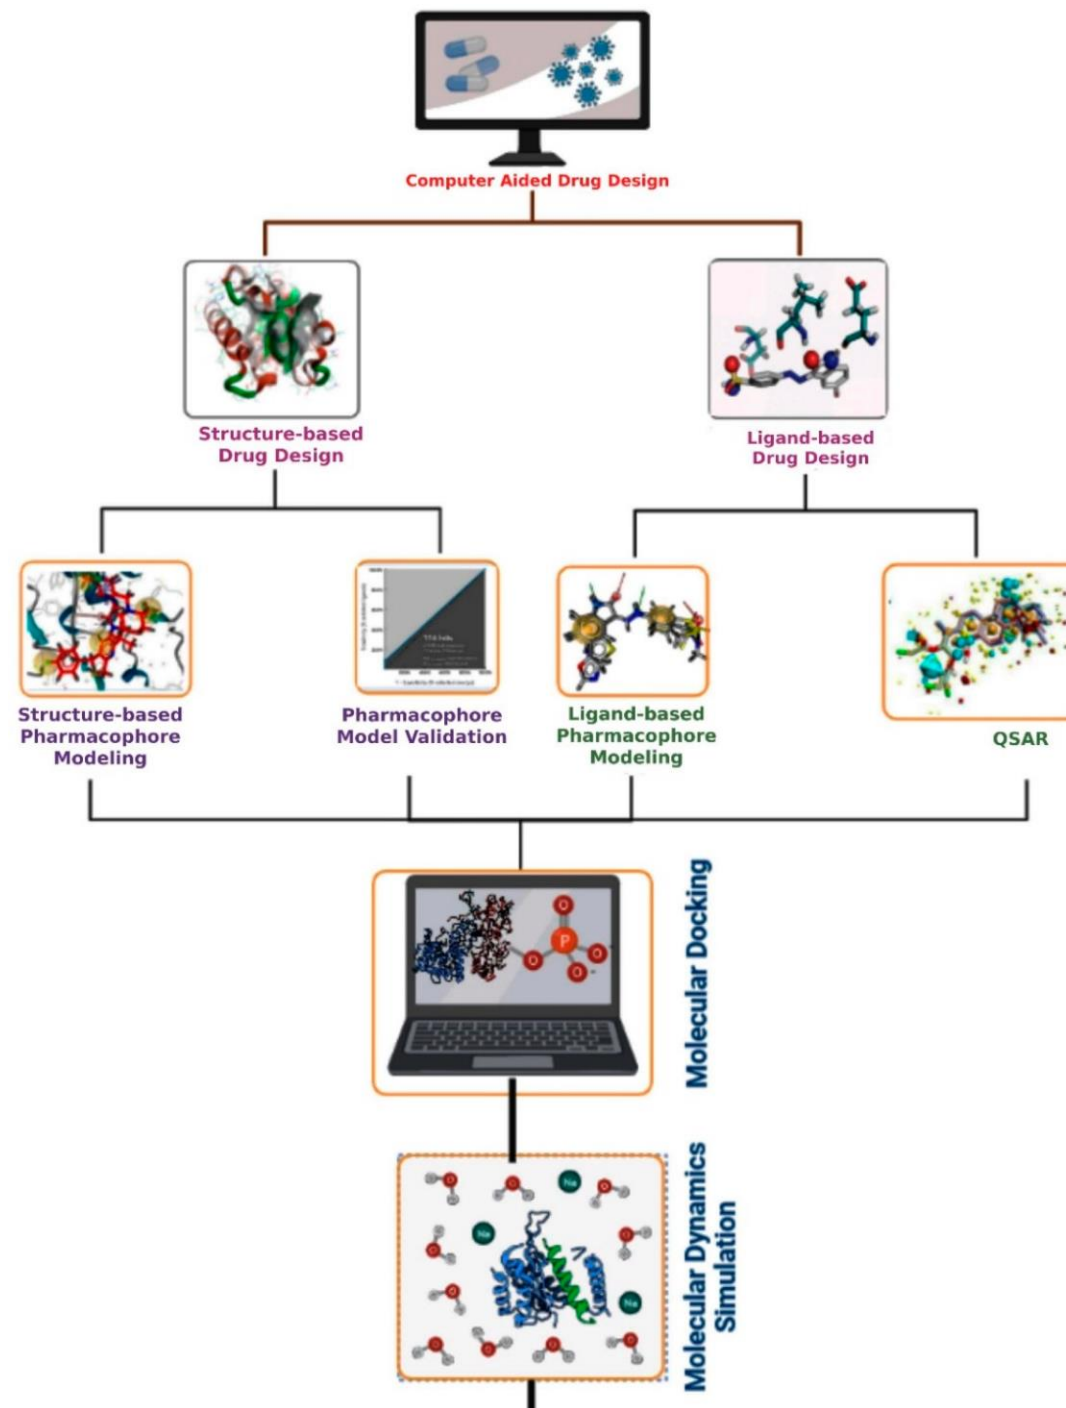
➤ *Pharmacophore Modeling*

➤ *Quantitative Structure-Activity Relationship (QSAR) Analysis*

➤ *Docking Studies*

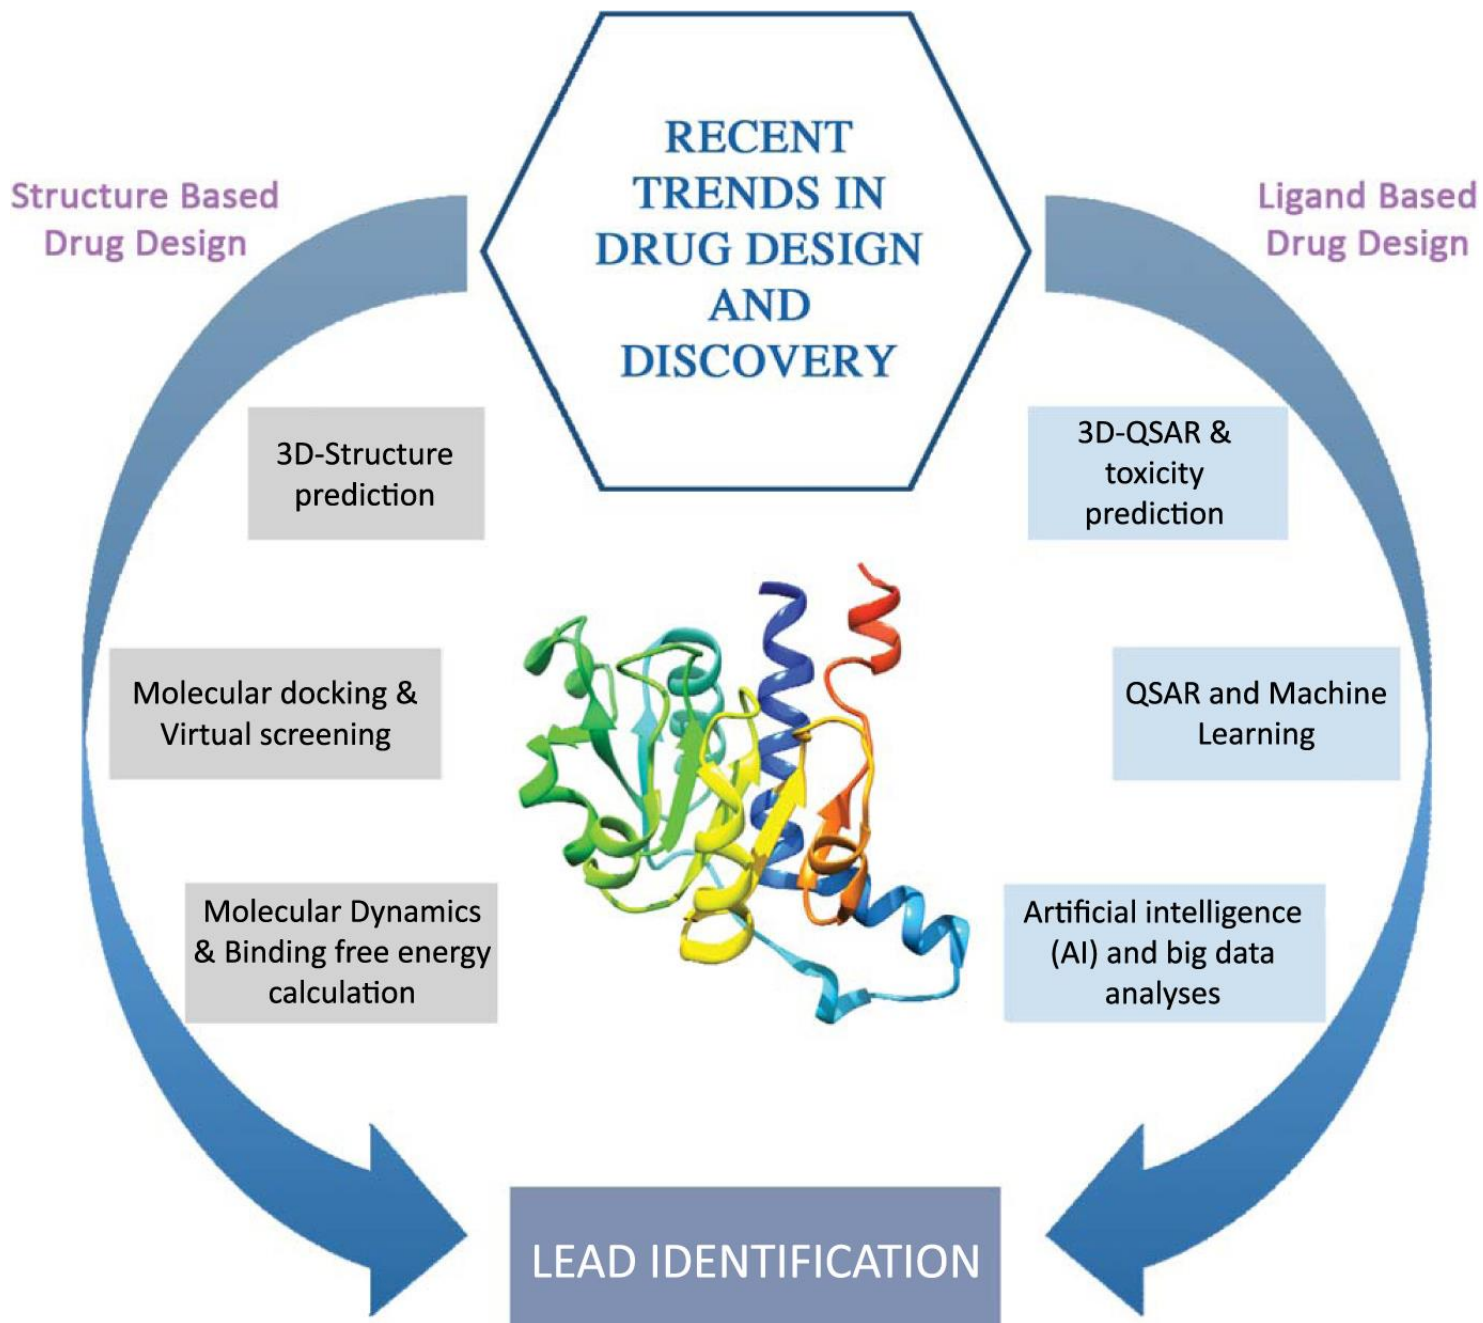
➤ *Molecular Dynamics*

➤ *ADME-Tox Prediction*



Method	Operation	Pros	Cons (Limitations)
<b>2D QSAR</b>	Uses molecular descriptors derived from 2D chemical structure (e.g., logP, molecular weight, atom counts) to model biological activity.	Simple, fast, low computational cost; interpretable descriptors.	Ignores 3D conformation, steric and spatial interactions; oversimplified representation.
<b>3D QSAR</b>	Uses 3D molecular alignment and steric/electrostatic fields to relate spatial features to biological activity (e.g., CoMFA, CoMSIA).	Captures spatial orientation; more predictive than 2D QSAR when alignment is optimal.	Highly sensitive to molecular alignment; fails if conformations are inaccurate; still lacks dynamics.
<b>Ligand-Based Pharmacophore</b>	Identifies common 3D features (e.g., hydrogen bond donor/acceptor, hydrophobic regions) from a set of active ligands to define essential pharmacophoric space.	No need for receptor structure; good for hit expansion; interpretable models.	Assumes similar ligands bind similarly; lacks receptor interaction info; static model.
<b>Structure-Based Pharmacophore</b>	Extracts pharmacophore features directly from ligand-receptor complex or protein binding site (e.g., via docking or crystallography).	Utilizes direct structural information from receptor; context-specific modeling.	Requires accurate receptor structure; limited by static snapshot; ignores molecular dynamics.

