JTLU

# Calculating place-based transit accessibility: Methods, tools and algorithmic dependence

**Christopher D. Higgins**
University of Toronto Scarborough
cd.higgins@utoronto.ca

**Matthew Palm**
University of Toronto
matthew.palm@utoronto.ca

**Amber DeJohn**
University of Toronto
amber.dejohn@mail.utoronto.ca

**Yang Luna Xi**
University of Toronto
luna.xi@mail.utoronto.ca

**James Vaughan**
University of Toronto
james.vaughan@utoronto.ca

**Steven Farber**
University of Toronto Scarborough
steven.farber@utoronto.ca

**Michael Widener**
University of Toronto
michael.widener@utoronto.ca

**Eric J. Miller**
University of Toronto
eric.miller@utoronto.ca

**Abstract:** To capture the complex relationships between transportation and land use, researchers and practitioners are increasingly using place-based measures of transportation accessibility to support a broad range of planning goals. This research reviews the state-of-the-art in applied transportation accessibility measurement and performs a comparative evaluation of software tools for calculating accessibility by walking and public transit including ArcGIS Pro, Emme, R5R, and OpenTripPlanner using R and Python, among others. Using a case study of Toronto, we specify both origin-based and regional-scale analysis scenarios and find significant differences in computation time and calculated accessibilities. While the calculated travel time matrices are highly correlated across tools, each tool produces different results for the same origin-destination pair. Comparisons of the estimated accessibilities also reveal evidence of spatial clustering in the ways paths are calculated by some tools relative to others at different locations around the city. With the growing emphasis on accessibility-based planning, analysts should approach the calculation of accessibility with care and recognize the potential for algorithmic dependence in their calculated accessibility results.

## 1    Introduction

Transportation accessibility describes the potential for reaching spatially distributed opportunities while considering the difficulty involved in travelling to them (Páez et al., 2012). Researchers advocate for accessibility-based planning over traditional mobility-based planning because the former better captures the complex relationships between transportation and land use (Proffitt et al., 2019). Accessibility

measures fall into four categories depending on whether they are infrastructure-based, person-based, utility-based, or place-based (Geurs & van Wee, 2004). Place-based metrics are the most used group of accessibility measures in transportation planning practice, particularly for regional planning (Boisjoly & El-Geneidy, 2017). They are also easier to interpret, operationalize, and communicate than other measures (Geurs & van Wee, 2004). To capture place-based accessibility, a variety of different metrics have been developed (see Siddiq & Taylor (2021) for a review of 54 place-based accessibility metrics and their potential for planning practice) and this versatility has led to their application in service of a broad range of planning goals, including public health (Giles-Corti et al., 2005), sustainability (Handy, 2008), social equity (Martens & Golub, 2018), and economic attainment (Blumenberg & Pierce, 2014), among others.

Although researchers and practitioners have been utilizing measures of place-based accessibility for several decades, there are a number of issues and considerations that should be taken into account when performing practical accessibility analysis. In terms of best practice, there is a rich history of theoretical development and a large body of academic and grey literature on applications to be considered. On the practical side, users are presented with a growing suite of software options, each of which entails a different workflow and relative strengths and weaknesses. The purpose of this study is to review the current best practice and state of the art in place-based accessibility measurement, explore how such measures are currently used in transportation analysis and planning, perform a comparative evaluation of software tool computational performance, and examine the role of algorithmic dependence in the calculated accessibility scores.

The paper begins with a conceptual overview of different place-based accessibility measures, highlighting their strengths and weaknesses. Then we examine how researchers have improved upon the technical limitations of standard metrics, exploring the state of art in research. After that, we introduce our framework for reviewing accessibility tools and test five software packages for multimodal routing (ArcGIS Pro, Emme 4.4, OpenTripPlanner in R and Python, and R5 in R) in local and regional analysis scenarios in the City of Toronto, Canada, and analyze their outputs in terms of travel time and calculated accessibilities.

## 2        Literature review

### 2.1        Place-based accessibility measures

Of the four types of accessibility measures described by Geurs and van Wee (2004), infrastructure-based measures describe the performance of the transportation system including metrics like time lost to congestion and travel time savings. Person-based measures calculate the number and duration of possible activities that individuals can participate in considering their typical time constraints (e.g., work, school), household commitments (e.g., chaperoning children), travel costs, home and work locations, etc. (Fransen et al., 2018). Utility-based measures estimate the consumer surplus, or benefits, that individuals derive from their travel, activity, and residential location choices, reflecting individual preferences, travel constraints, and built form impacts simultaneously (Miller, 2018).

Finally, place-based measures examine what is reachable from a location (e.g., a traffic analysis zone or census block) given land use and transport network constraints. Measures of place-based accessibility range in sophistication, from simple proximity measures to more complex indicators that account for competition for activities and time of day variation in their availability. This section begins with the most simplistic measures and moves to the most complex, highlighting their strengths, weaknesses, and data requirements.

## 2.2        Proximity and threshold measures

Proximity measures work in two ways. First, they can indicate whether a location is within a given distance of a destination of interest. The Toronto Transit Commission's service standard of ensuring that all Torontonians are within 400 meters of a transit stop is an example of this type of threshold (TTC, 2017). Proximity metrics cannot account for the range, number, and diversity of options accessible to travelers within reasonable travel times.

## 2.3        Dual or minimum travel time measures

Minimum travel measures can indicate the minimum distance or travel time to the nearest location of a service. For example, healthcare researchers sometimes use distance or travel time to the nearest healthcare facility as a measure of healthcare accessibility. In contrast to more 'primal' measures of access that consider how many of something can be reached in a given travel time, Cui and Levinson (2020) refer to this as the dual of accessibility – the travel costs involved in reaching a fixed number of something. These dual access measures are significantly correlated with patient utilization of healthcare services in most studies in which they are used (Kelly et al., 2016). They are most often deployed as policy goals, e.g., ensuring a minimum level of service coverage.

## 2.4        Cumulative opportunity measures

Cumulative measures estimate the number of opportunities within a given travel time distance of each spatial unit. For example, measures of the number of jobs accessible by transit within 30-, 45-, or 60-minute thresholds are the most common measures of transportation accessibility used in regional planning (Boisjoly & El-Geneidy, 2017). These measures benefit from their ease of interpretation and development (Geurs & van Wee, 2004), as they only require calculating origin-destination travel times and summing the destinations of OD times below the cut-off time. However, they include several drawbacks. First, they do not account for the relatively lower attractiveness of destinations that are further away, failing to reflect how distance influences travel choices, as illustrated in Figure 1. Second, the choice of which temporal threshold to use, e.g., 30 versus 45 minutes, can bias results, as documented in recent case studies (Pereira, 2019; Xi et al., 2018). These limitations have led researchers to embrace the next group of measures considered.

