Predicting Bike Traffic Using Graph Neural Networks: Integrating Residential Density, Amenity Distribution, and Street Networks

Alternative title: Pedaling Through the Data: Predicting Bike Traffic

Wen-Chuang Chou

Hey there, fellow bike enthusiasts and data adventurers! Have you ever wondered what makes people hop on their two-wheeled steeds and pedal away, leaving traffic jams and fuel costs in their rearview mirrors? Well, my friends, I've got some exciting news for you. I've been on a quest to unravel the mysteries of bike traffic prediction using the power of machine learning, specifically Graph Neural Networks (GNNs). Trust me, it's as cool as it sounds!

Now, I know what you're thinking. "Graph Neural Networks? That sounds like rocket science!" Well, fear not, my friends. I'm here to break it down for you in the most fun and engaging way possible. So, buckle up your helmet, because we're about to embark on a wild ride through the world of data and cycling.

First things first, let's talk motivation. We all know that cycling is not only a fantastic way to get around, but it's also better for our planet. Plus, with skyrocketing fuel costs and public transportation fares, more and more folks are turning to their trusty bikes as a mode of transportation. But here's the kicker: authorities want to understand what drives people to choose bicycles, how urban development can boost cycling demand, and how to promote biking through dedicated infrastructure. It's like solving a puzzle on wheels!

So, here's the task at hand. We're predicting bike traffic volume, and we're doing it by crunching numbers from a social and environmental dataset. We're looking at things like residential density (because let's face it, more people, more bikes), distribution of amenities (who wouldn't want to pedal to their favorite cafe or fitness center?), and street networks (the backbone of our cycling adventures).

But wait, where do we get all this glorious data? Well, a big shoutout to the Stadtradeln app users in Germany who recorded their cycling traces during a three-week period. You guys rock! And thanks to the MOVEBIS research project at TU Dresden, we were able to get our hands on this treasure trove of information. Now, keep in mind that the data represents the number of participants in the cycling campaign, not the total number of bicycle users in the cities. So, don't worry if you're not seeing every cyclist out there on the graph.

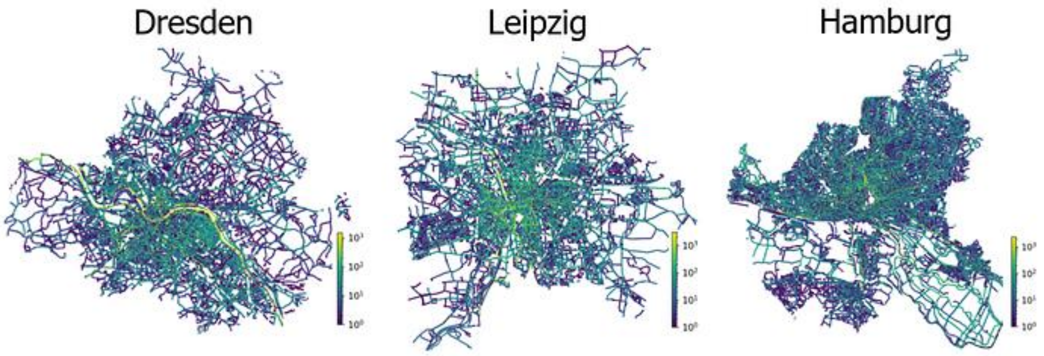

Fig.1 Pedal Power Unleashed: Peek into the Cycling Chronicles of Three Cities!

Speaking of graphs, let's dive into the nitty-gritty of how we wrangle all this data. We're using OpenStreetMap to get the scoop on street networks and types. From cycleways to main streets, residential streets to paths, we're covering it all. And when it comes to residential density, we're using the total floor area of buildings for approximation. Bigger buildings, more people, you get the picture.

Fig.2 From Maps to Graphs: The transformation of OpenStreetMap data into graph structures.

Now, hold on tight because here comes the fun part. We're employing some sneaky keywords from the OSM classification hierarchy to find all the amenities that make our cycling adventures even sweeter. Think department stores, supermarkets, cafes, and even biergartens. Yes, you read that right. We're not just predicting bike traffic; we're predicting the number of nearby spots where you can grab a cold one after a long ride. Talk about priorities, right?

But how do we piece all this together? Well, imagine a giant interconnected web of streets forming a graph-based data structure. Each street segment becomes a node, and each node is armed with 43 features. We're talking about segment length, street type, floor area, and the number of amenities. It's like giving each street its own personality and charm.

And guess what? Our secret weapon comes from Graph Neural Networks called Graph Attention Network. It's the powerhouse that captures all the mobility patterns associated with bicycle users and makes those magical predictions.