# Where is AI






## Software 1.0

Traditional programming




## Software 2.0

Machine learning




# *The Mathematical Background*

- |                              |                                     |                               |
|------------------------------|-------------------------------------|-------------------------------|
| 1. Probability               | 9. Inference algorithms             | 14. State-Space Modeling      |
| 2. Decision Theory           | 10. Message Passing Algorithms      | 15. Latent Space Modeling     |
| 3. Information Theory        | 11. Markov Chain Monte Carlo (MCMC) | 16. Reinforcement Learning    |
| 4. Linear Algebra            | 12. Hamiltonian Monte Carlo (HMC)   | 17. Deep Neural Networks      |
| 5. Calculus                  | 13. Generative Models               | 18. Causality                 |
| 6. Generalized linear models |                                     | 19. Beyond the iid assumption |
| 7. Optimization              |                                     | 20. ...                       |
| 8. Graphs   HyperGraphs      |                                     |                               |




# *Everything is a Graph*



Networks

VS.




Images




Text

## Convolutional




$$\mathbf{h}_i = \phi \left( \mathbf{x}_i, \bigoplus_{j \in \mathcal{N}_i} c_{ij} \psi(\mathbf{x}_j) \right)$$

## Attentional

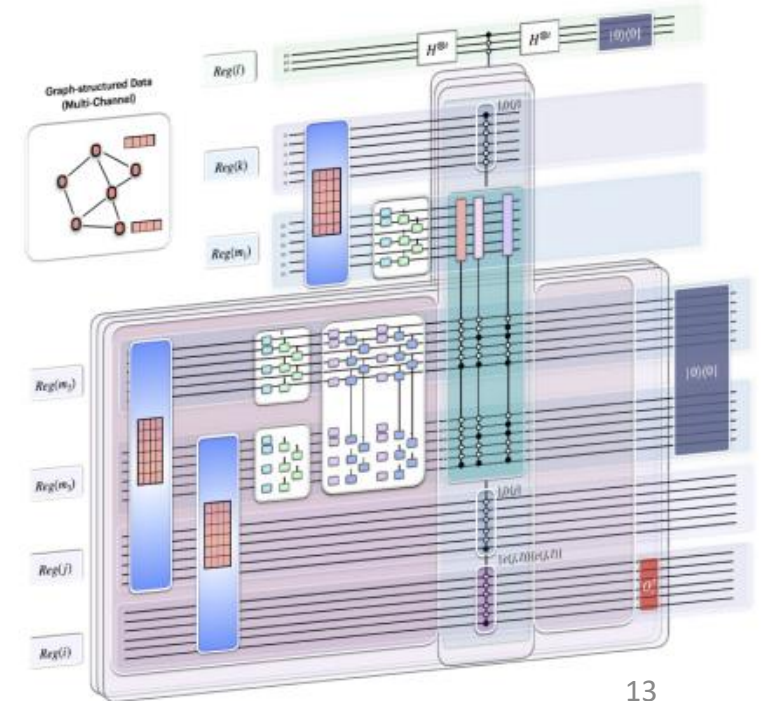




$$\mathbf{h}_i = \phi \left( \mathbf{x}_i, \bigoplus_{j \in \mathcal{N}_i} a(\mathbf{x}_i, \mathbf{x}_j) \psi(\mathbf{x}_j) \right)$$

## Message-Passing



$$\mathbf{h}_i = \phi \left( \mathbf{x}_i, \bigoplus_{j \in \mathcal{N}_i} \psi(\mathbf{x}_i, \mathbf{x}_j) \right)$$







Figure 16.10: Left: a multigraph can have different edge types. Right: a hypergraph can have edges which connect multiple nodes. From [Sanchez-lengeling2021]. Used with kind permission of Benjamin Sanchez-Lengeling.



# Recurrent Neural Networks (RNN)



- 1. Sequential data**
- 2. Order Matters**
- 3. Dependency**
- 4. Memory Needed:** internal hidden state that serves as a memory of previous inputs.
- 5. Time matters:** the same set of weights and biases across all time steps
- 6. Dynamic**



- **Sequential data examples:**

Time Series Data


Text Data

DNA Sequences

Speech Signals

Video Data

**SMILES (Simplified Molecular Input Line Entry System)**



N1CCN(CC1)C(C(F)=C2)=CC(=C2C4=O)N(C3CC3)C=C4C(=O)O



# Generating Focused Molecule Libraries for Drug Discovery with Recurrent Neural Networks

Marwin H. S. Segler,<sup>\*,†</sup> Thierry Kogej,<sup>‡</sup> Christian Tyrchan,<sup>§</sup> and Mark P. Waller<sup>\*,||</sup>

<sup>†</sup>Institute of Organic Chemistry & Center for Multiscale Theory and Computation, Westfälische Wilhelms-Universität Münster, 48149 Münster, Germany


<sup>‡</sup>Hit Discovery, Discovery Sciences, AstraZeneca R&D, Gothenburg, Sweden

<sup>§</sup>Department of Medicinal Chemistry, IMED RIA, AstraZeneca R&D, Gothenburg, Sweden


<sup>||</sup>Department of Physics & International Centre for Quantum and Molecular Structures, Shanghai University, Shanghai, China

## Supporting Information

**ABSTRACT:** In *de novo* drug design, computational strategies are used to generate novel molecules with good affinity to the desired biological target. In this work, we show that recurrent neural networks can be trained as generative models for molecular structures, similar to statistical language models in natural language processing. We demonstrate that the properties of the generated molecules correlate very well with the properties of the molecules used to train the model. In order to enrich libraries with molecules active toward a given biological target, we propose to fine-tune the model with small sets of molecules, which are known to be active against that target. Against *Staphylococcus aureus*, the model reproduced 14% of 6051 hold-out test molecules that medicinal chemists designed, whereas against *Plasmodium falciparum* (Malaria), it reproduced 28% of 1240 test molecules. When coupled with a scoring function, our model can perform the complete *de novo* drug design cycle to generate large sets of novel molecules for drug discovery.



1. Valence Information
2. Implicit Hydrogens
3. Lack of 3D Geometry Information
4. Isomer Representation



5-methyl-4-(2-methylpropyl)-4-(propan-2-yl)hept-1-ene

CC(C)CC(CC=C)(C(C)C)C(C)CC

$$PE_{(pos,2i)} = \sin(pos/10000^{2i/d_{model}})$$

$$PE_{(pos,2i+1)} = \cos(pos/10000^{2i/d_{model}})$$

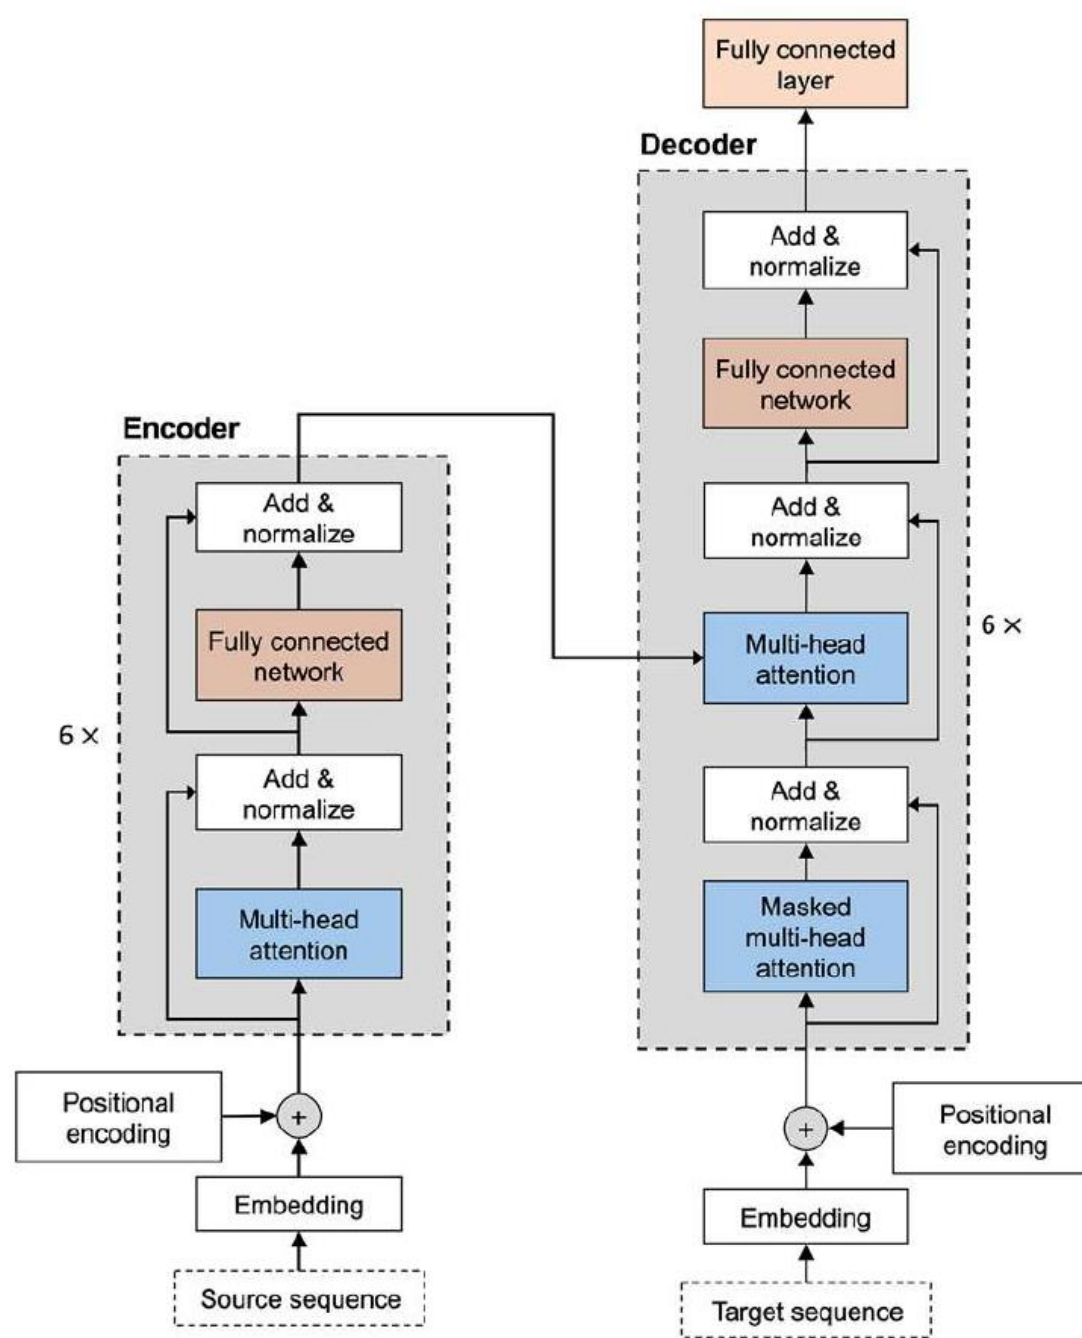


Figure 16.6: The original transformer architecture

- 1. Generative Pre-trained Transformer (GPT)**
- 2. Bidirectional Encoder Representations from Transformers (BERT)**
- 3. Bidirectional and Auto-Regressive Transformer(BART): The best of both worlds**





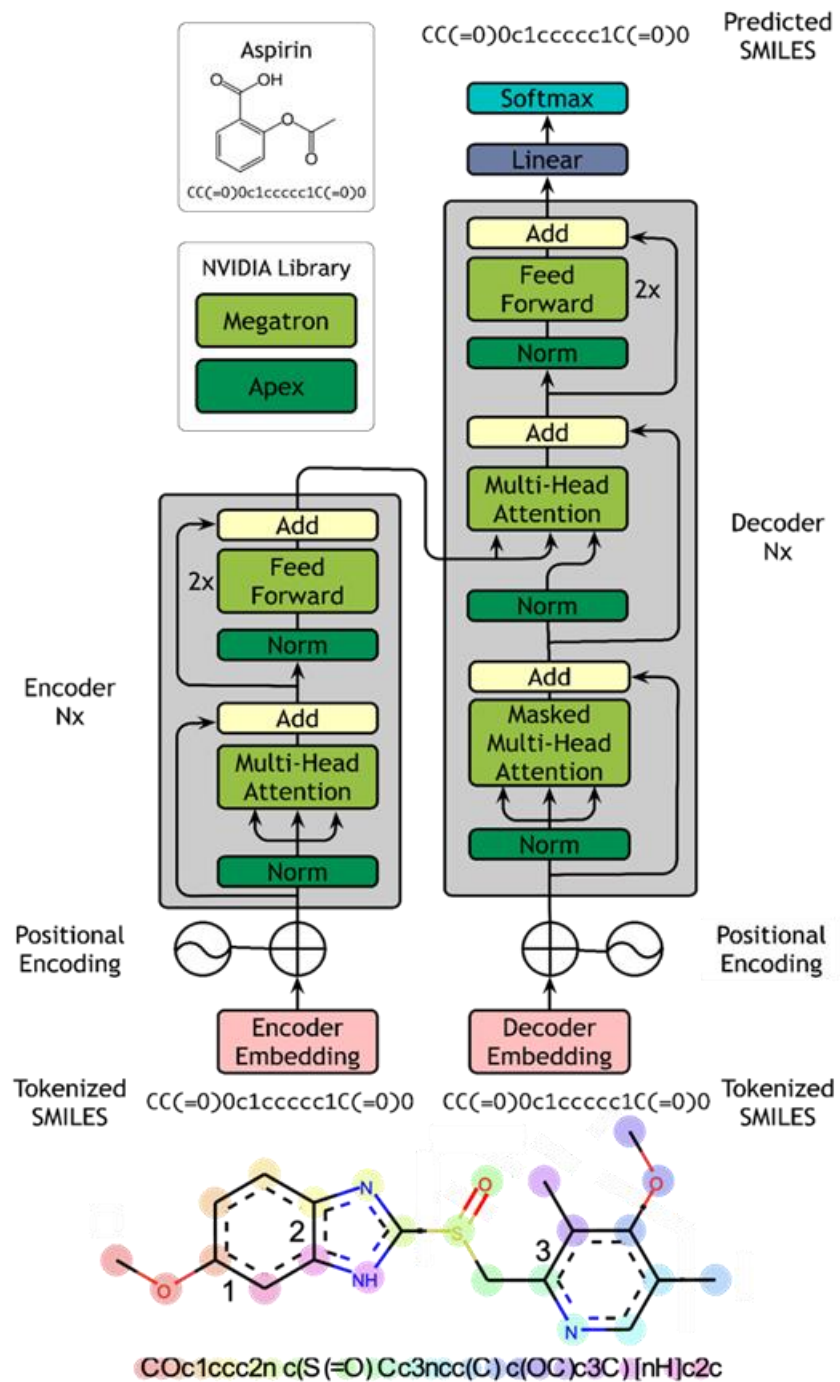
***DrugGPT***



# *MegaMolBART*

- MegaMolBART molecular sequence, based upon known molecular sequences, is an autoencoder trained on small molecules in the form of SMILES that can be used for molecular representation tasks, molecule generation, and retrosynthesis. It was developed using the BioNeMo framework. MegaMolBART has eight layers, four attention heads, a hidden space dimension of 256, and contains 45M parameters. This model is ready for commercial/non-commercial use.







