

Parking Tour Overview

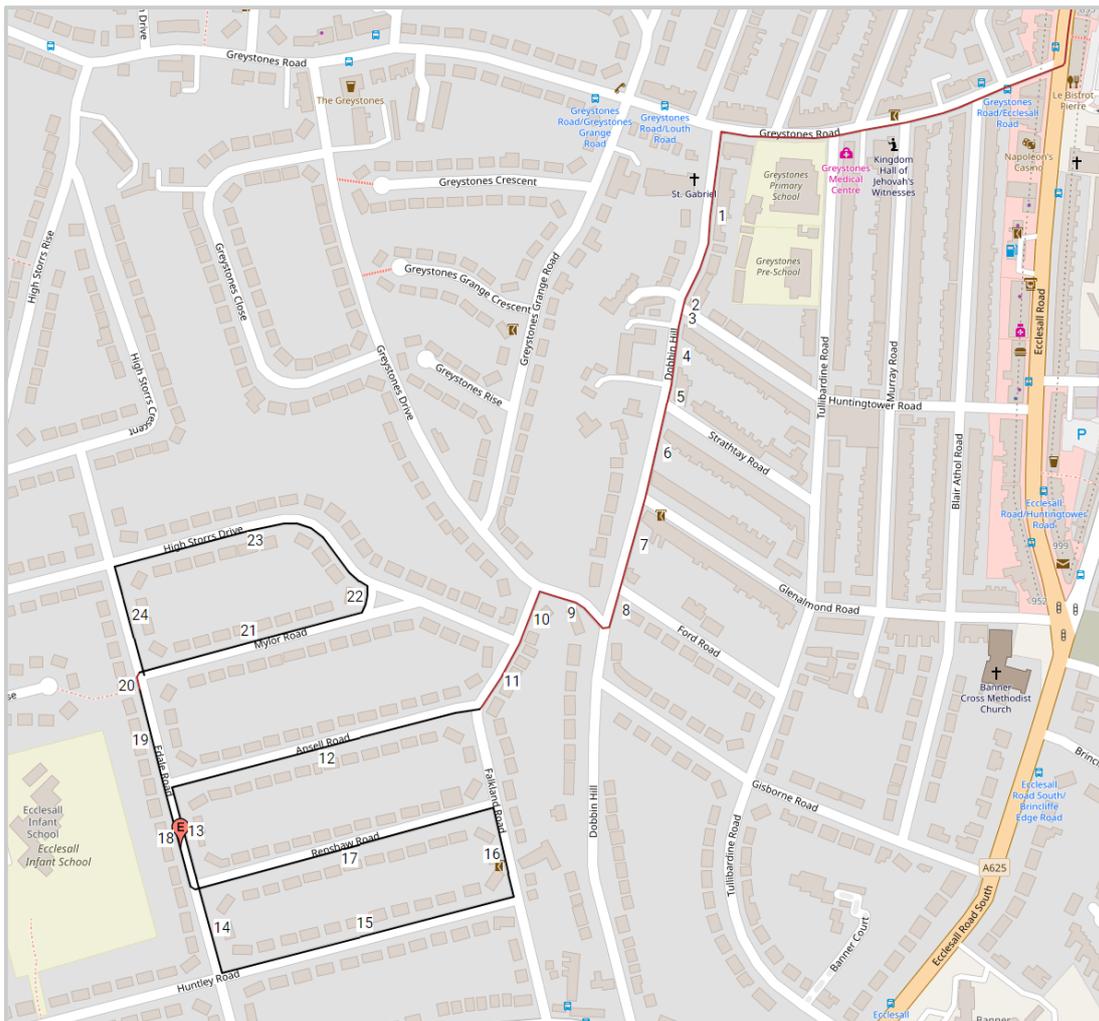
Cotares Limited

August 2018

A "Parking Tour" is the sequence of roads that an expert driver might explore when looking for an on-street parking place. The intention is that a driver will be guided along this tour in identical fashion to their usual navigational guidance, but with additional information suggesting that they should park, or not park, if they see an available space.

Our Parking Tour Generator (PTG) is given a start location (road link), a destination (LatLon on a walking map), and then generates a route (potentially inter-continental) that ends in a tour around the vicinity of the destination, as defined by the maximum walking time.

The image below is computed using data © [OpenStreetMap contributors](#). The Tour is shown in red (don't park yet) and black (park if you see a space). Note that these tours in our images are offset to the driving side, left in the case of the UK. The destination is shown as red teardrop marked "E". The link numbered "1" is the first one that is within the maximum walking time (ten minutes in this case):



The Tour differs from a "Route" in several important ways:

- 1) It does not have a specific road that has to be reached. It might recommend parking on any roads within easy walking distance of the ultimate destination. The destination is shown as a red teardrop above.
- 2) It does not know on which part of the tour (shown in red and black above) the driver might find an available parking space, but it does know the probabilities.
- 3) It may involve visiting the same road two or more times because that is the easiest way to explore the area.
- 4) It may involve visiting a series of roads two or more times because they have had time for spaces to be vacated.
- 5) It has a point in the tour where the driver is told to start looking to park: link 12 where the road line colour changes from red to black. Before that, even though we are within the maximum walking time (10 minutes in this case) the driver is advised NOT to take available spaces because there are better ones to come.
- 6) It may have additional points in the tour where the driver is told not to park, either because the walking route is unusually difficult, or because of time restrictions or cost of available spaces.

Cost Function

The "Parking Tour Generator" (PTG) itself is very different from a "Router", but is still driven by a "cost function" that is trying to model the tradeoffs that a driver makes. It includes a weighted sum of driving time, walking time, cost of parking, permitted parking times, ease of driving, ease of walking etc. We do not always know all of this information, but the algorithm will more accurately reflect the driver's preferences if we can provide it.

Probabilistic Cost Function

Another difference from a conventional "Router" is that we do not know what the costs will be for a particular tour, because we do not know where the parking spaces will be available. Consequently, we do not minimise the component "walking distance", but we minimise the component "*expected* walking distance" given the probabilities of parking at each stage of the Tour.

Using this probabilistic cost-function approach, we can easily accommodate features such as walking less if it is raining (at the expense of a longer search time), or using less time overall (at higher cost) if running late, and the Tours adapt naturally to changes in parking probabilities.

Many expected behaviours emerge automatically from the combination of cost function and the probabilistic approach. For example, where parking is easily available, the tour will head straight to the vicinity of the destination before asking the driver to park. This avoids a driver walking unnecessarily far past many empty spaces. Conversely, when parking is rare, it may first explore roads further away from the destination to maximise the area covered in the shortest time. This avoids the driver being taken past a valuable space before a long frustrating exploration to find one closer.

Some of these important aspects are discussed further in our technical documents.

Data Input

The generator is written to use as much information as can be provided, and uses reasonable defaults if there are unknowns. For example, the historical parking data for a given time of day is reduced to a probability of parking on each road link, and a "churn interval" so that we know how soon it might be worth re-visiting a link. If the churn interval is unknown, a default of two minutes works well, and if the parking probabilities are unknown, a default of 0.1 per 100m on appropriate roads generates tours that are reasonable in many circumstances.

It is when parking is rare that it becomes more important to import measured data, so that the tours will concentrate in the roads that have been observed to have more spaces available in the past, and minimise the search times.