Now, it's time for the big reveal. Drumroll, please! We trained our model on the data from three magnificent cities: Dresden, Leipzig, and Hamburg. These cities became our training ground for bike traffic prediction excellence. And boy, did our model deliver! We managed to minimize the discrepancies between predicted and actual bike traffic, all while having a blast with our graphs and GNNs.

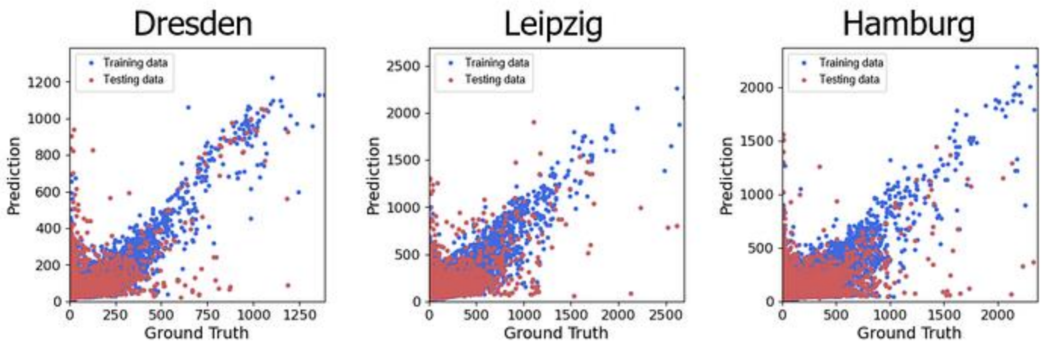

Fig.3 Unveiling Predictive Insights: The predictions derived from models trained on three diverse cities.

But that's not all, my friends. We took it a step further and tested the transferability of our model across different cities. And you know what? It aced the test! Our model proved its versatility and showed that it's ready to conquer new territories, one city at a time.

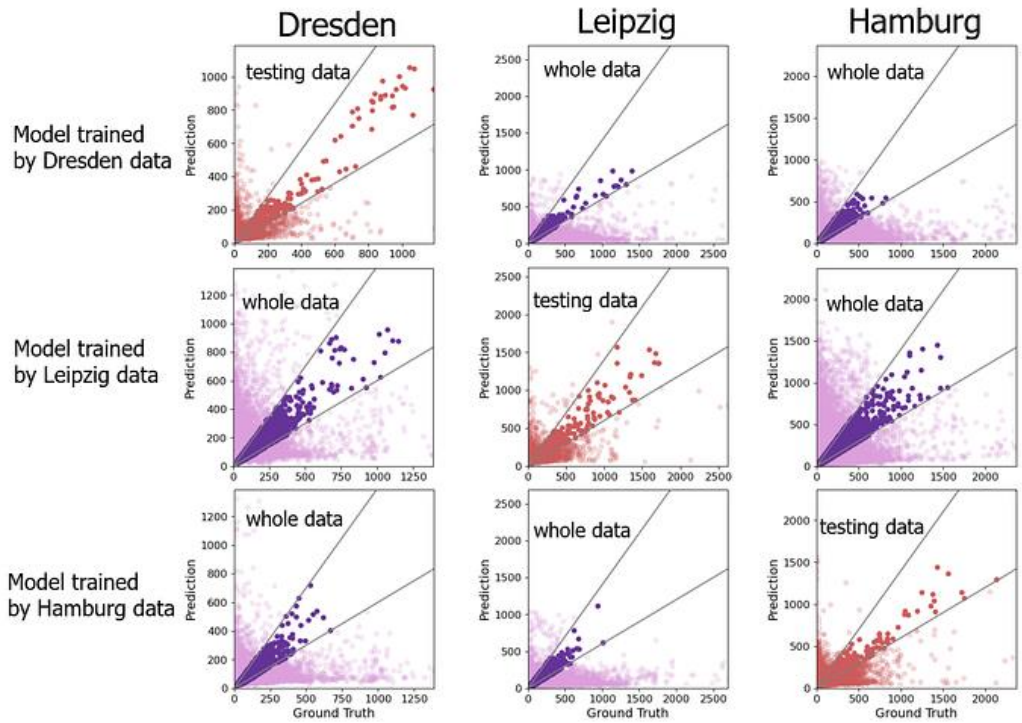

Fig.4 City-Hopping Success: Explore the transferability of our trained models across different urban landscapes.

Now, I know you're itching to dig deeper into the technique details. And guess what? I've got you covered. Just head over to my GitHub pages for all the juicy goodness.

So, there you have it, folks. We've cracked the code on predicting bike traffic using Graph Neural Networks. We've explored residential density, amenity distribution, and street networks, all while having a blast with our data and some seriously cool algorithms. It's time to hop on your bikes, embrace the power of data, and pedal your way to a greener and more exciting future.

Ride on, my friends, ride on!