# Graph Neural Networks (GNN)

- **GNNs** have been an area of rapid development in recent years. According to the **State of AI report from 2021**, GNNs have evolved **“from niche to the hottest fields of AI research.”**
- **GNNs** have been applied in a variety of areas, including the following:
  - Text classification (<https://arxiv.org/abs/1710.10903>)
  - Recommender systems (<https://arxiv.org/abs/1704.06803>)
  - Traffic forecasting (<https://arxiv.org/abs/1707.01926>)
  - **Drug Discovery** (<https://arxiv.org/abs/1806.02473>)


Directed graph




Undirected graph




Edge-labeled undirected graph



Node-labeled undirected graph



Node- and edge-labeled undirected graph



# Dynamic vs Static Graphs

## Temporal Nature:

**Static Graphs:** Static graphs represent a **snapshot** of a network or system at a single point in time. They do not capture changes or interactions over time.

**Dynamic Graphs:** **Dynamic graphs explicitly capture changes and interactions over time.** They consist of **multiple snapshots** (timestamps), and edges or nodes can appear, disappear, or change attributes between snapshots.



# Dynamic vs Static Graphs

## Use Cases:

**Static Graphs:** Static graphs are suitable for modeling systems or networks that **do not change or evolve significantly over time**. They are commonly used for social networks, citation networks, and many other applications where the underlying structure remains relatively constant.

**Dynamic Graphs:** Dynamic graphs are used when modeling systems or networks that **exhibit temporal dependencies, where interactions, events, or relationships change over time**. Examples include communication networks, transportation systems, and epidemiological models.

# Dynamic vs Static Graphs

## Representation:

**Static Graphs:** Static graphs are typically represented using a **single adjacency matrix**, where each entry represents the presence or absence of an edge between two nodes. Node and edge attributes are **constant**.

**Dynamic Graphs:** Dynamic graphs are represented as a **sequence of static graphs, each associated with a specific timestamp**. Edges and nodes can have associated timestamps and attributes that **evolve over time**.

# Dynamic vs Static Graphs

## Analytical Challenges:

**Static Graphs:** Analyzing static graphs is often **simpler**, as they do not involve temporal dynamics. Traditional graph algorithms and metrics are commonly applied.

**Dynamic Graphs:** Analyzing dynamic graphs can be more complex due to the need to consider temporal aspects. Researchers use **specialized algorithms** for tasks like tracking node or edge changes, detecting patterns over time, and predicting future states.




# Dynamic vs Static Graphs


## Storage and Processing:

**Static Graphs:** Storing and processing static graphs are often more **straightforward** since the graph structure remains constant.

**Dynamic Graphs:** Handling dynamic graphs requires **more advanced data structures** and algorithms to efficiently manage changes over time.




Friend graph



Molecular graph  
of caffeine

# Caffeine molecule



$$A = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Adjacency matrix

$$X =$$

Node label matrix

One-hot encoding of the atom type

$$\begin{bmatrix} \text{C} & \text{N} & \text{O} \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

$$X_E =$$

Edge label matrix

One-hot encoding of the bond type


$$\begin{array}{cc} \text{SB} & \text{DB} \\ (1,2) & \begin{bmatrix} 1 & 0 \end{bmatrix} \\ (2,3) & \begin{bmatrix} 1 & 0 \end{bmatrix} \\ (2,6) & \begin{bmatrix} 1 & 0 \end{bmatrix} \\ (3,4) & \begin{bmatrix} 0 & 1 \end{bmatrix} \\ (4,5) & \begin{bmatrix} 1 & 0 \end{bmatrix} \\ (5,6) & \begin{bmatrix} 0 & 1 \end{bmatrix} \\ (5,12) & \begin{bmatrix} 1 & 0 \end{bmatrix} \\ (6,7) & \begin{bmatrix} 1 & 0 \end{bmatrix} \\ (7,8) & \begin{bmatrix} 0 & 1 \end{bmatrix} \\ (7,9) & \begin{bmatrix} 1 & 0 \end{bmatrix} \\ (9,10) & \begin{bmatrix} 1 & 0 \end{bmatrix} \\ (9,14) & \begin{bmatrix} 1 & 0 \end{bmatrix} \\ (10,11) & \begin{bmatrix} 0 & 1 \end{bmatrix} \\ (10,12) & \begin{bmatrix} 1 & 0 \end{bmatrix} \\ (12,13) & \begin{bmatrix} 1 & 0 \end{bmatrix} \end{array}$$

Product C:


```
[[[ 57  21  48   0]
  [ 71  27  60   0]
  [111  51  96   0]]
```

```
[[ 66  16  36  18]
 [ 95  22  58  29]
 [ 83  20  46  23]]
```


```
[[ 44  34  72  44]
 [ 34  32  72  34]
 [ 13  25  64  13]]], shape=(3, 3, 4)
```




3D Array




shape : (4, 3, 2)



Node-level embedding




Edge-level embedding




Graph-level embedding


# Dynamic Graph

## Molecular Graph




## Reconstruction / Generation





**CNN**

**In Euclidean Space**




**GNN**

**In Non-Euclidean Space**





Surfaces



Distributions



Graphs / Networks




Functions on Manifolds




Hyperbolic spaces



Hyper-surfaces

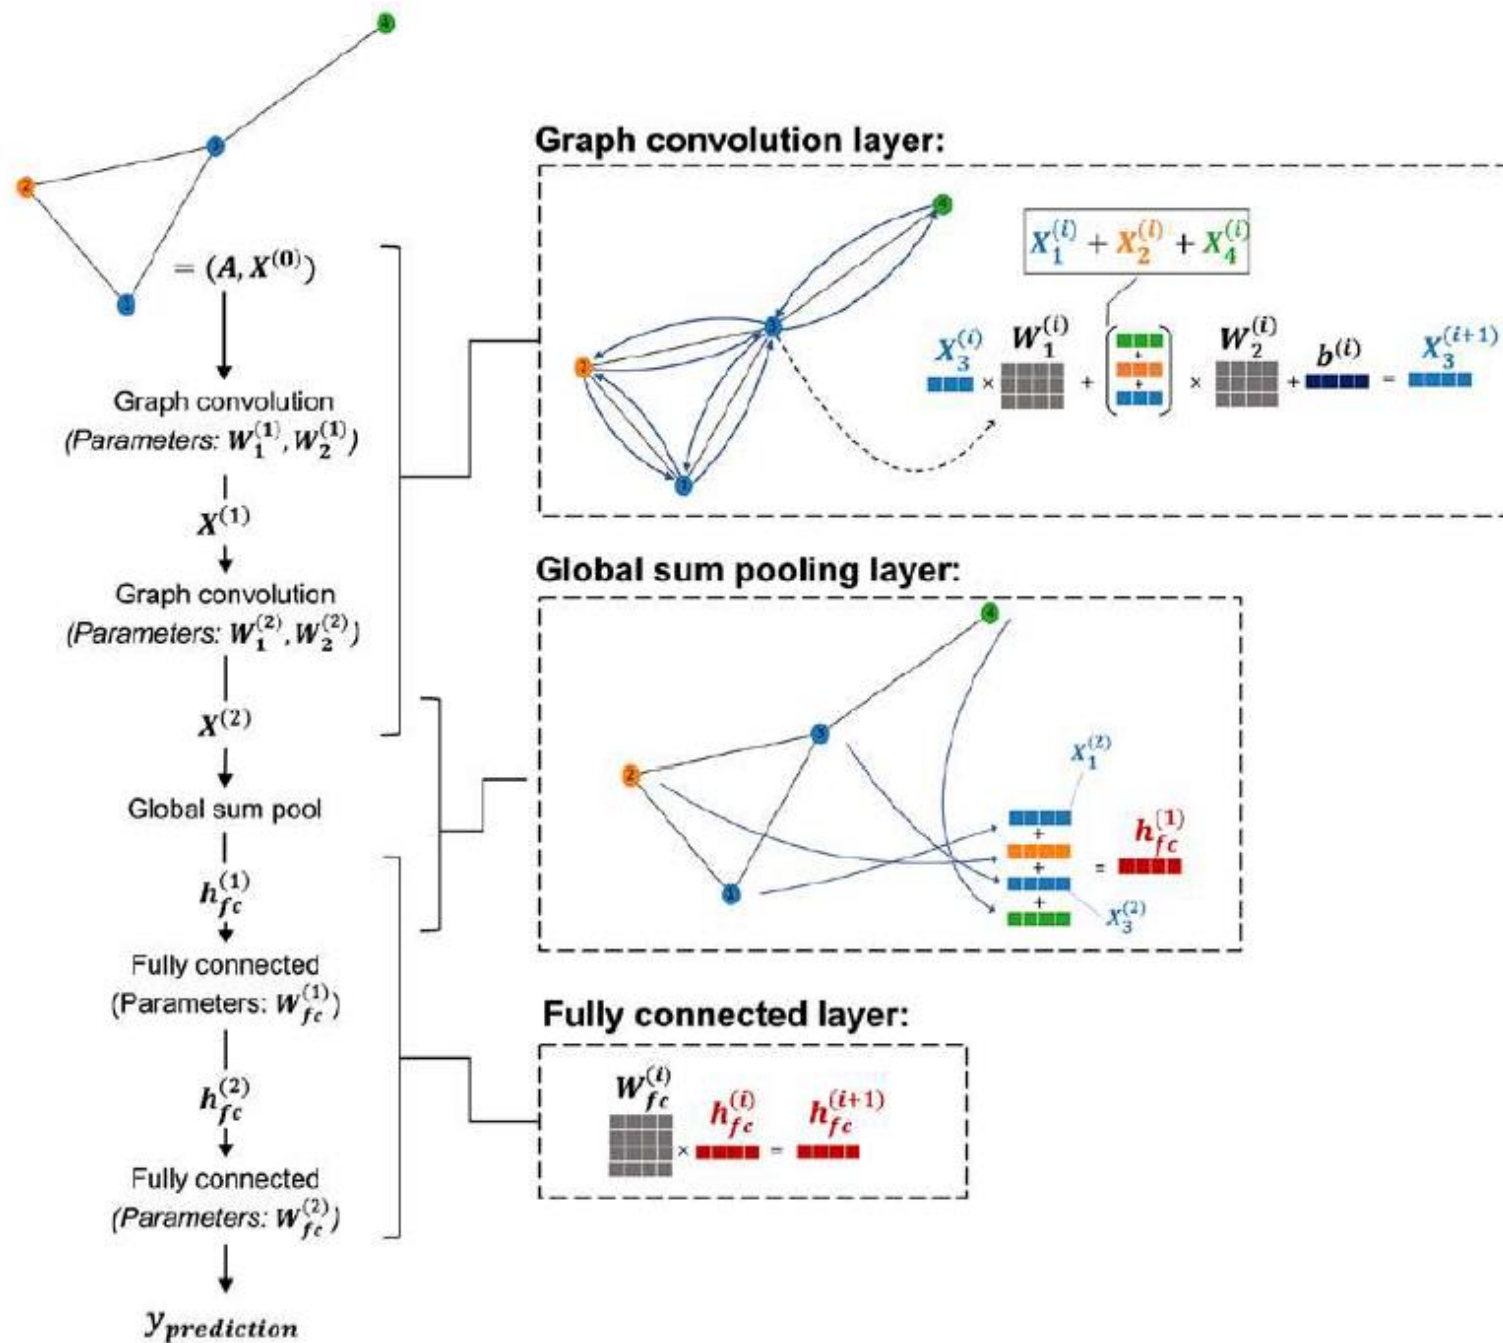


Molecules



General manifolds





## Computer Science &gt; Machine Learning

[Submitted on 4 Apr 2017 (v1), last revised 12 Jun 2017 (this version, v2)]

# Neural Message Passing for Quantum Chemistry

Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, George E. Dahl

Supervised learning on molecules has incredible potential to be useful in chemistry, drug discovery, and materials science. Luckily, several promising and closely related neural network models invariant to molecular symmetries have already been described in the literature. These models learn a message passing algorithm and aggregation procedure to compute a function of their entire input graph. At this point, the next step is to find a particularly effective variant of this general approach and apply it to chemical prediction benchmarks until we either solve them or reach the limits of the approach. In this paper, we reformulate existing models into a single common framework we call Message Passing Neural Networks (MPNNs) and explore additional novel variations within this framework. Using MPNNs we demonstrate state of the art results on an important molecular property prediction benchmark; these results are strong enough that we believe future work should focus on datasets with larger molecules or more accurate ground truth labels.

