Inductive Graph-based Surrogate Models for Transportation Policy Evaluation across Cities

Master‘s Thesis of Ankit Basu

Mentoring:
Elena Natterer M.Sc.
Dr. Antonios Tsakarestos

External Mentoring:
Dominik Fuchsgruber M.Sc.
TUM School of Computation, Information and Technology

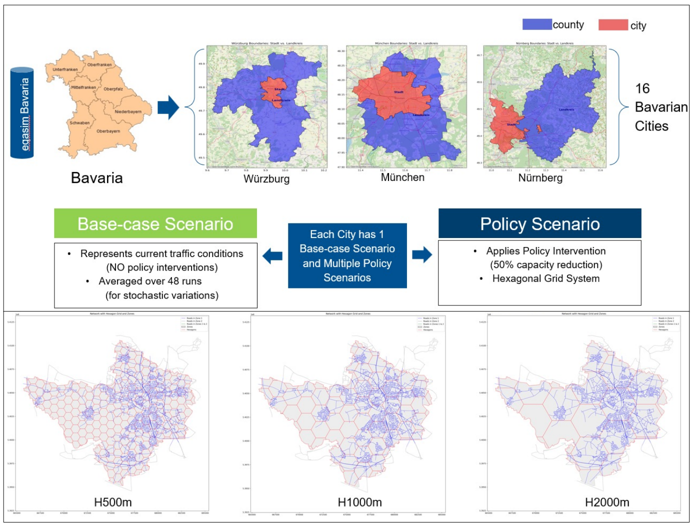

The growing complexity of urban mobility systems has made transportation policy design increasingly dependent on simulation-based approaches. Among these, Agent-Based Models (ABMs) stand out for their ability to capture individual travel behaviors at high fidelity. However, their use in city-scale policy analysis is severely constrained by computational and scalability constraints, making extensive scenario exploration impractical. Recent research demonstrates that **Graph Neural Networks (GNNs)** can serve as effective **surrogate models** for Agent-Based Models (ABMs) in transportation policy evaluation, achieving high accuracy and computational efficiency when trained and tested within the same city. Building upon current research, this thesis tries to investigate whether such surrogate models can **generalize inductively** across unseen and heterogeneous urban areas, which could enhance the scalability and computational constraints of agent-based transportation policy analysis.

In the transductive setting, Surrogate Models build with attention based and convolutional architectures (GATv2, Trans-formerConv) achieved R2 scores in the mid-60s, while GraphSAGE with pooling reached the high-70s. The best results were obtained with the hybrid Transformer Encoder using a convolutional backbone, achieving R2 ≈ 0.9. In the inductive setting (unseen cities), the same hybrid model showed limited but better generalization than neighborhood aggregation models like GraphSAGE, suggesting that its learned representations can transfer partially across different network structures.

Further experiments explored model fine-tuning with sparse graph samples to assess the impact of adaptation on generalization. When fine-tuned using only a small fraction of graphs (approximately 2%) from Landshut, the inductive model demonstrated a notable improvement in predictive accuracy, reaching $R^2 \approx 0.95$, nearly matching transductive performance.

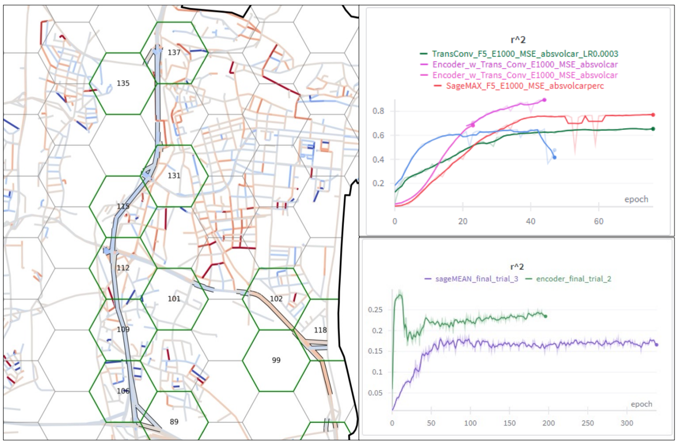

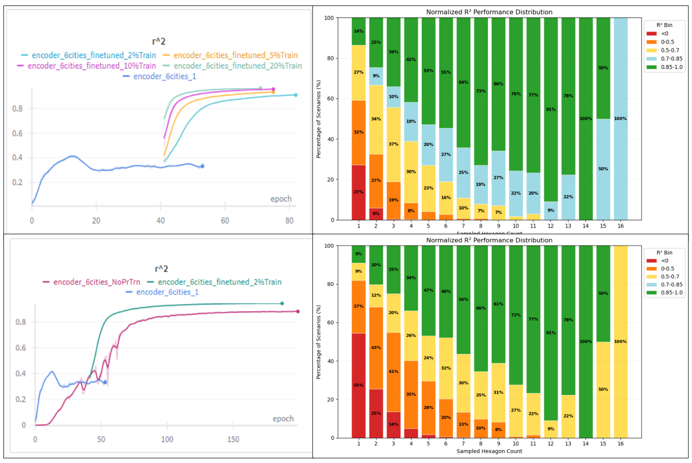

Conversely, in the sparse transductive model, where the model was trained and tested on identically distributed sparse graphs as for the fine-tuned inductive model, the R2 achieved was approximately 0.88. While this is a reasonable level of accuracy, the performance remained below that of the fine-tuned inductive model. This difference was also reflected in the overall test set performance, where the sparse transductive model underperformed across most test scenarios. The comparison indicated that fine-tuning provides an edge over capturing transferable spatial dependencies, allowing the model to generalize more effectively even under limited data conditions.

This demonstrates that targeted fine-tuning, even with limited data, enables the model to efficiently emulate the agent-based model simulation output with near equal predictive accuracy and also in much less time, making it computationally efficient for exhaustive policy scenario analysis and optimization.