## 2.5        Gravity based measures

These metrics also require the calculation of origin-destination travel times but instead of assigning an arbitrary cut-off time, the analyst selects a function that weights the number of opportunities at each destination based on their distance from the origin.

$$A_i = \sum_j O_j f(t_{ij}) \tag{1}$$

Equation 1 is a generic example of a gravity measure in which the zone i's accessibility ($A_i$) is the sum of the size (or other measure of the attractiveness) of opportunities at each zone j ($O_j$), weighted by an "impedance function", $f(t_{ij})$, where f decreases in value as the distance or time between zone i and zone j ($t_{ij}$) increases.

   The impedance function weights closer opportunities as greater in value than opportunities further away and ideally is calibrated using observed travel behavior from the study area. When such data are unavailable, analysts can draw upon functions established in the literature as reasonable predictors of

travel behavior. For example, Kwan (1998) details several families of impedance functions including the particular specifications of the inverse power (POW), modified Gaussian (MGAUS), negative exponential (EXP), cumulative rectangular (CUMR), and cumulative linear (CUML) functions illustrated in Figure 1 from Higgins (2019).
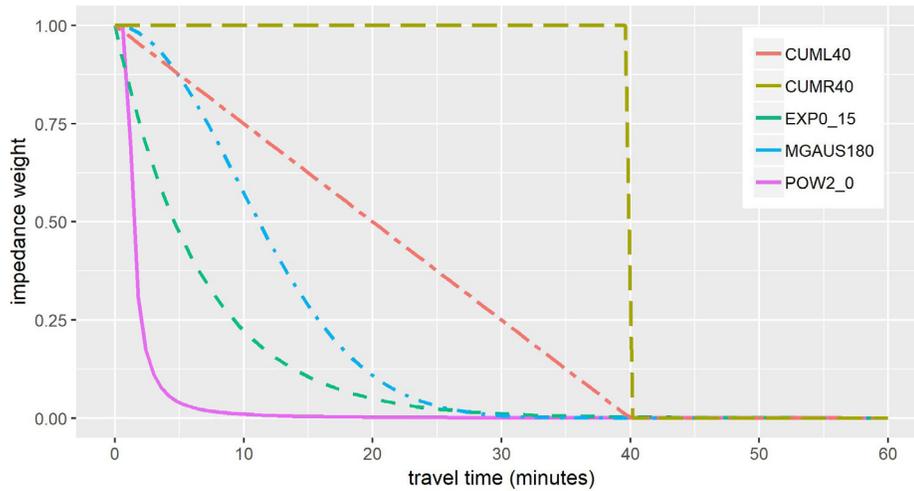


**Figure 1**. Example impedance functions

Note that the cumulative rectangular opportunities measure is a special case of Equation 1 in which $f(t_{ij})$ = 1 for all $t_{ij}$ less than or equal to the threshold travel time. Continuous gravity measures do not suffer from the drawbacks of simpler cumulative distance measures; however, the choice of impedance means they are more challenging to specify and more difficult to interpret as the output values are not straightforward counts of reachable destinations.

Transportation and planning agencies across the continent have applied place-based accessibility measures to establish policy goals and predict a broad range of likely impacts from their investments and plans. We provide a recent example of each type of place-based metric deployed in project or plan evaluation in Table 1. These examples reflect the growing number of essential activities that planners seek to improve access to when devising plans and evaluating projects and suggest that accessibility calculation is an increasingly important part of the transportation planning process. In response, the remainder of this paper seeks to better understand how accessibility calculations can be shaped by the data and algorithms employed by some of the most popular software packages available to analysts for large-scale multimodal routing operations.

**Table 1.** Recent use of place-based accessibility measures by planning agencies

| Project/Plan Evaluated | Policy Area | Evaluation Metric | Type | Outcome/Impact |
|---|---|---|---|---|
| Ontario Line, Toronto, Canada (Farber & Allen, 2019) | Economic opportunity, education, equity | Percent change in distance-decay weighted job and educational opportunities accessible to both the total population and disadvantaged households | Gravity-model | The analysis demonstrated that the Ontario Line will benefit disadvantaged households within the City of Toronto. |
| Mobility 2045, Dallas, USA (North Central Texas Council of Governments, 2018) | Public health, economic opportunity | Measures impact of the plan on:<br>• Shares of the population within 15 minutes of a hospital, by transit and auto<br>• Shares within 30 minutes of regional shopping generators, by transit and auto | Proximity or threshold | NCTCOG adopted a final Regional Transportation Plan that ensured protected populations did not experience a disproportionate loss of hospital and shopping access due to projected increases in congestion. |
| Regional Transportation Plan, Portland, USA (Metro, 2018) | Education, health, community services | Measures impact of the plan on:<br>• The change in the number of "community places" reachable by auto (20 min), transit (30 min), and active travel (20 min). Community places included schools, childcare, healthcare, religious orgs and banks.<br>• The number of low wage and medium wage jobs reachable by auto (30 min), transit (45 min), bicycle (30 min), and walking (20 min). | Cumulative opportunities | METRO's proposed 10-year constrained investment strategy is estimated to increase the number of community places the average household can access by transit from 78 to 100, and by car from 33 to 57. The agency highlights this evidence in justifying its proposal. |
| Transport 2025, London, UK (TfL, 2006) | Employment | Measures impact of the plan on the change in the number of jobs accessible by public transit within 45 minutes | Cumulative opportunities | The justification of the adopted scenario highlights that it will increase accessibility across the population by an average of 25%. |
| Analysis of impacts of population growth on service access, British Columbia, Canada (Mayaud, et al., 2019) | Education and health services | Competition-based accessibility to walk-in clinics, hospitals and schools in Surrey, B.C. Analysts accounted for competition from other residents using the 2016 Census for a baseline. They replicated the analysis under a 2022 population growth scenario and compared differences in scores to track loss of access. | Competitive | The analysis highlights areas where residents are likely to experience lower access to services due to population growth, helping the city identify where to locate future services. |
| Southern California Association of 2016 Regional Transportation Plan, Southern California, USA (SCAG, 2016) | Parks and open-space | Distance decay weighted share of the region's parks and open space reachable within 45 minutes by any transit, 45 minutes by local bus, and 30 minutes by auto. Tested how a proposed Regional Transportation Plan improved these metrics for different socio-demographic groups. | Gravity | The agency's equity summary highlights that the adopted plan will improve residents' access to parks and open space, with benefits accruing to all demographic groups. |
| The Atlanta Region's Plan (ARC, 2019) | Entry level employment | The analysis models the impact of the region's proposed regional transportation plan on the number of low-wage jobs accessible by transit within 60 minutes for disadvantaged communities identified by the agency as "Equity Target Areas" (ETAs). | Cumulative Opportunities | In justifying the final plan, the agency highlights that it significantly increases jobs accessible by transit from ETAs. |

## 3        Methods and analysis

To review tools available for the calculation of place-based accessibility, this study adopts a two-step methodology. The first step involves a preliminary appraisal of five popular place-based accessibility tools according to several criteria. The second step consists of evaluating the relative performance of the tools using a case study of transit accessibility in the City of Toronto.