Comments: 14 pages

Subjects: **Machine Learning (cs.LG)**

ACM classes: I.2.6

Cite as: arXiv:1704.01212 [cs.LG]

(or arXiv:1704.01212v2 [cs.LG] for this version)

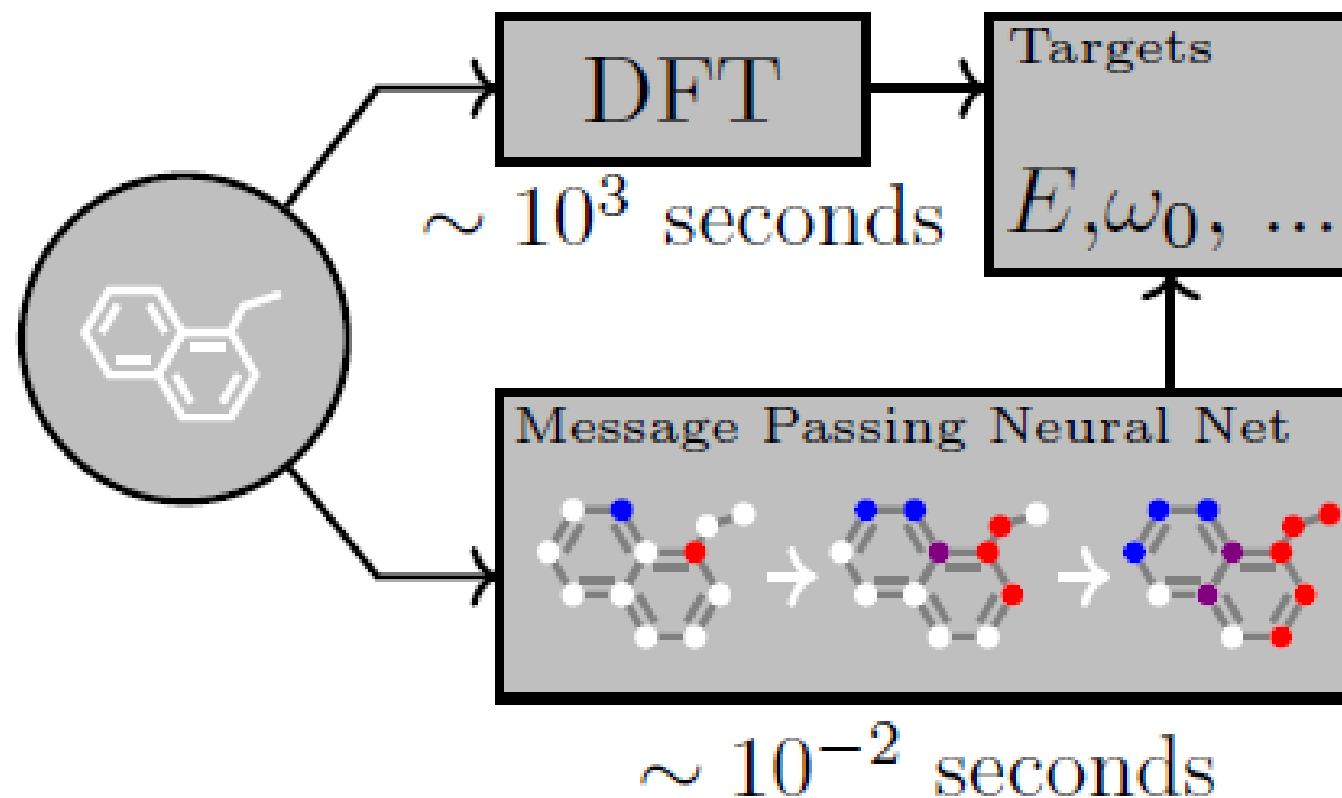
<https://doi.org/10.48550/arXiv.1704.01212> 

## Submission history

From: Justin Gilmer [\[view email\]](#)

[\[v1\]](#) Tue, 4 Apr 2017 23:00:44 UTC (140 KB)

[\[v2\]](#) Mon, 12 Jun 2017 20:52:56 UTC (118 KB)



*Figure 1.* A Message Passing Neural Network predicts quantum properties of an organic molecule by modeling a computationally expensive DFT calculation.



Open-Source Cheminformatics  
and Machine Learning

# RDKit

1. *Molecule Creation*
2. *Property Calculation*
3. *Chemical Reactions*
4. *Molecular fingerprints*
5. *3D Conformer Generation*
6. *QSAR modeling*
7. *Substructure Searching*
8. *Force Field Optimization*
9. *Pharmacophore Modeling*
10. *Constrained Embedding*
11. *Shape and Volume Calculations*
12. *...*

# RDKit Chem Module : Atom-Level Functions

- **atom.GetSymbol()**: Returns the atomic symbol (e.g., "C", "O").
- **atom.GetChiralTag()**: Retrieves the atom's chirality (R/S or achiral).
- **atom.GetTotalDegree()**: Gets the total degree of the atom, including bonds to hydrogen.
- **atom.GetFormalCharge()**: Provides the atom's formal charge.
- **atom.GetTotalNumHs()**: Counts explicit and implicit hydrogens on the atom.
- **atom.GetNumRadicalElectrons()**: Returns the number of radical electrons.
- **atom.GetHybridization()**: Retrieves the atom's hybridization (e.g., sp2, sp3).
- **atom.GetIsAromatic()**: Checks if the atom is aromatic.
- **atom.IsInRing()**: Determines if the atom is part of a ring.
- **atom.GetOwningMol()**: Retrieves the molecule object that owns the atom.






# RDKit Chem Module : Bond-Level Functions

- **bond.GetBondType()**: Gets the bond type (single, double, triple, aromatic).
- **bond.GetBondDir()**: Retrieves the bond direction (e.g., BEGINWEDGE).
- **bond.GetStereo()**: Provides stereochemistry information (cis/trans).
- **bond.GetIsConjugated()**: Checks if the bond is conjugated.
- **bond.IsInRing()**: Determines if the bond is part of a ring.
- **bond.GetBeginAtomIdx()** and **bond.GetEndAtomIdx()**: Fetch the indices of the bonded atoms.
- **bond.GetOwningMol()**: Retrieves the molecule object that owns the bond.



# RDKit Chem Module : Molecule-Level Functions

- **mol.GetConformer()**: Retrieves the conformer for accessing 3D positions.
- **mol.GetNumConformers()**: Checks if the molecule has any conformers.
- **mol.Compute2DCoords()**: Generates 2D coordinates for the molecule.





# Unified Graph: N-Dimensional Non-Linear Probabilistic Graph Model



Model	Pros	Cons / Limitations
<b>GCN</b> (Graph Convolutional Network)	Fast, scalable; good for basic property prediction.	No edge/bond awareness; poor long-range expressiveness; over-smoothing in deep layers.
<b>*MPNN</b> (Message Passing Neural Network)	Accurate in molecular tasks; good chemical realism.	Higher computational cost; complex architecture tuning.
<b>GAT</b> (Graph Attention Network)	Interpretable; effective on heterogeneous graphs.	Attention overhead increases complexity; limited scalability to large graphs.
<b>*RGCN</b> (Relational GCN)	Ideal for molecular graphs with heterogeneous bonds.	Model size grows with relation types; overfitting risk in small datasets.
<b>GIN</b> (Graph Isomorphism Network)	Strong representational capacity; good for classification.	Ignores edge types; overfitting possible on small data; less interpretable.
<b>GraphSAGE</b>	Works well on large molecular graphs; generalizable to unseen data.	Less sensitive to chemical structure; weaker granularity in molecular context.
<b>*D-MPNN</b> (Directed MPNN)	More chemically accurate; great for reaction/property prediction.	More complex implementation; requires directed graph data.
<b>ChebNet</b>	Efficient and deeper local context modeling.	Less intuitive than spatial GCNs; not widely adopted in cheminformatics.



*[Submitted on 8 Mar 2018]*

## Learning Deep Generative Models of Graphs

Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, Peter Battaglia

Graphs are fundamental data structures which concisely capture the relational structure in many important real-world domains, such as knowledge graphs, physical and social interactions, language, and chemistry. Here we introduce a powerful new approach for learning generative models over graphs, which can capture both their structure and attributes. Our approach uses graph neural networks to express probabilistic dependencies among a graph's nodes and edges, and can, in principle, learn distributions over any arbitrary graph. In a series of experiments our results show that once trained, our models can generate good quality samples of both synthetic graphs as well as real molecular graphs, both unconditionally and conditioned on data. Compared to baselines that do not use graph-structured representations, our models often perform far better. We also explore key challenges of learning generative models of graphs, such as how to handle symmetries and ordering of elements during the graph generation process, and offer possible solutions. Our work is the first and most general approach for learning generative models over arbitrary graphs, and opens new directions for moving away from restrictions of vector- and sequence-like knowledge representations, toward more expressive and flexible relational data structures.

Comments: 21 pages

Subjects: **Machine Learning (cs.LG)**; Machine Learning (stat.ML)Cite as: [arXiv:1803.03324](#) [cs.LG](or [arXiv:1803.03324v1](#) [cs.LG] for this version)<https://doi.org/10.48550/arXiv.1803.03324> 

### Submission history

From: Yujia Li [[view email](#)]**[v1]** Thu, 8 Mar 2018 22:20:00 UTC (1,145 KB)

[Submitted on 24 May 2022 (v1), last revised 30 Jan 2023 (this version, v3)]

# Graph Neural Networks Intersect Probabilistic Graphical Models: A Survey

Chenqing Hua, Sitao Luan, Qian Zhang, Jie Fu

Graphs are a powerful data structure to represent relational data and are widely used to describe complex real-world data structures. Probabilistic Graphical Models (PGMs) have been well-developed in the past years to mathematically model real-world scenarios in compact graphical representations of distributions of variables. Graph Neural Networks (GNNs) are new inference methods developed in recent years and are attracting growing attention due to their effectiveness and flexibility in solving inference and learning problems over graph-structured data. These two powerful approaches have different advantages in capturing relations from observations and how they conduct message passing, and they can benefit each other in various tasks. In this survey, we broadly study the intersection of GNNs and PGMs. Specifically, we first discuss how GNNs can benefit from learning structured representations in PGMs, generate explainable predictions by PGMs, and how PGMs can infer object relationships. Then we discuss how GNNs are implemented in PGMs for more efficient inference and structure learning. In the end, we summarize the benchmark datasets used in recent studies and discuss promising future directions.

