Scalable Graph Neural Networks for Dynamic Molecular Docking

Geometric Equivariance & Spatial-Temporal Graph Attention

Candidate: Alex Mercer

Ph.D. Dissertation Defense

Dissertation Committee: Prof. Eleanor Vance (Chair), Dr. Marcus Thorne, Dr. Sarah Lin
Department of Computer Science & Computational Biology | Vertex Institute of Science

Defense Agenda & Road Map

Structure of Dissertation Presentation

Presentation Overview

1. **Motivation & Problem Context**
   Challenges in 3D molecular binding pose prediction
2. **Theoretical Foundations**
   $\mathrm{SE}(3)$-equivariant message passing frameworks
3. **Proposed SynapseGNN Model**
   Architecture, attention mechanisms, loss functions
4. **Empirical Benchmarks**
   Performance on PDBbind, BindingMOAD, runtime analysis
5. **Conclusions & Outlook**
   Key takeaways, publications, future research

Core Thesis Contribution

Thesis Statement:

"Incorporating strict geometric equivariance with spatio-temporal graph attention reduces conformational search latency by $140\times$ *while achieving sub-angstrom RMSD accuracy in ligand docking."*

- **High Precision:** < 1.2 Å average RMSD on complex binding pockets.
- **Scalability:** Linear time complexity $\mathcal{O}(|V| + |E|)$.
- **Generalization:** Zero-shot transfer to unseen protein families.

SECTION 01

Motivation & Context

Computational Bottlenecks in Molecular Docking & Drug Discovery

The Challenge of Molecular Docking

Limitations of Physics-Based Algorithms

Classical Physics-Based Methods

AutoDock Vina, Glide, Gold:

- Search space exponentially grows with rotatable bonds ($3^{N_{\text{rot}}}$).
- High computational cost (10–60 minutes per target ligand).
- Rigid or semi-flexible protein backbone approximations.
- Frequent false positives due to empirical scoring function limits.

Deep Learning Gaps

Standard Graph Neural Networks:

- Lack strict 3D physical symmetry ($\mathrm{SE}(3)$ rotational invariance).
- High sensitivity to arbitrary coordinate frame transformations.
- Over-smoothing in deep message-passing architectures.
- Poor generalization to novel side-chain conformations.

Goal of dissertation: Bridge the gap between rigid physical symmetry constraints and rapid neural inference to produce real-time, high-fidelity binding predictions.

Key Dissertation Contributions

Three Pillars of SynapseGNN

Pillar I

$\mathrm{SE}(3)$-Equivariant Vectors

- Mathematically rigorous rotational and translational equivariance.
- Direct coordinate updates without frame canonicalization.

Pillar II

Spatial-Temporal Attention

- Multi-scale distance encoding across protein-ligand interfaces.
- Adaptive message pruning for distant atomic pairs.

Pillar III

Multi-Task Scoring Objective

- Simultaneous pose (RMSD) and affinity (ΔG) estimation.
- End-to-end differentiable loss landscape.

Target Impact: Screening 10M+ Virtual Compound Libraries in under 4 hours on standard GPU cluster.

SECTION 02

Theoretical Framework

Equivariant Representations & Convergence Analysis

Problem Formulation & Symmetry Constraints

$\mathrm{SE}(3)$-Equivariant Message Passing

Consider a molecular graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with atomic features $h_i \in \mathbb{R}^d$ and 3D positions $x_i \in \mathbb{R}^3$.

Equivariance Definition

A transformation $\check{} : \mathcal{X} \to \mathcal{Y}$ is $\mathrm{SE}(3)$-equivariant if for all rotation matrices $\mathbf{R} \in \mathrm{SO}(3)$ and translations $g \in \mathbb{R}^3$:

$$\check{}(\mathbf{R} \cdot \mathbf{x}_i + g) = \mathbf{R} \cdot \check{}(\mathbf{x}_i) + g$$

Scalar Feature Update:

$$m_{ij} = \phi_m\left(h_i^{(l)}, h_j^{(l)}, \|x_i^{(l)} - x_j^{(l)}\|^2, e_{ij}\right)$$

$$h_i^{(l+1)} = \phi_h\left(h_i^{(l)}, \sum_{j \in \mathcal{N}(i)} m_{ij}\right)$$

Vector Position Update:

$$x_i^{(l+1)} = x_i^{(l)} + C \sum_{j \in \mathcal{N}(i)} (x_i^{(l)} - x_j^{(l)})\, \phi_x(m_{ij})$$

Guarantees strict physical covariance under arbitrary 3D rigid body motions.

Theoretical Guarantees

Convergence & Expressive Power Bounds

Theorem 1 (Equivariant Approximation Bound, Mercer 2026)

Let $f^* : \mathbb{R}^{N\times 3} \to \mathbb{R}^{N\times 3}$ be any smooth $\mathrm{SE}(3)$-equivariant force-field operator. For any $\epsilon > 0$, there exists a L-layer SynapseGNN network with hidden dimension $d \geq \mathcal{O}(\log(1/\epsilon))$ such that:

$$\sup_{\mathbf{X}\in\mathcal{K}} \|f^*(\mathbf{X}) - \mathrm{SynapseGNN}(\mathbf{X})\|_F \leq \epsilon$$

where $\mathcal{K}$ is any compact domain of atom coordinates.