### 3.1        Preliminary identification and appraisal of accessibility tools

#### 3.1.1        Evaluation Criteria: Technical Dimensions of Place-based Accessibility Measures

The place-based accessibility metrics described above can be further refined to better account for other influences based on our judgement of the state-of-the-art in accessibility research. We provide an overview of recent improvements to these metrics here and utilize them to conduct an appraisal of the five accessibility tools.

*Multi-modal Routing:* Modeling the shortest path between an origin and destination for walking or driving involves minimizing travel time or distance across a network of edges and junctions using a solver or router (see Bast et al. (2016) for an in-depth review of different solvers). In this case, a key consideration is the ability to solve multi-modal network routing problems (e.g., walking and transit, cycling, driving) that reflect the travel options and complex networks of cities. One popular solution for modelling transit networks is through General Transit Feed Specification (GTFS) files, which consist of a system of text files corresponding to transit trips, service calendars, route geometry, and other attributes. While the GTFS system can be cumbersome, it appears to have become a de-facto data standard for many transit agencies around the world to share their network information and is currently the foundation of many of the available tools for analyzing transit accessibility.

*Traveler perception factors:* Travelers do not perceive the time costs of different components of their travel in the same way. Using the example of transit, travelers perceive walking to stations (access), waiting times at stations, transfer times and waits, and travel to destinations (egress), differently than time spent in transit vehicles (Iseki & Taylor, 2009; Tilhaun & Li, 2015; Wardman, 2004). To reflect this, analysts often apply factors that re-weight these times to reflect these perceptions. This requires analysis tools that can accurately estimate these components of the transit trip and weight them accordingly.

*Generalized Cost:* Costs associated with transportation, like fares and fuel, can also limit accessibility. Analysts can devise place-based accessibility measures accounting for these costs by constructing generalized cost metrics that convert travel times into monetary costs using a value-of-time metric and adding this to the monetary costs of travel (El-Geneidy et al., 2016). Analysts can then estimate the number of destinations accessible within generalized cost thresholds specified in dollar amounts. This process can also work in reverse, with analysts converting transit fares or vehicle costs into an equivalent time cost that is added onto travel times. The incorporation of monetary costs is important in contexts where transit fares vary by zone, or where roads are tolled. Analysts can also combine generalized costs across different modes to calculate the maximum accessibility from a given point at varying cost or time levels (Ford et al., 2015). Beyond generalized costs, recent research by Conway and Stewart (2019) utilizes Pareto sets of paths that jointly minimize fare costs and travel time in route selection and accessibility calculation.

*Impedance functions:* As discussed above, analysts commonly weight the value placed on opportunities by the travel time required to reach them using an impedance function. For this criteria, tools are evaluated based on whether they directly implement the concept of impedance.

*Competition:* For some destinations, such as hospitals or employment locations, accessibility mea-

sures may be biased if they do not account for the fact that many people may be competing for a limited number of opportunities at those destinations. Examples of this include the relationship between the demand for and supply of hospital beds (Páez et al., 2019) or job openings (Merlin & Hu, 2017). If this competition is not accounted for, accessibility may be over-stated. Recent research suggests that competitive job accessibility measures better predict individuals' employment outcomes than non-competitive metrics (Merlin & Hu, 2017). Analysts have three options for accounting for competition. The first option controls for the competition of other travelers at origins (Geurs & van Wee, 2004). The second option, in contrast, controls for competition at destinations (Shen, 1998). A third approach accounts for competition at both origins and destinations simultaneously (Merlin & Hu, 2017). Each of these approaches adds a level of complexity that makes resulting values less intuitive to policymakers, however.

*Flexibility of inputs/outputs:* Analysts often need to engage in customized workflows involving importing and exporting data in a variety of file formats. For this analysis, a key consideration relates to the flexibility of inputs such as different transportation networks and files (e.g., road network shapefiles, OSM data) and outputs including the ability to export non-proprietary file formats such as travel time matrices and accessibility results in comma-separated values (CSV) files for analysis.

*Temporal variation:* Analysts building place-based access measures must select a departure time, or range of departure times, when building accessibility metrics. Morning and/or afternoon peak periods are often chosen for analysis, given the importance of these periods for work and school commuting, as well as being the periods of peak congestion in the system. Alternatively, to construct a measure that accounts for temporal variability over the entire day, average accessibilities can be calculated from departure times repeated at different intervals (Farber et al., 2016; Owen & Levinson, 2014). Others suggest calculating the OD travel times at one-minute intervals and using the median time for each OD pair to determine if the destination is reachable (e.g., if that median is below the cumulative accessibility cut-off time) (Conway et al., 2018). Selecting a range of departure times can help to offset the effects of the Modifiable Temporal Unit Problem in accessibility analysis (Pereira, 2019) and the software tools will be evaluated based on whether they offer a built-in way to perform this type of analysis.

*Spatial refinements:* Finally, the spatial aggregation of accessibility origins and destinations into zones can introduce bias into model results. This is known as the Modifiable Areal Unit Problem in geography (Openshaw, 1984). For example, reliance on spatial unit centroids for zonal geographies implicitly assumes that accessibility for residents within the zone will be constant over the area of the zone when in reality it will vary spatially. This can also occur when aggregating more detailed address-level data to zonal boundaries. Analysts can respond to these biases in several ways. First, increases in computing power mean that disaggregate origins and destinations can be used for accessibility analysis, although such analysis remains very computationally intensive when conducted at the city or metropolitan scale. Second, instead of aggregating address-level destination data to a geography like the Dissemination Area/Block or Census Tract, they can divide a region into a regular grid of small square or hexagonal cells and assign address-level data to the resulting grid centroids, preserving detailed destination data and minimizing spatial bias (Pereira, 2019). Third, analysts with access to detailed land-use data can utilize dasymetric mapping to assign populations and opportunities in an aggregated geography to smaller sub-geographies (Hu & Downs, 2019). Tools will be evaluated based on whether they natively support any of these strategies.

### 3.1.2 Assessment of place-based accessibility tools

In line with the increase in the number of open-source tools available for transportation planning (Lovelace, 2021), there are a variety of routing engines that can be used for accessibility analysis. This includes unimodal (e.g., 'gtfsrouter' (Padgham et al., 2021)) and multimodal (e.g., 'UrbanAccess'

(Blanchard & Waddell, 2017)), 'osrm' (Huber & Rust, 2016), GraphHopper, and the 'access' (Saxon et al., 2021) package for Python) free and open-source routing tools. Other tools are paid (e.g., TransCAD, Conveyal Analysis, CUBE Access) or offer "freemium" tiers (e.g., the 'hereR' (Unterfinger & Possenriede, 2021) package's R interface to HERE's routing API or the 'googleway' (Cooley & Barcelos, 2020) package's R interface to the Google Maps routing API).