Subjects: **Artificial Intelligence (cs.AI)**; Machine Learning (cs.LG)

Cite as: [arXiv:2206.06089 \[cs.AI\]](#)

(or [arXiv:2206.06089v3 \[cs.AI\]](#) for this version)

<https://doi.org/10.48550/arXiv.2206.06089> 

## Submission history

From: Chenqing Hua [[view email](#)]

[\[v1\]](#) Tue, 24 May 2022 03:36:25 UTC (3,376 KB)

[\[v2\]](#) Fri, 18 Nov 2022 04:05:34 UTC (3,376 KB)

[\[v3\]](#) Mon, 30 Jan 2023 10:47:31 UTC (3,376 KB)

Encoder

Latent Space Optimization

Decoder

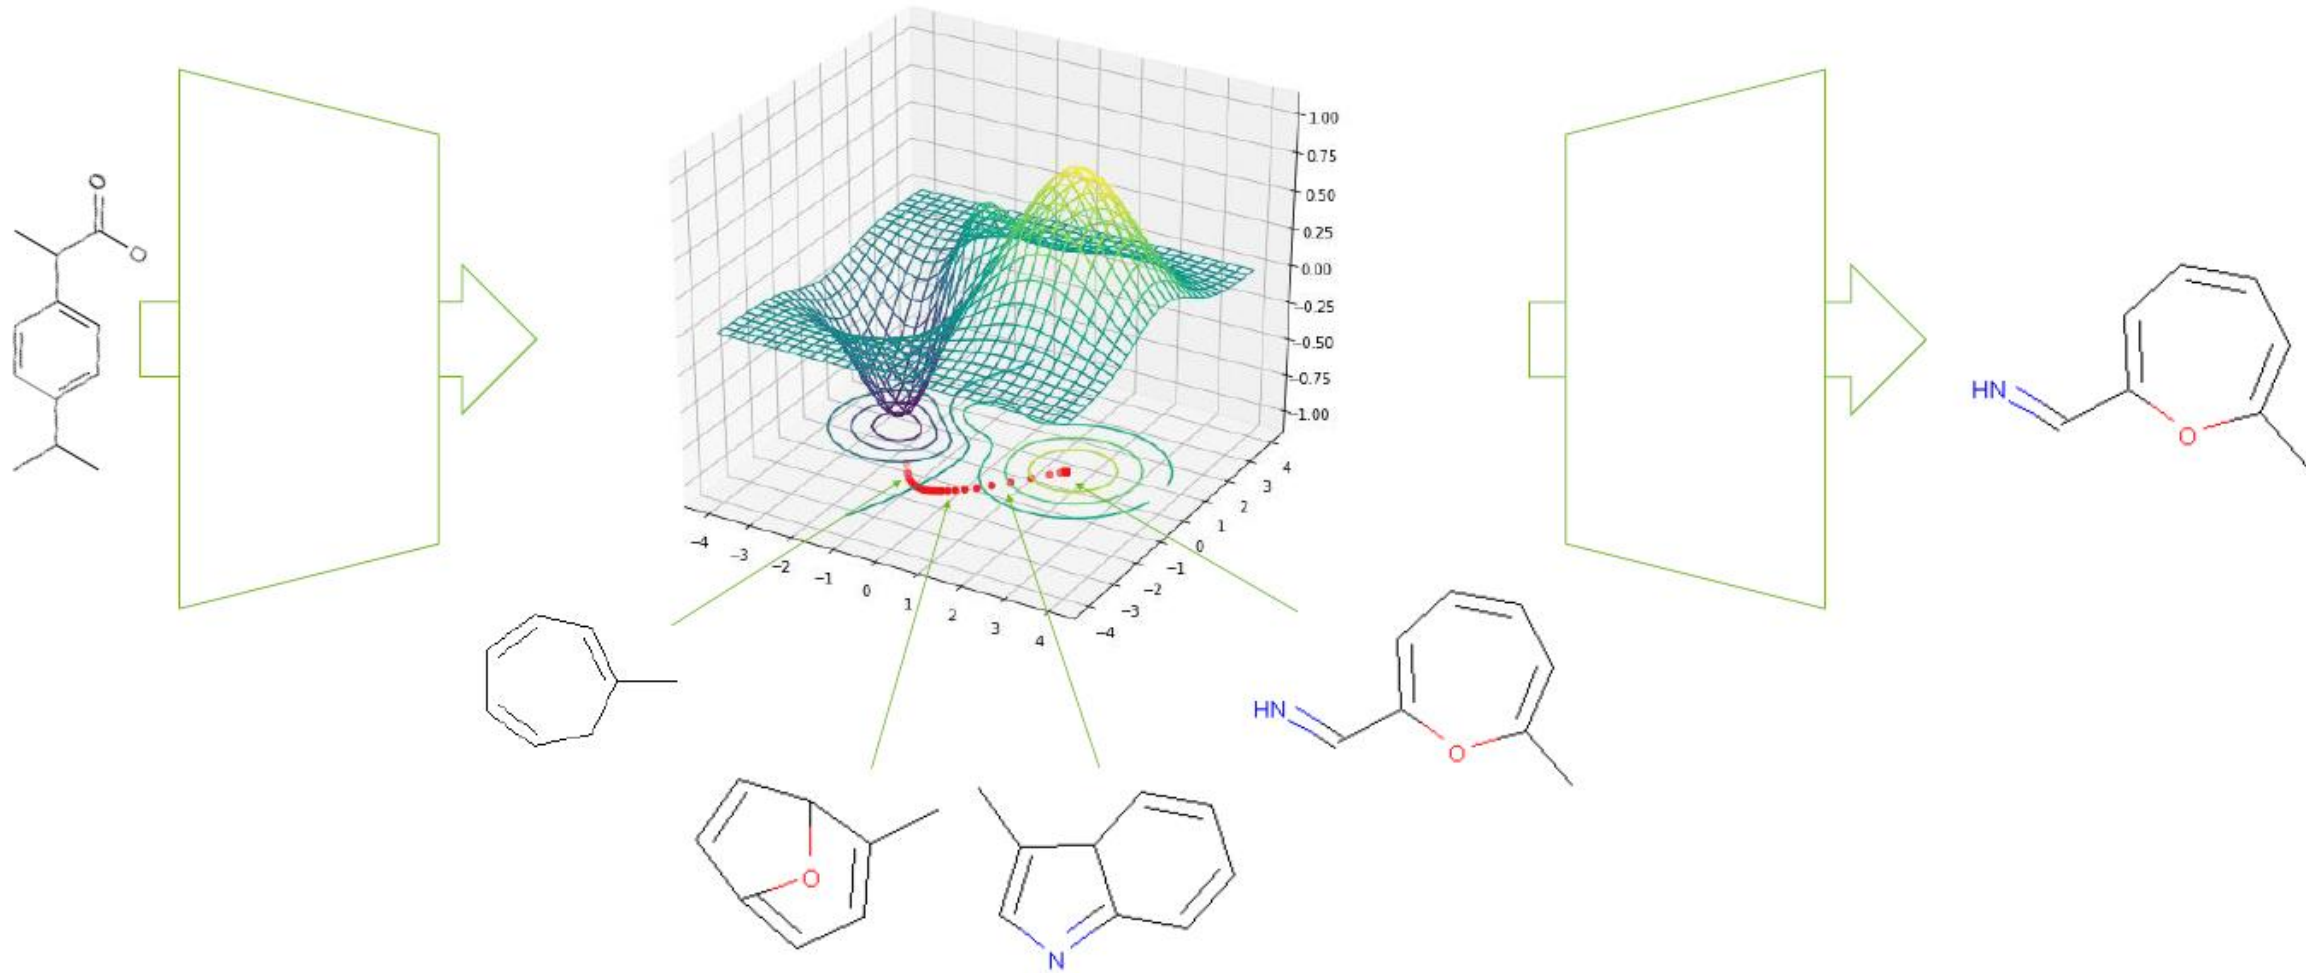


Table 5. Molecular Generation and Reconstruction Performance in ZINC250k

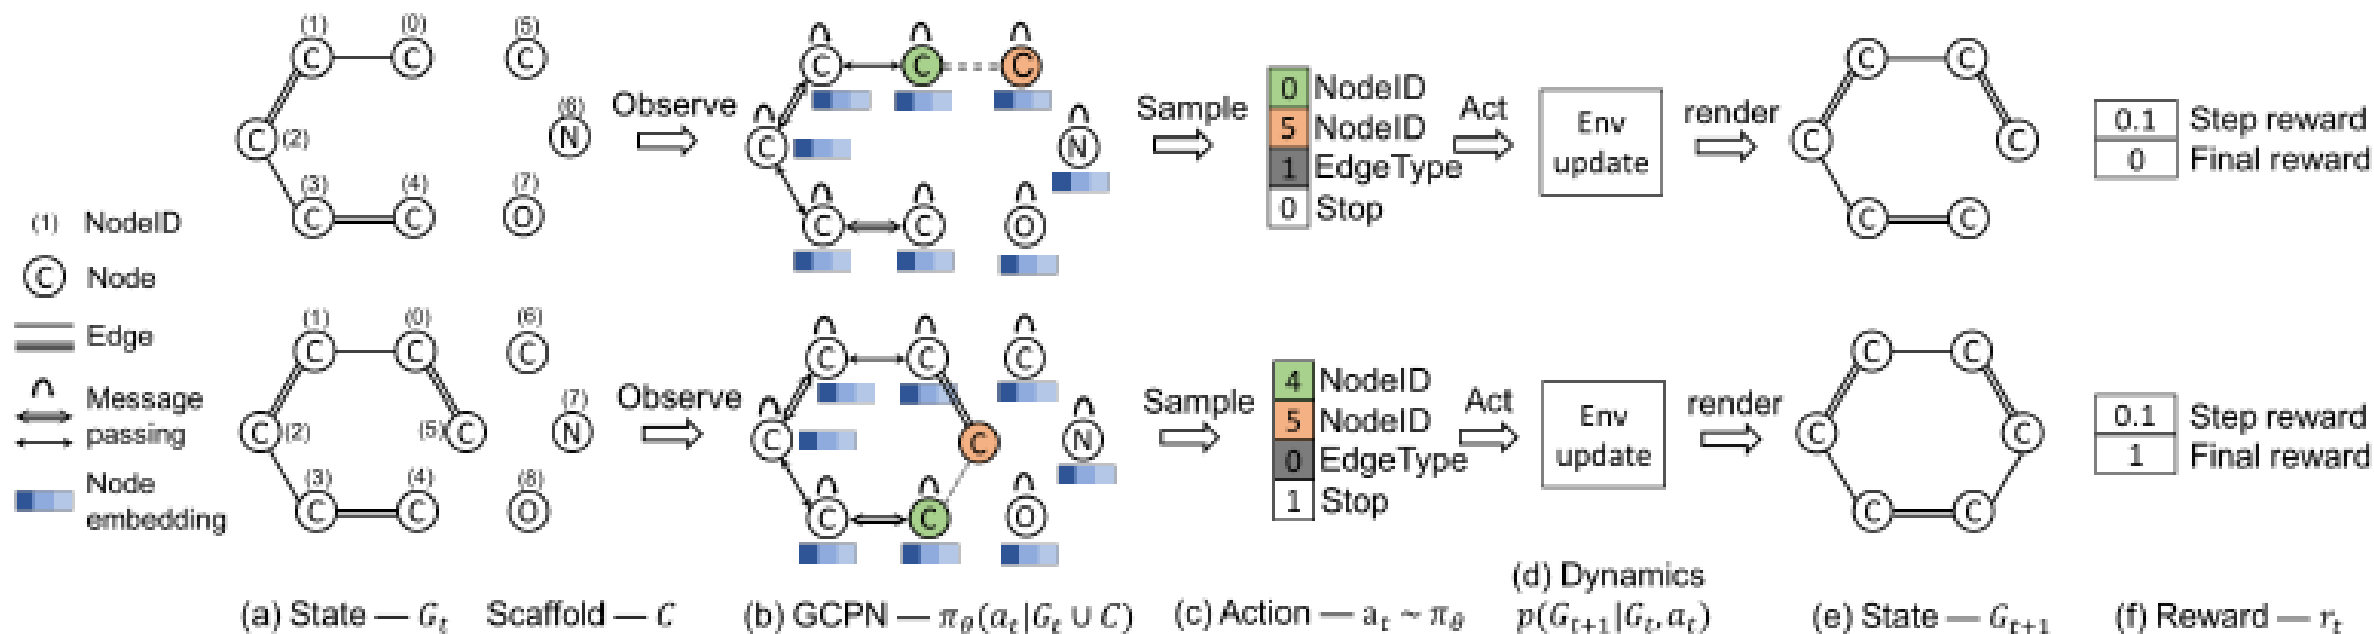
	%Validity	%Validity w/o check	%Novelty	%Uniqueness	%Reconstruction
SDVAE <sup>c</sup>	41.40	41.40	<b>100</b>	<b>100</b>	76.22
CGVAE <sup>b</sup>	<b>100</b>	n/a	<b>100</b>	99.82	n/a
JT-VAE <sup>c</sup>	<b>100</b>	n/a	<b>100</b>	99.96	76.54
GraphNVP <sup>a</sup>	42.60	42.60	<b>100</b>	94.80	<b>100</b>
GraphAF <sup>c</sup>	<b>100</b>	71.40	<b>100</b>	99.10	<b>100</b>
MoFlow <sup>c</sup>	<b>100</b>	45.61	<b>100</b>	99.92	<b>100</b>
GraphDF <sup>b</sup>	<b>100</b>	<b>89.03</b>	<b>100</b>	99.16	<b>100</b>
GCPN <sup>c</sup>	<b>100</b>	21.04	<b>100</b>	99.93	n/a
MRNN <sup>a</sup>	<b>100</b>	65.00	<b>100</b>	99.89	n/a