Key Proof Takeaways

1. **No Information Loss:** Invariant radial basis functions preserve pairwise metric topology.
2. **Resolution of Over-Smoothing:** Layer-wise residual vector shortcuts ensure non-vanishing gradient flow across $L = 16$ message passing iterations.
3. **Bounded Variance:** Normalization layer keeps positional updates $\|x_i^{(l+1)} - x_i^{(l)}\| \leq \delta_{\max}$.

SECTION 03

SynapseGNN Architecture

Deep Equivariant Network & Spatial-Temporal Flow

SynapseGNN Architecture

End-to-End Equivariant Pipeline

Spatial Attention Module

- Multi-head attention over pairwise distance kernels.
- Soft-cutoff radial envelope at $r_c = 10.0$ Å.
- Weight $\alpha_{ij} = \mathrm{softmax}\left(\frac{q_i^T k_j}{\sqrt{d_k}} + \psi(r_{ij})\right)$.

Optimization & Loss Function

- Loss: $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{RMSD}} + \lambda_1 \mathcal{L}_{\text{affinity}} + \lambda_2 \mathcal{L}_{\text{clash}}$.
- Steric clash penalty prevents unphysical atomic overlaps.
- AdamW optimizer with cosine annealing ($\eta_0 = 5 \times 10^{-4}$).

SECTION 04

Empirical Evaluation

Benchmarks, Accuracy Metrics, & Screening Throughput

Quantitative Docking Performance

Evaluation on PDBbind v2020 & BindingMOAD Core Sets

Model / Algorithm	Top-1 RMSD (< 2Å %)	Median RMSD (Å)	pK_d R^2	Runtime / Complex
AutoDock Vina (Physics)	48.2%	2.45	0.42	14.2 min
Glide SP (Physics)	52.6%	2.10	0.51	8.5 min
SchNet (Baseline GNN)	41.3%	2.89	0.38	1.2 sec
EGNN (Equivariant)	64.5%	1.62	0.64	0.8 sec
EquiBind (Direct Pose)	68.1%	1.48	0.67	0.4 sec
SynapseGNN (Ours)	**83.4%**	**0.92**	**0.84**	**0.06 sec**

Key Performance Highlights

- **Sub-Angstrom Accuracy:** 83.4% of predicted poses within < 2Å RMSD.
- $140\times$ **Speedup:** Reduces docking time from minutes to milliseconds.

Dataset Splitting

- Rigorous time-split and 30% sequence identity cutoff.
- Zero training overlap with test set protein structures.

Ablation Studies & Generalization

Understanding Module Contributions

Component Sensitivity Analysis

Variant Configuration	RMSD ($<$ 2Å %)	Δ Perf.
Full SynapseGNN Model	**83.4%**	–
w/o Equivariant Vector Layer	59.1%	-24.3%
w/o Spatial Attention $\psi(r_{ij})$	71.8%	-11.6%
w/o Steric Clash Penalty	76.2%	-7.2%
w/ Single-Scale Radius	74.0%	-9.4%

Generalization Out-of-Distribution

Tested on Novel Kinase Targets:

- Evaluated on 150 cryo-EM structures from 2025–2026.
- Maintained **79.2%** success rate ($<$ 20% sequence homology).
- Robust performance on highly flexible loop regions.

Result: Equivariant coordinate updating is the primary driver of pose accuracy.

SECTION 05

Conclusions & Outlook

Dissertation Summary, Publications, & Future Horizons

Dissertation Summary

Core Findings & Defense Takeaways

Summary of Achievements

1. Developed **SynapseGNN**, an end-to-end $\mathrm{SE}(3)$-equivariant graph neural network.
2. Achieved **sub-angstrom (**0.92 **Å)** pose accuracy at **60ms per complex**.
3. Proven theoretical convergence and universal approximation bounds.
4. Deployed in virtual screening pipeline at **Nexa BioPharma**.

Future Research Directions

- **Full Protein Flexibility:** Extending message passing to ensemble side-chain movements.
- **Covalent Binding:** Incorporating chemical reaction energy barriers into loss function.
- **Generative Design:** Coupling SynapseGNN with diffusion models for de novo design.

Publications & Academic Output

Peer-Reviewed Contributions

Key Publications Derived from Dissertation

1. **A. Mercer**, M. Thorne, E. Vance. "Equivariant Spatial-Temporal Graph Attention for Molecular Docking." *Journal of Chemical Information and Modeling (JCIM)*, 2025.
2. **A. Mercer**, S. Lin, E. Vance. "Ultra-Fast Virtual Screening via SE(3)-Equivariant Neural Operators." *NeurIPS Computational Biology Workshop (Oral)*, 2025.
3. **A. Mercer**, E. Vance. "Convergence Guarantees for Deep Geometric Message Passing on Molecular Manifolds." *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, under review, 2026.

Acknowledgments & Funding

- **Advisory Committee:** Prof. Eleanor Vance, Dr. Marcus Thorne, Dr. Sarah Lin.
- **Research Grants:** Vertex Doctoral Fellowship (Grant #VDF-2023-882) & Nexa AI Research Grant.

Thank You!

Questions & Defense Discussion

Contact: Alex Mercer | a.mercer@vertex.edu
Code & Benchmarks: github.com/vertex-lab/synapse-gnn