Several constraints limited our ability to consider all these tools. First, several were not publicly available when this analysis began in December 2019. Software compatibility with local travel demand model workflows and the project team's time and resource constraints further narrowed our scope. These constraints, combined with our focus on transit and multimodal routing, led us to preliminarily evaluate five tools: *ArcGIS Pro*, *Emme 4*, implementations of *OpenTripPlanner* in *R* and *Python*, and the *R5R* package for R. Some of these tools are proprietary commercial products while others are open source. For some of these tools, analysts have also developed open-source packages and add-ons that support accessibility analysis that we detail below.

We assessed each for its ability to address the technical dimensions of place-based accessibility described earlier. We considered how an analyst might adapt their workflow if the software lacked a particular function and documented the flexibility of their inputs and outputs—e.g., can the tool easily input and output data as comma-separated values (.csv) files and GIS shapefiles. In each category, we gave the tool a score that ranged from zero to two (0-2). A two (2) means the tool can compute a given function internally, requiring no outside software inputs or outputs. A one (1) means the tool can accommodate the function as part of a larger workflow. One can also mean that consultants or academics have built plug-ins or add-ons enabling the tool to complete the given task. A score of zero (0) means that the tool cannot assist with a given task due to compatibility issues with other software, or due to some other issue. These results are summarized in Table 2 and discussed below. No tool received a perfect score in this appraisal. More detailed justifications can be found in Higgins et al. (2020).

ArcGIS's Network Analyst is a component of ArcGIS and enables users to generate a transportation network using shapefiles and GTFS data. Analysts can execute a range of functions with the resulting network dataset in the ArcGIS environment via Python script or GUIs that allow for detailed specification of network attributes and considerations. Network Analyst possesses tools that can account for spatial, temporal, and perception factors. However, researchers have needed to develop additional add-ons to generate generalized cost measures with ArcGIS (Ford et al., 2015) and to incorporate gravity-measure impedance functions into ArcGIS workflows (Higgins, 2019).

**Table 2**. Assessment of place-based accessibility measurement tools

|  | ArcGIS Network Analyst | Emme 4 | R-OTP | Python-OTP | R5R |
|---|---|---|---|---|---|
| *Multi-modal routing* | 2 | 2 | 2 | 2 | 2 |
| *Perception factors* | 2 | 2 | 2 | 2 | 2 |
| *Generalised costs* | 1 | 2 | 2 | 2 | 1 |
| *Impedance functions* | 1 | 1 | 1 | 1 | 2 |
| *Competition* | 1 | 1 | 1 | 1 | 1 |
| *Flexibility of inputs/ outputs* | 2 | 1 | 1 | 1 | 1 |
| *Temporal variation* | 1 | 1 | 1 | 1 | 2 |
| *Spatial refinements* | 2 | 2 | 2 | 2 | 2 |
| *Total* | 12 | 12 | 12 | 12 | 13 |

Emme is a multimodal transportation forecasting software used for zonal-aggregate travel demand modelling. The latest version of Emme allows computation of accessibility in terms of generalized cost measures (INRO, 2019). The tool is sensitive to perception factors and can be deployed to measure travel at different times of day and at different spatial scales. Emme does not allow for direct in-software computation of cumulative opportunity, competition, or gravity measures. Emme is included in this analysis as being representative of commercially available transportation network modelling packages and was chosen since it is used extensively within the Toronto case study region.

Open Trip Planner (OTP) is an open-source multimodal trip planner that was originally developed for TriMet in Portland, OR. The OTP tools utilize OpenStreetMap (OSM) for their underlying street networks and employ a Java runtime environment for their routing backend that can be utilized for larger analyses through additional tools such as R and Python. The ubiquity of their data inputs makes OTP implementations extremely flexible. They can calculate travel times at various scales and repeatedly over multiple time periods. They allow easy application of time perception factors and rely on publicly available, standardized data. As of version 1, OTP uses an algorithm called A-star for routing walk and cycle trips while any trips that involve transit use a mix of A-star and the Tung-Chew heuristic (OTP, 2020). Interfacing with OTP is facilitated through the *opentripplanner* package for R (Morgan et al., 2019) and in Python through the use of a customized script based on Allen (2020) and Pereira et al. (2019).

It is important to note that OTP's original purpose was to provide a way for suggesting itineraries for transit riders based on the optimization of several different trip cost parameters and the platform has been steadily updated based on feedback on transit rider's travel preferences from transit agencies and other stakeholders (Stewart, 2020). As such, the algorithm is designed to minimize generalized costs, which include travel costs but also factors like interchange penalties and may not return paths that minimize travel time alone.

Another consideration unique to OTP relates to how travel time is minimized if one utilizes options for "initial wait time clamping". This feature is designed to make OTP route selections less sensitive to the stated departure time by allowing the analyst to account for how travelers will adjust their departure time to align with the arrival of their first transit vehicle. For example, if the analyst has selected an 8:00AM departure time but the first bus for an individual arrives on the network at 8:30AM, the calculation of the trip duration will not begin until the individual needs to leave to catch the bus. In this case, the initial wait clamp parameter can be used to subtract initial waiting times up to some limit for the first transit trip from the total trip duration since the start time.

More accessibility analysis-focused functions related to the explicit minimization of travel time have been developed within OTP1 but were not utilized by default due to inefficiencies in the underlying router design (Stewart, 2020). Instead, analysis functions are better implemented within the newer R[5] tool discussed below. The second-generation OTP2 currently under development utilizes some of R[5]'s routing innovations but is focused on producing itineraries for transit users and does not contain any analysis functions (OTP, 2020).

The R[5] routing engine uses OSM and GTFS data to build a network graph for a region and a server infrastructure for parallelizing workflows and rapid accessibility calculation. R[5] utilizes the RAPTOR and McRAPTOR routing algorithms (Pereira et al., 2021) to find the earliest arrival path that minimizes travel time from departure. While not tested here, R[5]'s use of range-RAPTOR is designed to enable "profile routing" that finds the shortest overall trip duration over a departure time window compared to the more traditional approach employed in other routers like OTP that aggregate a large number of individual start time queries (Conway et al., 2018; Pereira et al., 2021). R[5] is the core of Conveyal Analysis (Conway et al., 2017; Conway et al., 2018), which constitutes an entire multi-modal scenario planning platform that grew out of OTP and enables the analyst to not only model place-based accessibility from

existing GTFS and street networks, but to alter the transit network and run comparative analyses. The R5R package (Pereira et al., 2021) provides an R interface to the R$^5$ solver.