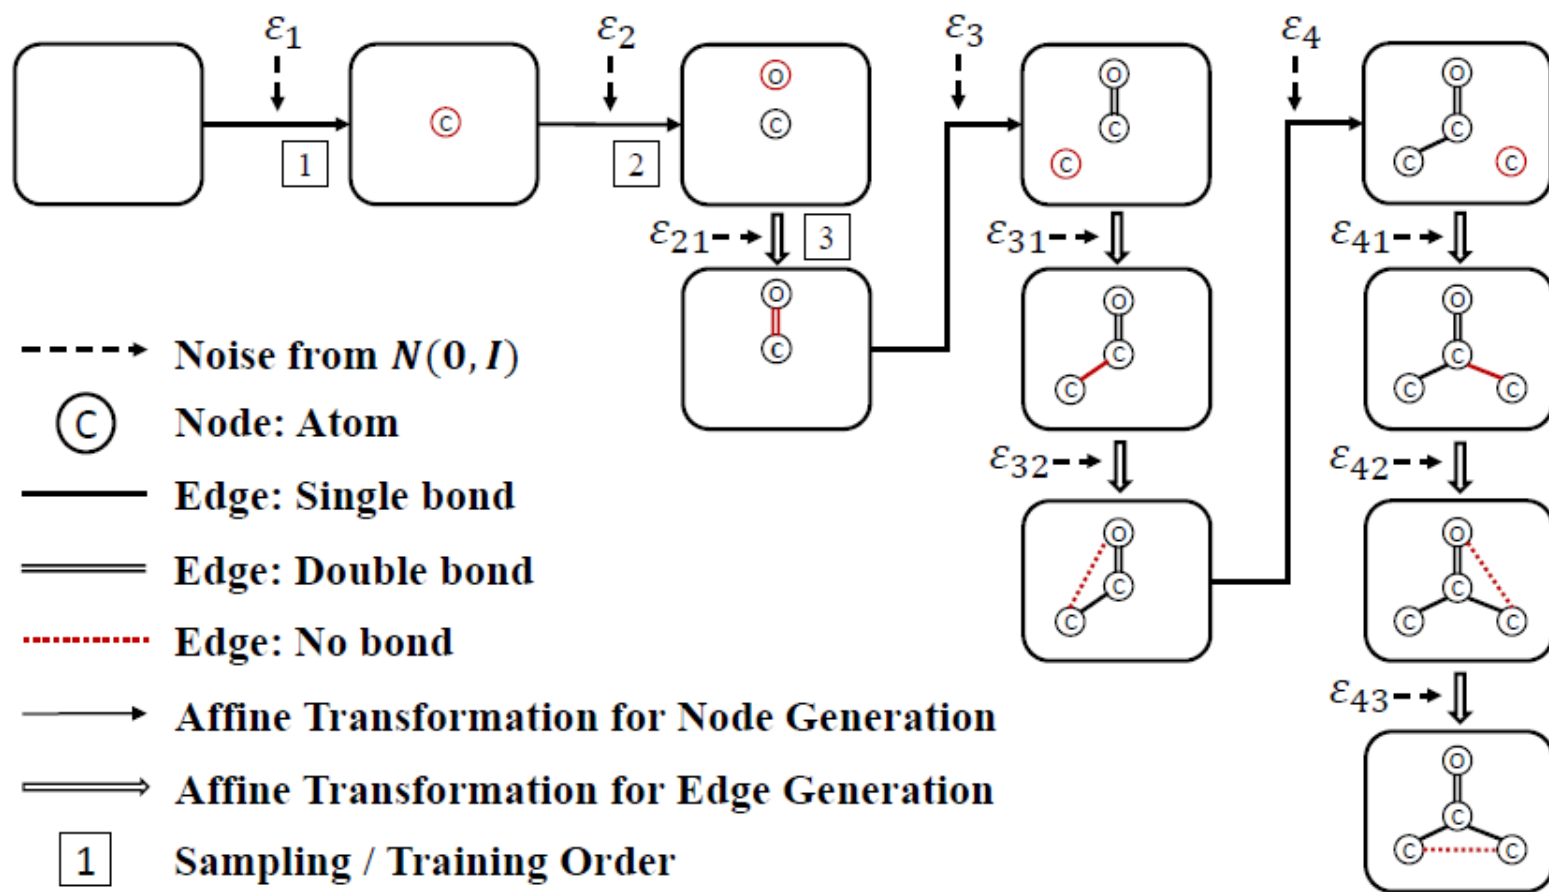
<sup>a</sup>Data is cited from MoFlow.<sup>106</sup> <sup>b</sup>Data is cited from the corresponding papers. <sup>c</sup>Data is obtained by running its official source code.

Table 6. Molecular Generation and Reconstruction Performance in QM9

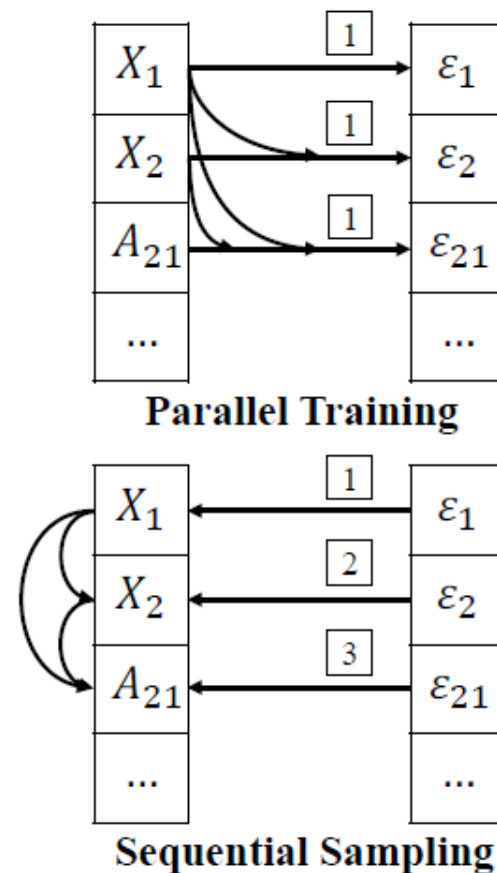
	%Validity	%Validity w/o check	%Novelty	%Uniqueness	%Reconstruction
CGVAE <sup>a</sup>	<b>100</b>	n/a	94.35	98.57	n/a
JT-VAE <sup>b</sup>	99.86	n/a	<b>100</b>	96.32	68.53
GraphNVP <sup>b</sup>	50.86	50.86	88.46	97.52	<b>100</b>
GraphAF <sup>b</sup>	<b>100</b>	46.30	91.54	99.15	<b>100</b>
MoFlow <sup>b</sup>	<b>100</b>	81.14	97.3	<b>99.26</b>	<b>100</b>
GraphDF <sup>a</sup>	<b>100</b>	<b>82.67</b>	98.1	97.62	<b>100</b>
GCPN <sup>b</sup>	<b>100</b>	18.23	<b>100</b>	87.13	n/a

<sup>a</sup>Data is cited from the corresponding papers. <sup>b</sup>Data is obtained by running its official source code.






























(a) Sampling Phases





























(b) Framework














# Compute Unified Device Architecture (CUDA)

C	OpenACC, CUDA
C++	Thrust, CUDA C++
Fortran	OpenACC, CUDA Fortran
Python	PyCUDA, PyOpenCL





## Accelerating AutoDock4 with GPUs and Gradient-Based Local Search

Diogo Santos-Martins, Leonardo Solis-Vasquez, Andreas F Tillack, Michel F Sanner, Andreas Koch\*, and Stefano Forli\*



Access Through Your Institution

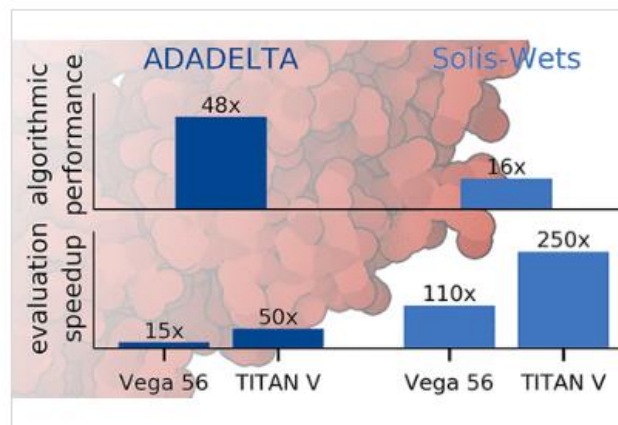
Other Access Options

Supporting Information (1)

### Abstract

AutoDock4 is a widely used program for docking small molecules to macromolecular targets. It describes ligand–receptor interactions using a physics-inspired scoring function that has been proven useful in a variety of drug discovery projects. However, compared to more modern and recent software, AutoDock4 has longer execution times, limiting its applicability to large scale dockings. To address this problem, we describe an OpenCL implementation of AutoDock4, called AutoDock-GPU, that leverages the highly parallel architecture of GPU hardware to reduce docking runtime by up to 350-fold with respect to a single-threaded process. Moreover, we introduce the gradient-based local search method ADADELTA, as well as an improved version of the Solis-Wets random optimizer from AutoDock4. These efficient local search algorithms significantly reduce the number of calls to the scoring function that are needed to produce good results. The improvements reported here, both in terms of docking throughput and search efficiency, facilitate the use of the AutoDock4 scoring function in large scale virtual screening.

Copyright © 2021 American Chemical Society

**Journal of Chemical Theory and Computation**Cite this: *J. Chem. Theory Comput.* 2021, 17, 2, 1060–1073<https://doi.org/10.1021/acs.jctc.0c01006>

Published January 6, 2021

Copyright © 2021 American Chemical Society

[Request reuse permissions](#)

Get e-Alerts

Article Views

5173

Altmetric

17

Citations

158

[Learn about these metrics](#)

### Recommended Articles

**AutoDock Vina 1.2.0: New Docking Methods, Expanded Force Field, and Python Bindings**July 19, 2021 | *Journal of Chemical Information and Modeling*

Jerome Eberhardt\*, Diogo Santos-Martins, Andreas F. Tillack, and Stefano...

**Vina-GPU 2.0: Further Accelerating AutoDock Vina and Its Derivatives with Graphics Processing Units**March 20, 2023 | *Journal of Chemical Information and Modeling*

Ji Ding, Shidi Tang, Zheming Mei, Lingyue Wang, Qinqin Huang, Haifeng Hu, ...

[Show more +](#)

# Vina-GPU 2.0: Further Accelerating AutoDock Vina and Its Derivatives with Graphics Processing Units

Ji Ding, Shidi Tang, Zheming Mei, Lingyue Wang, Qinqin Huang, Haifeng Hu, Ming Ling, and Jiansheng Wu\*



Access Through Your Institution

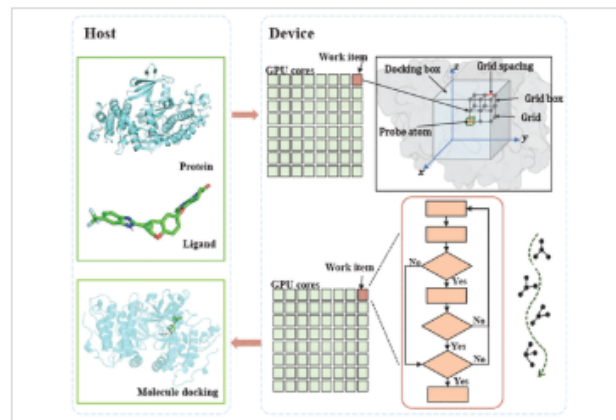
Other Access Options

Supporting Information (1)

## Abstract

Modern drug discovery typically faces large virtual screens from huge compound databases where multiple docking tools are involved for meeting various real scenes or improving the precision of virtual screens. Among these tools, AutoDock Vina and its numerous derivatives are the most popular and have become the standard pipeline for molecular docking in modern drug discovery. Our recent Vina-GPU method realized 14-fold acceleration against AutoDock Vina on a piece of NVIDIA RTX 3090 GPU in one virtual screening case.

Further speedup of AutoDock Vina and its derivatives with graphics processing units (GPUs) is beneficial to systematically push their popularization in large-scale virtual screens due to their high benefit–cost ratio and easy operation for users. Thus, we proposed the Vina-GPU 2.0 method to further accelerate AutoDock Vina and the most common derivatives with new docking algorithms (QuickVina 2 and QuickVina-W) with GPUs. Caused by the discrepancy in their docking algorithms, our Vina-GPU 2.0 adopts different GPU acceleration strategies. In virtual screening for two hot protein kinase targets, RIPK1 and RIPK3, from the DrugBank database, our Vina-GPU 2.0 reaches an average of 65.6-fold, 1.4-fold, and 3.6-fold docking acceleration against the original AutoDock Vina, QuickVina 2, and QuickVina-W while ensuring their comparable docking accuracy. In addition, we develop a friendly and installation-free graphical user interface tool for their convenient usage. The codes and tools of Vina-GPU 2.0 are freely available at <https://github.com/DeltaGroupNJUPT/Vina-GPU-2.0>, coupled with explicit instructions and examples.



**Journal of Chemical Information and Modeling**  
Cite this: *J. Chem. Inf. Model.* 2023, 63, 7, 1982–1998

<https://doi.org/10.1021/acs.jcim.2c01504>

Published March 20, 2023

Copyright © 2023 American Chemical Society

[Request reuse permissions](#)

Get e-Alerts

Article Views

4526

Altmetric

17

Citations

4

[Learn about these metrics](#)

## Recommended Articles

AutoDock Vina 1.2.0: New Docking Methods, Expanded Force Field, and Python Bindings

July 19, 2021 | *Journal of Chemical Information and Modeling*

Jerome Eberhardt\*, Diogo Santos-Martins, Andreas F. Tillack, and Stefano...

