Project description AI master thesis

Project name	Predictive Parking
Description	The project Predictive Parking is geared towards predicting the so-called parking pressure (fraction of parking spots occupied) per region for the near future. Last year, two students worked on imputing unobserved data - the results of which could be directly put to use in reports used by management and *stadsdeel* advisors - and predicting the parking pressure per neighbourhood for the next hour(s). We are currently busy bringing the pipelines generating the ground truth, as well as the most promising model into production on our new data platform. This will enable us to train/impute automatically based on the latest available data, allowing for a faster iteration over model versions. A potential thesis project could be related to bringing said processes into production and optimising the model. Previous work has shown that particular data sources, as well as different modelling approaches greatly improve the model quality. Synthesising the benefits of these existing models, combined with said data sources is a promising avenue of research. Due to time, data-availability and compute constraints, we have thus far had to restrict ourselves to predictions at the neighbourhood/hour scale. Given the increase in compute from our new data platform, another potential thesis project could investigate modelling parking at the road segment level, using the Amsterdam road network graph. Available data sources include, but are not limited to, scan car data, data on parking permits, road segments and parking spots (and their historic locations throughout the city), weather/precipitation data, etc.
Work environment	The Parking department of the City of Amsterdam where the expert knowledge regarding parking (policies, data interpretation, etc) resides, combined with the [redacted] for more in-depth technical support (data engineering/infrastructure, ML, etc). Both environments are very informal, and aim to achieve directly applicable results (i.e. not merely an academic paper). Last year, the COVID-19 situation led to both students doing the better part of their thesis projects remotely with an occasional face-to-face meeting. This worked well in combination with a scrum-like way of working (using [redacted] and the ease of checking in with everyone involved (via [redacted]).
Expectations	The results of this project should be directly applicable and relevant to the Parking department. In addition to the delivery of useful predictive models, relevant results also include byproducts of the whole data analysis/processing pipeline that could be turned into insightful reports. Supervision will be shared between the academic institution (UvA / VU / …) and data scientists/analysts at the City of Amsterdam.