## 3.2    Comparative analysis

### 3.2.1    Study design

To evaluate the strengths and weaknesses of ArcGIS Pro, Emme, R5R, and OTP using R (R-OTP) and Python (Python-OTP) in terms of computational performance and accessibility results, we present a general algorithm – a term we loosely define as a sequence of operations for solving a problem – for accessibility analysis. The calculation of accessibility scores consists of several steps or sub-routines: assembling inputs such as road networks and GTFS files, creating network objects for analysis, determining shortest paths using different route solvers, and calculating accessibility outputs.

1.  Creating a combined street and transit network for analysis.
    The first step is to gather the required input files to assemble routable multi-modal networks. This includes street networks and GTFS transit schedules.

2.  Calculating origin-destination (OD) matrices for trips.
    In the second step, each tool calculates the shortest path between origin-destination pairs using a solver or routing algorithm. Origins and destinations are generated as random points within the boundary of the City of Toronto using ArcGIS Pro, consisting of 100, 1,000, 10,000, and 100,000 points. In addition, a single input point was manually created at the approximate center of Toronto City Hall. From these points, the analyses are structured around a "local" single-origin-based scenario that estimates an OD matrix from the single City Hall point to all other points and a "regional" scenario that estimates a full OD matrix for all origin-destinations pairs.

3.  Calculating accessibility for each origin place.
    The final step is to utilize the returned OD matrices to calculate place-based accessibility for the origin points per Equation 1. This typically involves joining the opportunities available at the destinations to the computed origin-destination cost matrix through a common ID field, multiplying the destination opportunities by some weighting function based on the travel time involved in reaching them from the origins, and then summing these weighted values over the origins. To simplify this calculation in the present case, the "opportunities" available at each destination point are set equal to 1 and we utilize a cumulative accessibility measure where the accessibility for a given origin point is the sum of all points reachable within 45 minutes:

$$f(t_{ij}) = \begin{cases} 1, & \text{if } t_{ij} \leq 45 \\ 0 & \text{if } t_{ij} > 45 \end{cases} \qquad (2)$$

    While such cumulative measures tend to be an arbitrary simplification of actual travel behavior compared to the more continuous weighting functions of gravity specifications, it allows the project team to implement a simple accessibility use case in its code across the packages. Implementing other impedance functions is not expected to meaningfully affect compute time. Accessibility calculation is completed in ArcGIS Pro for that tool while R was used to calculate accessibility for R-OTP and R5R after running their respective analysis and for Emme 4 and Python-OTP after reading in the output OD matrix files in .csv format.

### 3.2.2    Hardware configuration

The computer selected for this runtime analysis is a server hosted by the University of Toronto's Transportation Research Institute. This server is configured with an Intel Xeon E7-8890 v3 CPU with 18 cores/36 threads running at a base/turbo frequency of 2.5GHz/3.3GHz with 64GB of memory. Versions for the individual software packages are listed below:

- ArcGIS Pro version 2.4.2
- Emme version 4.4
- R-OTP uses:
  - OTP version 1.4.0
  - Java 8 Runtime Environment (64-bit)
  - R version 3.6
  - RStudio Desktop 1.2.5033
- Python-OTP uses:
  - OTP version 1.4.0
  - Java 8 Runtime Environment (64-bit)
  - Python 2.7 with built-in IDLE for code
  - Windows command prompt
  - Jython 2.7.1 for linking Python with Java
- R5R uses:
  - R5 version 6.2
  - R5R version 0.3-4
  - Java 11 Runtime Environment (64-bit)
  - R version 3.6
  - RStudio Desktop 1.2.5033

### 3.2.3    Analysis settings

Because GTFS transit networks are schedule-based, a (pre-COVID) start time/date of Monday December 30, 2019 8:00:00 AM was used for the ArcGIS Pro, OTP, and R5R implementations. The maximum travel time used for these tools was 300 minutes. Both OTP and R5R support an additional maximum walking distance parameter to influence access and egress to transit; in both cases this was set to 5,000m, which is a distance greater than people tend to travel to/from transit stations in Toronto (Crowley et al., 2009). On the other hand, ArcGIS Pro and Emme have no limit on walking distances. Default walk speeds in OTP, R5R, and Emme differ slightly but for comparability, we alter all walk speeds to 5kph and use the same value in the creation of the ArcGIS Pro walking network. For OTP, we set the initial wait clamp to zero so that the trip duration is calculated from 8:00AM inclusive of waiting.

For Emme, a custom GTFS-based transit network can be created within the program itself based on original GTFS packages. Such custom networks can utilize factors such as traffic conditions and are more general in terms of temporal dependence. For the Emme workflow, a typical morning-peak period (6AM to 9AM) in 2016 was assumed based on the available data (e.g., travel demand from the Transportation Tomorrow Survey). Because Emme can only perform a one-hour assignment to get an OD travel time matrix, a peak-hour factor was applied to the peak-period demand to get the one-hour demand. In this sense, the network inputs and output from Emme 4 are more reflective of surface traffic conditions but are not directly comparable with the other tools.

### 3.2.4    Tool configuration

Although we have taken care to make the workflows comparable across tools, each of the five tools implements the general accessibility algorithm in a different way. In this sense, the idiosyncrasies of each offer different parameters that can customize the analysis but necessarily impact results.

*ArcGIS Pro*: The 2.4 version of ArcGIS Pro offers an updated Network Analyst that stores the network and associated input-output layers in-memory, which can speed up a network analysis. However, the core Network Analyst extensions are natively single-threaded, which greatly limits the performance of the software in a large-scale analysis workflow. ArcGIS Pro's Python backend does accept custom multi-process workflows and this research utilizes the toolbox published in Higgins (2019). For some comparability, an OSM network is also used as the underlying network for the ArcGIS scenario; in this case, a walking network was collected through the OSMnx tool (Boeing, 2017), exported as a shapefile, and used as an input alongside the GTFS files in the creation of a transit and walking network.

*Emme 4.4:* Emme itself is a product of INRO Consultants Inc. and a TMG Toolbox has been developed by the Travel Modeling Group (TMG) at the University of Toronto to facilitate advanced analyses of travel demand forecasting models implemented in Emme. Both auto and transit assignments need to be performed in Emme to estimate auto and transit OD travel times. An existing GTAModel (Miller et al., 2015) Emme network that contains both road and transit information was used instead of importing the GTFS and OSM files directly due to restrictions on the number of links supported by the current software license. The transit network was previously built from the 2016 GTFS files, which should generally be similar to the current 2019 GTFS files used by the other tools. However, because Emme is built on top of an existing traffic network, it supports routing based on congested transit travel times for surface routes. On the other hand, the road network only consists of major streets across the region, which is a sparser network as compared to the OSM networks above. A "shortest path" tool that is newly available in Emme 4.4 was utilized to generate the auto OD matrices, and the eXtensible Travel Modelling Framework (XTMF) travel model system built by TMG (Miller et al., 2015) was used to generate the transit OD matrices. One limitation of this comparison concerns Emme's ability to estimate OD matrices for transit in the single-origin scenario (e.g., a 1 x 100 OD matrix). The Level-of-Services (LOS) for transit can only be extracted at the centroid level in Emme, therefore, the points are assigned as new centroids in the network and a transit assignment is performed to obtain the travel time between the 2,298 centroids/points. Due to the nature of the algorithm used in the transit assignment, only full OD matrices can currently be generated (e.g., 100 x 100 OD matrix). Essentially, this means a transit OD matrix cannot be computed for the local scenario.