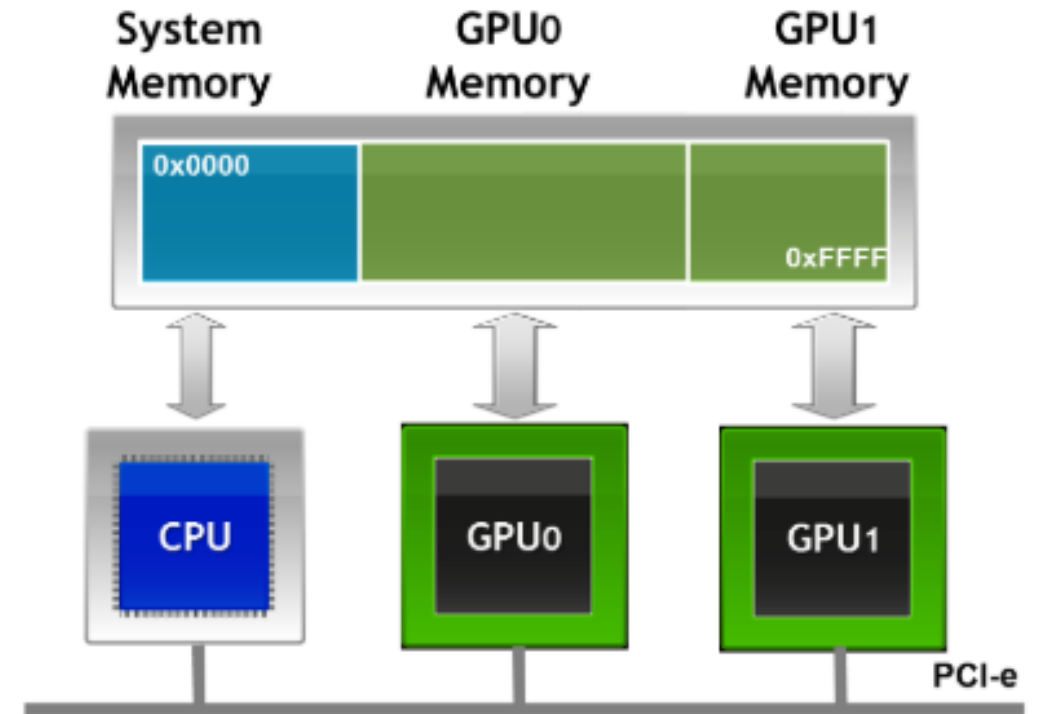
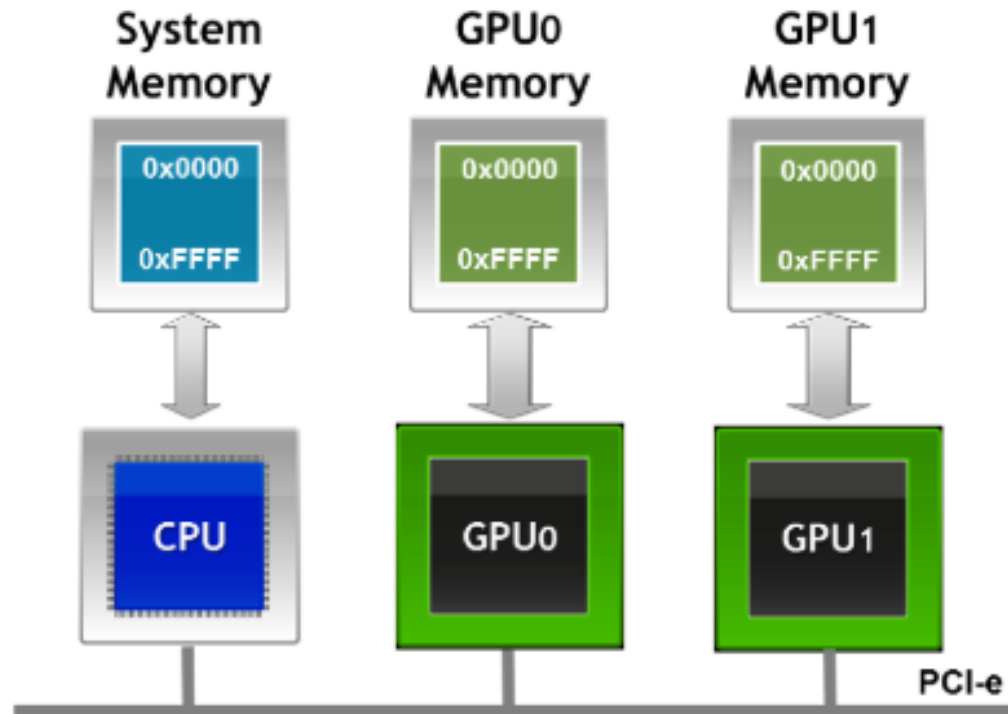
Uni-Dock: GPU-Accelerated Docking Enables Ultralarge Virtual Screening

April 26, 2023 | *Journal of Chemical Theory and Computation*

Yuejiang Yu, Chun Cai, Jiayue Wang, Zonghua Bo, Zhengdan Zhu\*, and Hang...

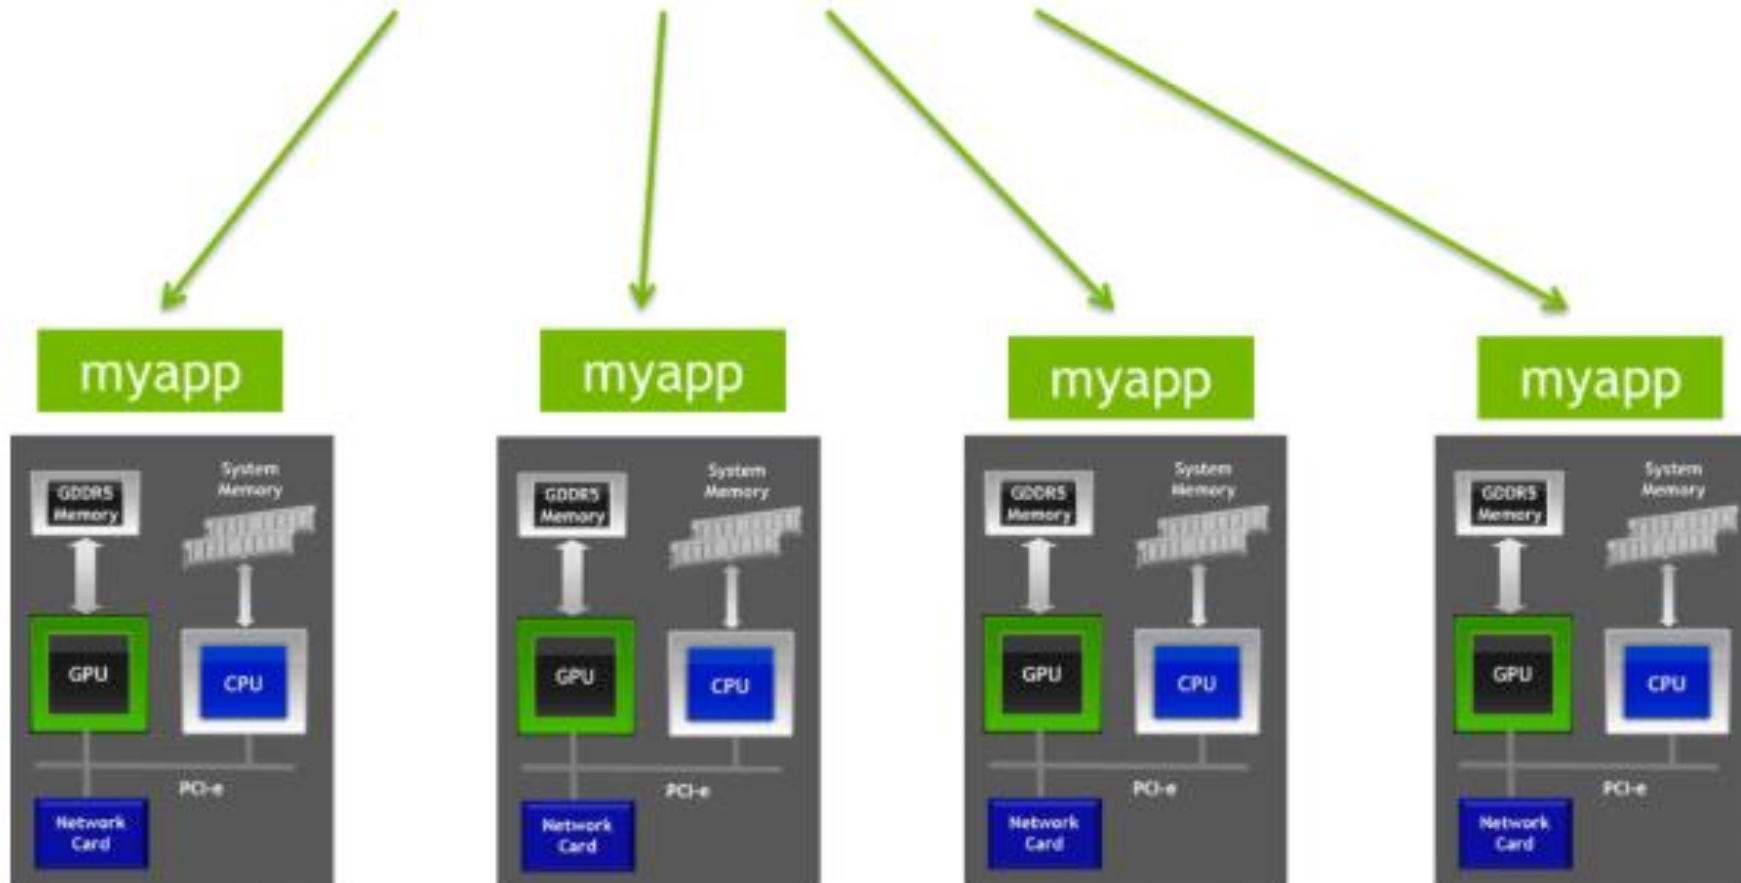
[Show more +](#)