*OpenTripPlanner*: To set up the OTP tool for the runtime analysis, the team first collected a metropolitan-level OSM extract for Toronto using the free extract service from interline.io. Alongside the relevant GTFS files, an OTP server was created for use with R and Python. The opentripplanner package (Morgan et al., 2019) allows the analyst to build and set up an OTP network graph using OSM and GTFS data and connects it to R as a local server. From this, HTTP requests are made to the OTP Java server and it returns trip itineraries for the given modal options. For Python, the implementation is less streamlined than R-OTP. Commands are used to call an OTP Java executable to build an OTP graph referencing relevant GTFS and OSM files and Java. A Python script (e.g., Allen, 2020; or Pereira et al., 2019) can also be used to load the graph, parameterize the router, and obtain origin-destination information using OTP's scripting API library. Both R-OTP and Python-OTP allow for multiprocessing. In any implementation of OTP, the analyst can customize the routing algorithm in a number of ways. For example, trip modes can include many combinations of walking, cycling, transit, and driving. Wait and walk reluctance weighting parameters can also be set alongside the maximum distance individuals are expected to walk when using transit. OTP also has the option of returning several trip itineraries

between a given origin and destination pair on a set time and date and can account for flexibility in a trip's start time (in this case we minimize costs from the start time).

*R5R*: With their similar routing foundations, R5R uses many of the same inputs as the OTP tools above, including an OSM network for streets, GTFS files for transit routing, and Java for the R$^5$ routing server (Pereira et al., 2021). Like OTP, R5R offers options for multimodal routing as well as a maximum walking distance parameter to control aspects of access and egress behavior for trips. R5R supports returning several trip itineraries over a time window and utilizes multiprocessing in the computation of results. For computation, this work utilizes a custom R Notebook available via Higgins (2019) that employs a batching algorithm that subsets the origins into smaller batches for routing in the larger regional analyses.

### 3.2.5  Performance results: Local analysis

Results for the local analysis scenario are presented in Table 2. For ArcGIS and Emme, there are small computational time penalties associated with calculating or snapping the location of input points on the network. These are one-time costs if the same input points are re-used. Comparable network location time costs are not reported in the R/Python OTP and R5R tools and cannot be separated out from the larger solve time. They also do not carry over to future analyses and must be re-calculated each time.

For calculating the OD matrix, increasing the number of destinations dramatically reduces the performance of the R-OTP package. The other software packages perform much better with R5R exhibiting the best performance followed by Python-OTP and ArcGIS Pro. In ArcGIS Pro, the custom multiprocessing script is based on batching origins; with only a single origin in the local scenario, this workflow is limited to a single thread. Computational time involved in calculating and summarizing accessibility for the local scenario is rapid for all tools once the OD matrix calculation is finished. For Emme, the larger 1 x 100,000 analysis exceeds the license available at the university, and as such was not run. Note that point-based calculations within Emme are only available within command-line operations and so are not a feasible proposition for any automated procedure. Thus, all Emme-based calculations must occur at the traffic zone level of spatial aggregation.

**Table 3.** Performance analysis results

| Activity | Scenario | ArcGIS Pro | R-OTP | Python-OTP | R5R | Emme |
|---|---|---|---|---|---|---|
| Make/Build Network | All | 11m 26s | 8m 19s | 7m 24s | 2m 39s | 1m 25s |
| Calculate OD Matrix and Accessibility | 1 x 100 | 1m 37s | 20s | 23s | <1s | 2m 23s |
| | 1 x 1,000 | 1m 47s | 2m 49s | 24s | <1s | 4m 46s |
| | 1 x 10,000 | 2m 51s | 22m 37s | 28s | 1s | 28m 36s |
| | 1 x 100,000 | 12m 59s | 7h 12m 15s | 1m 4s | 7s | *not run* |
| Calculate OD Matrix and Accessibility | 100 x 100 | 3m 24s | 25m 30s | 2m 2s | 1s | 5m 5s |
| | 1,000 x 1,000 | 6m 56s | >24h (DNF) | 13m 42s | 9s | 18m 33s |
| | 10,000 x 10,000 | 49m 53s | *not run* | 4h 49m 18s | 7m 50s | 13h+ |
| | 100,000 x 100,000 | *not run* | *not run* | *not run* | 18h 15m 43s | *not run* |

### 3.2.6    Performance results: Regional analysis

Results for the full *n* x *n* regional analysis matrices are presented in the bottom half of Table 3. In line with the increase in OD size, each of the tools now exhibits greater computational time. Interestingly, the computational time for the 100 x 100 scenario is greater than that for the 1 x 10,000 scenario above, which we attribute to the extra time involved in finding the shortest street network and transit schedule paths between points in a much more geographically varied sample. This type of analysis magnifies some of the shortcomings associated with each workflow. First, despite making use of the multiprocessing resources of the PC, the R-OTP implementation exhibited the worst performance and was not able to finish the 1,000 x 1,000 scenario within 24 hours. This is likely due to R-OTP making one-to-one routing calls to underlying OTP servers compared to a shortest-path tree approach to calculating travel times to destinations utilized in the Python-OTP code. The Python-OTP and Emme tools show high performance for large accessibility calculations but generally finish the 10,000 x 10,000 analysis after many hours. In contrast, the ArcGIS Pro tool finishes the same accessibility analysis in roughly 50 minutes while the R5R tool finishes the analysis in under 8 minutes, suggesting its one-to-many routing solver scales very well to large transit analysis tasks. With a calculation time of 18 hours and 15 minutes, the R5R tool was the only one to complete the 100,000 x 100,000 scenario in under 24 hours.

### 3.2.7    Routing results: 1,000 x 1,000 analysis

To compare each tool's routing results, we report descriptive statistics for the travel time results for each tool for the 1,000 x 1,000 scenario (Table 4) and perform a scatterplot comparison (Figure 2). OTP results are reported from the Python-OTP tool only. While the correlations in Figure 2 suggest the computed travel times are generally consistent across tools, there are interesting differences and trends in the travel times for origin-destination pairs. ArcGIS Pro tends to estimate shorter trips than all the other tools. While built on a similar underlying OSM network with the same GTFS files, the ArcGIS results have a mean travel time of approximately 70 minutes with a maximum trip duration of 211 minutes. In contrast, mean travel times from R5R and OTP are 6 min and 7 min longer respectively while Emme returns the longest mean travel time of 89 min. The tails of the other packages are longer, with R5R and Emme returning maximum durations of approximately 285 min.

**Table 4.** Descriptive statistics: OD matrix travel times

|              | ArcGIS Pro | Emme     | OTP           | R5R           |
|--------------|-----------:|---------:|--------------:|--------------:|
| **Minimum**      | 0.0        | 0.0      | 0.0           | 0.0           |
| **1st Quartile** | 50.2       | 62.9     | 55.0          | 54.0          |
| **Median**       | 68.0       | 86.4     | 75.2          | 73.0          |
| **Mean**         | 69.6       | 89.0     | 77.8          | 75.6          |
| **3rd Quartile** | 87.3       | 111.4    | 97.9          | 95.0          |
| **Maximum**      | 211.1      | 285.8    | 251.3         | 285.0         |
| **Missing (%)**  | 0 (0.0%)   | 0 (0.0%) | 27,768 (2.8%) | 28,773 (2.9%) |

These factors suggest an overall tendency towards shorter travel times for the same trips in ArcGIS Pro. Because the transit solver algorithm is closed-source, the reasons for these differences are not clear. At the high end, the longer travel times from Emme are no doubt in part due to the congested transit network, which arguably offers more realistic travel times for surface transit and its sparse street network

that alters the location and "last mile" connection between different origins and destinations. In contrast, both OTP and R5R generally fall in the middle of the estimated travel time statistics. However, both packages do not return a full n x n OD matrix with no connections between roughly 28,000 origin-destination pairs. Reasons for this require further investigation, but with each tool using OSM data for their underlying street network, the similar results could be attributable to disconnected links or "islands" (while ArcGIS Pro uses the same OSM data source, the OSMnx package pre-processes the data to clean up nodes and links and network topology).
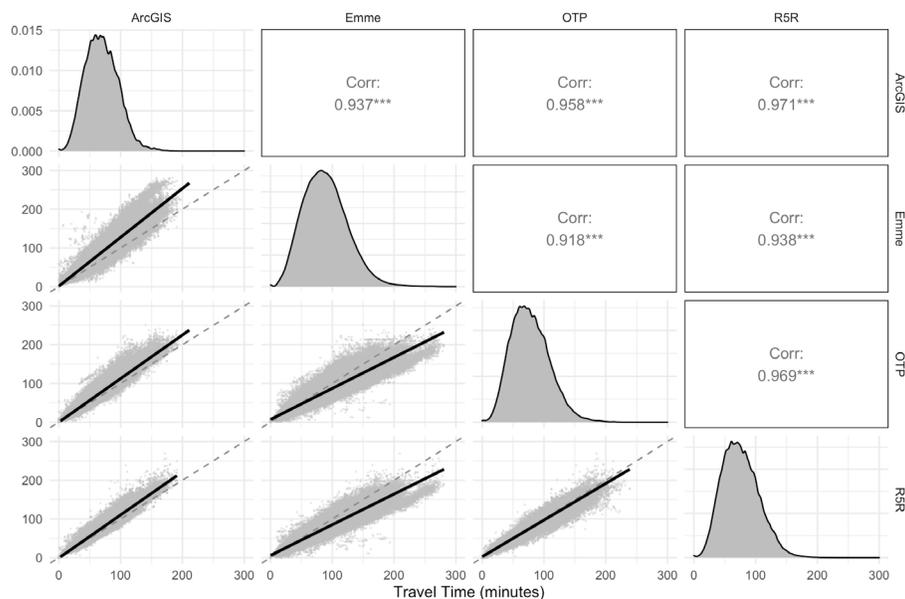


**Figure 2.** Scatterplots of calculated travel times

### 3.2.8    Accessibility results: 1,000 x 1,000 analysis

Considering the differences in travel time in the full OD matrix, it is not surprising that each of the tools returns differences in accessibility using the binary impedance function with a 45 min cut-off (Table 5). With its tendency towards shorter travel times, the ArcGIS tool returns the highest mean accessibilities with points able to reach an average of 188 destinations. Similarly, the Emme analysis returns the lowest levels of accessibility overall followed by OTP and R5R.

**Table 5.** Descriptive statistics: Accessibility

|  | **ArcGIS Pro** | **Emme** | **OTP** | **R5R** |
|---|---|---|---|---|
| **Minimum** | 1 | 1 | 0 | 1 |
| **1st Quartile** | 139 | 69 | 103 | 112 |
| **Median** | 192 | 97 | 140 | 153 |
| **Mean** | 188 | 100 | 139 | 153 |
| **3rd Quartile** | 237 | 127 | 179 | 196 |
| **Maximum** | 457 | 314 | 316 | 348 |

The spatial distributions of accessibility results are shown in Figure 3. In this case, the individual points are first converted to Voronoi polygons and the city's rapid transit network is layered on top for visualization. The general pattern across all maps is one with higher levels of accessibility around rapid transit lines and along the major bus routes that they connect to. Still, compared to more traditional accessibility analyses that consider opportunities like the number of people or jobs in a zone, it is important to note that the points were generated randomly and may not be areas traditionally considered as origins or destinations for transit service.
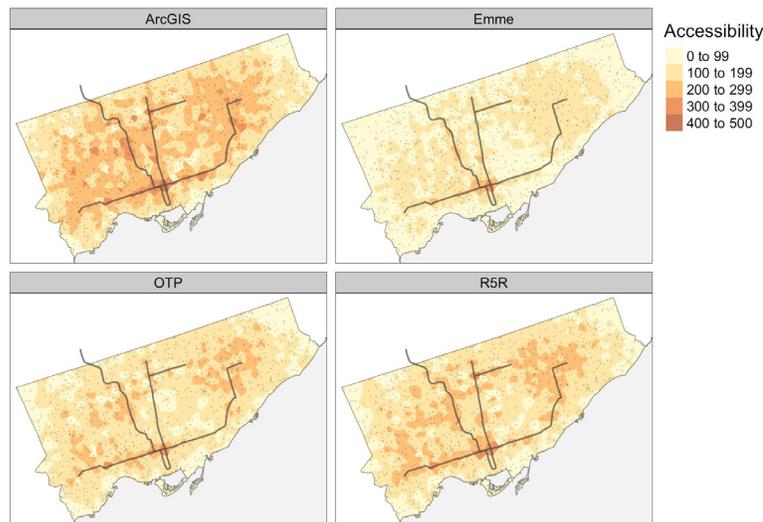


**Figure 3.** Accessibility comparison

To examine whether there are any spatial patterns of systemic bias across the tools, Figure 4 plots the differences in accessibility for the Voronoi polygons. Beyond the differences associated with ArcGIS and Emme producing the highest and lowest accessibilities respectively, there appears to be evidence of some spatial clustering in the ways each tool calculates travel times and accessibilities. After defining a spatial weights matrix based on queen contiguity for the Voronoi polygons, Moran's I tests confirm statistically significant moderate positive spatial autocorrelation in the calculated accessibility differences. To better highlight these clusters, Figure 5 displays the Voronoi polygons classified using local indicators of spatial association (Anselin, 1995) with p≤0.05 and without correcting for multiple testing. The results reveal some interesting patterns in the calculated accessibility differences. The most consistent result appears to be the calculation of lower accessibilities by OTP along the Yonge Street portion of the city's north-south subway line compared to the other tools. While such findings require further investigation, one reason could be the way OTP selects routes by minimizing generalized costs rather than explicitly minimizing travel time that could result in fewer transfers being made.