# Message Passing Interface (MPI)

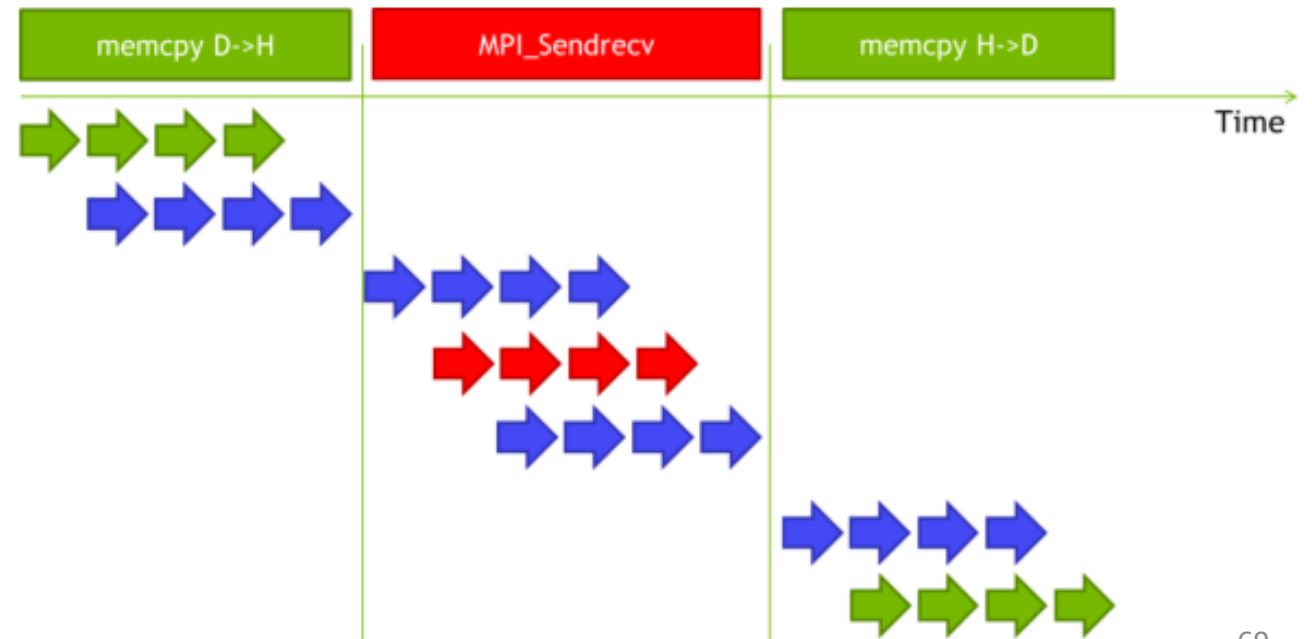



# Message Passing Interface (MPI)

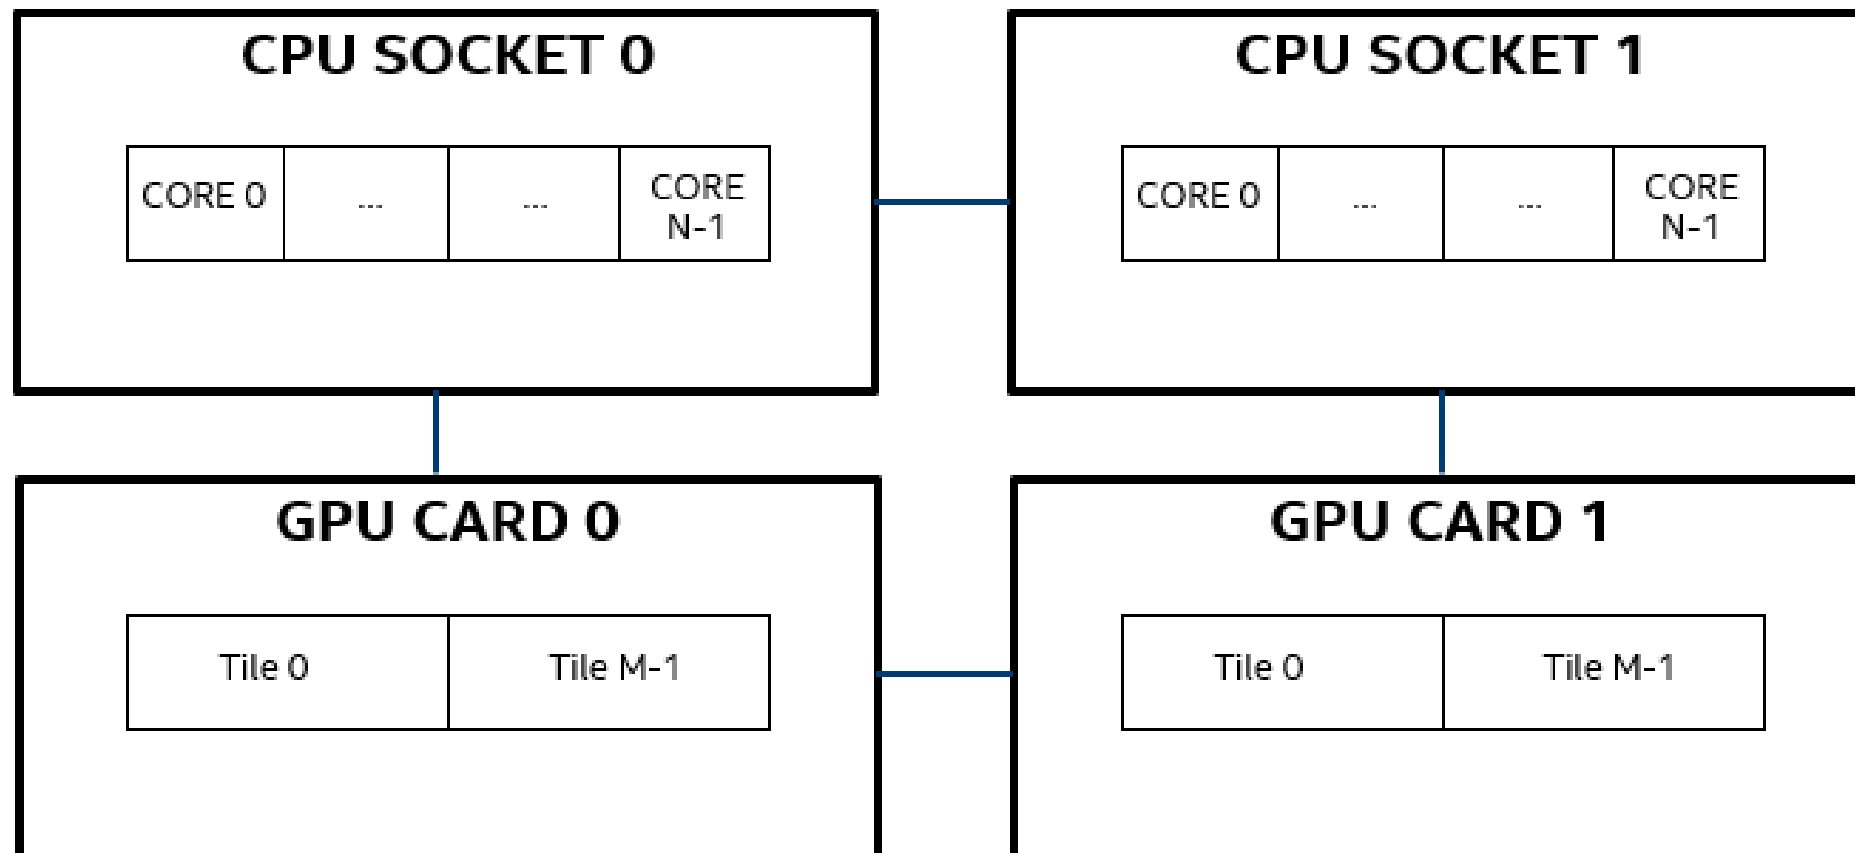
```
mpirun -np 4 ./myapp <args>
```



# Message Passing Interface (MPI)






## HETEROGENEOUS COMPUTE NODE



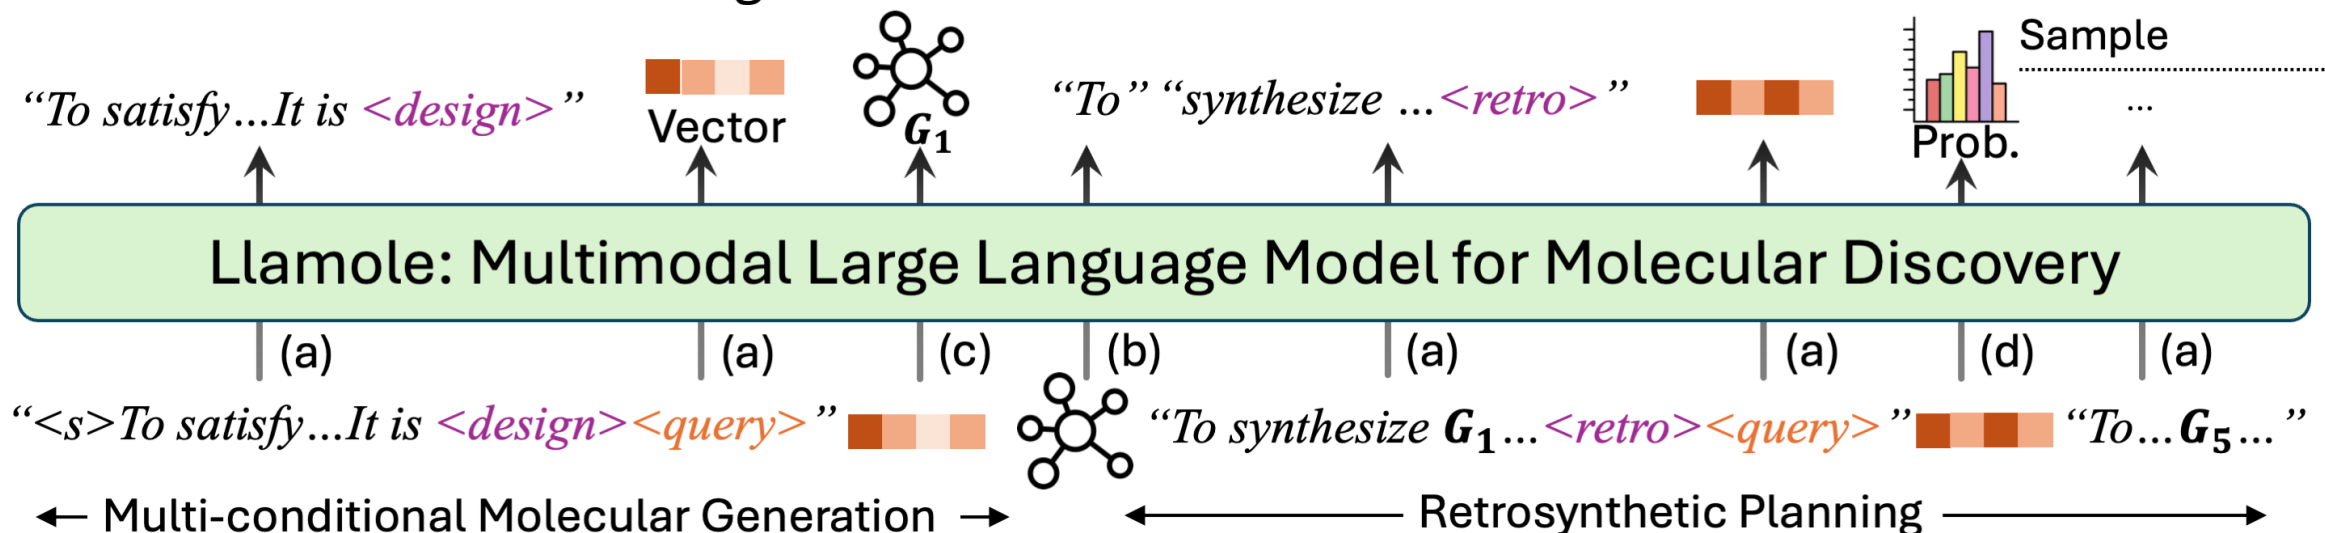


# GROMACS\_MPI (GMX\_MPI)

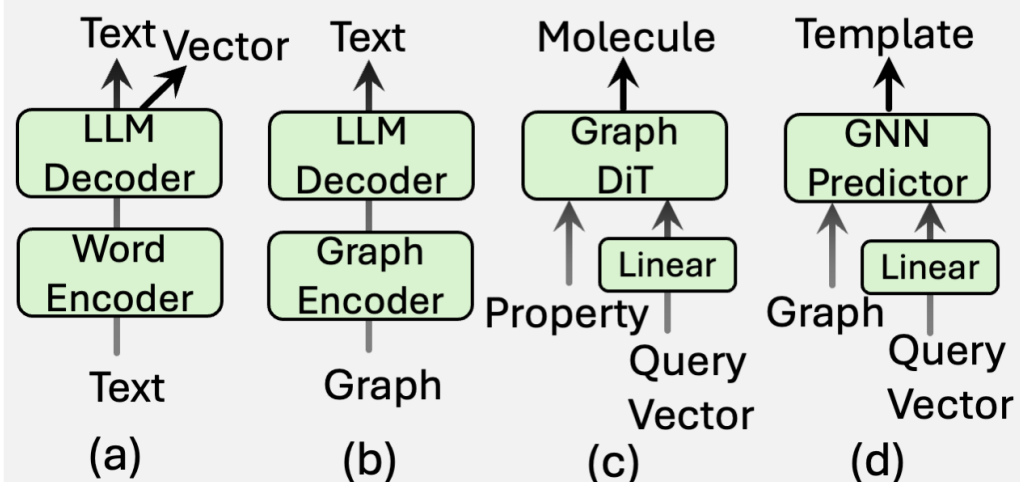




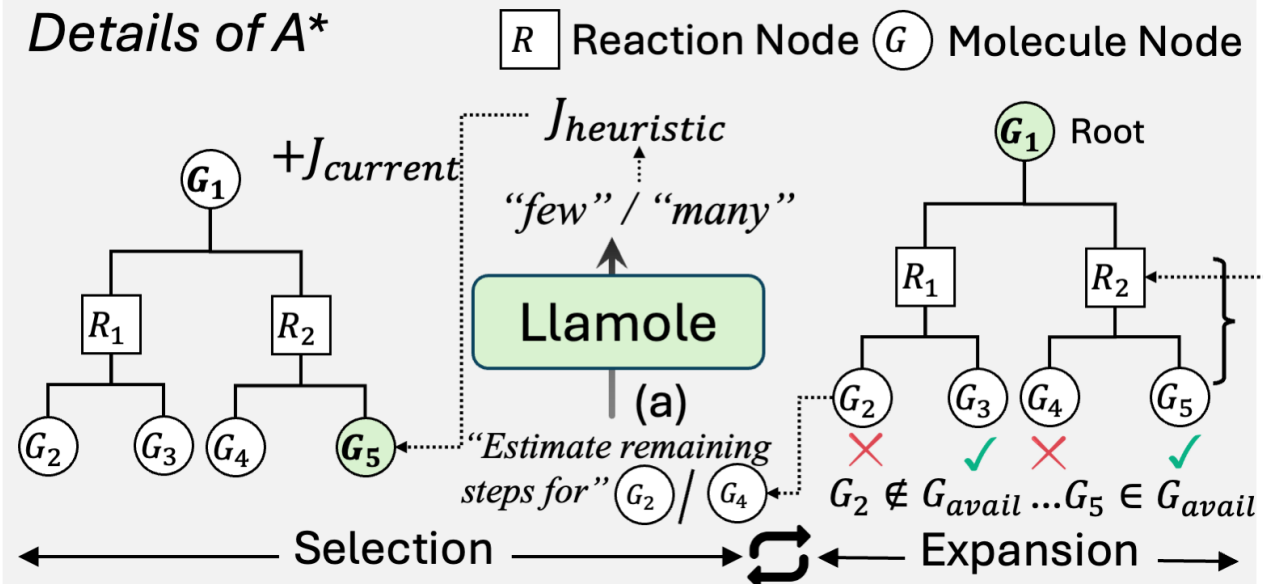
# Llamole's Multimodal Autoregressive Framework



## Details of Active Modules



## Details of $A^*$



### Reaction Prediction (Forward and Retrosynthesis)




Products

Encoder

Decoder

Reactants




### Property Prediction

LogP = 4.23


Property

Encoder

Molecule



### Molecular Optimization




New Molecule


Encoder

Decoder

Molecule




### De Novo Design



New Molecule

Decoder



Smoothed and  
Organized  
Latent Space

THANK  
YOU!