Compared to the other tools, Emme appears to calculate lower accessibilities in the central-west part of the city, perhaps because the Emme tool is utilizing information on estimated levels of traffic congestion in its calculation of surface transit trips. Emme also tends to estimate higher accessibilities for a cluster of locations in the city's north-east, an area that is home to the Rouge National Urban Park and Toronto Zoo. Higher results in this case can likely be attributed to the road network used in the Emme analysis that simplifies the streets and their connections to transit in this area. On the other hand, ArcGIS consistently returns lower accessibilities for this location as well as a cluster of points in the city's north-west around the Pearson International Airport. Comparisons between ArcGIS and R5R generate

the fewest number of statistically significant clusters with only smaller pockets of higher or lower accessibilities calculated beyond the more consistent trends noted above.
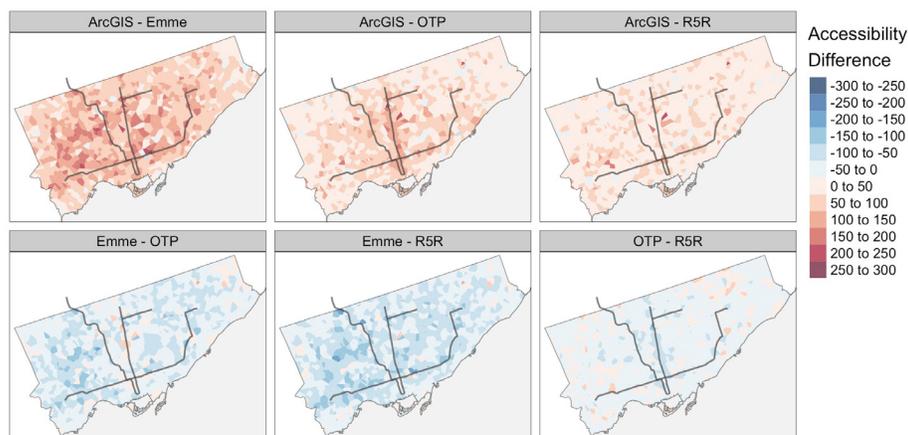


**Figure 4.** Accessibility differences
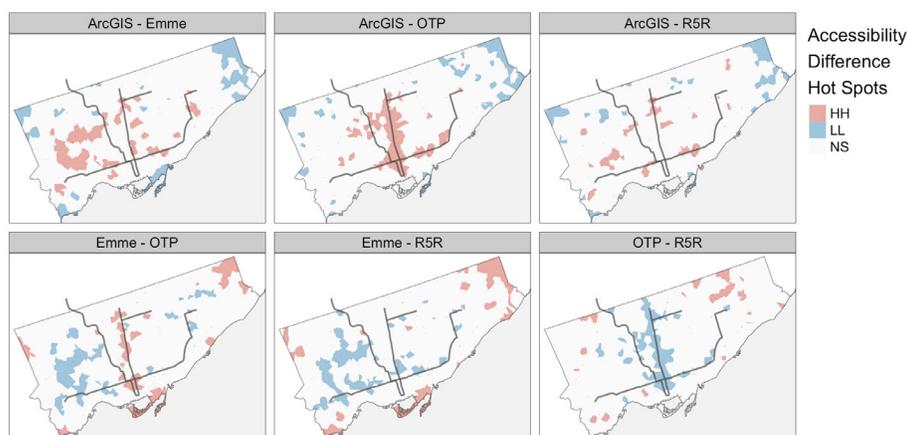


**Figure 5.** Accessibility difference hot spots

## 4     Discussion and conclusions

There is an increasingly rich selection of software tools available for applied accessibility analysis and the functions available within them continue to evolve. This research reviewed the state-of-the-art in place-based accessibility analysis, including different types of accessibility measures and the technical dimensions that inform them. We then conducted a comparative evaluation of ArcGIS Pro, Emme, R5R, and implementations of OTP in R and Python for accessibility analysis with a focus on differences in computation time and accessibility results. From this, we draw conclusions related to each tool followed by a discussion of overall implications for the state of practice.

First, we offer practitioners some conclusions about the relative strengths and weaknesses of each tool for the given place-based transit accessibility workflows. Although ArcGIS Pro's built-in Network Analyst is only single-threaded and does not offer the ability to calculate accessibility, tests revealed that custom tools utilizing Python multiprocessing can allow network analyses to scale. The tool estimates

the shortest travel times and highest calculated accessibilities within this study area. However, the software's public transit routing algorithm is closed-source and generally opaque. The solver may also return impractical results, such as a high number of transfers or long walking times.

Emme is a powerful tool for accessibility analysis whose support for congested transit travel times is unique amongst the other tools considered (although users can manually create "retrospective" GTFS files based on actual vehicle movements for use with other packages, e.g., Wessel and Farber (2019)). This feature is likely impacting the calculation of lower accessibilities along some surface transit routes in the case study area compared to the other results. However, accessibility analysis requires the integration of several custom additional tools, it relies on a sparse network and traffic analysis zone centroids that greatly simplify and abstract the walking portion of transit trips, and license restrictions in the present case limit the size of the analysis.

OTP in general allows for very flexible multi-modal routing with realistic travel behaviour, such as reluctance weights and limits on transfers and walking distances. The tool and data inputs are also free and open-source but require additional tools or scripts to utilize for accessibility analysis. Although the Python-OTP scripts performed well overall, the R-OTP tool generally performed poorly when conducting large transit accessibility analysis tasks. One potential issue for accessibility analyses based on travel time alone is that OTP's routing algorithm minimizes generalized costs in its search for optimal paths. The use of OSM for streets may also raise issues of data quality and the software lacks the network editing capabilities of ArcGIS and Emme.

Finally, the newer $R^5$ multimodal routing tool from Conveyal builds on OTP's foundations in flexible and realistic transit routing but with a focus on analytical operations that minimize travel time. When utilized through the R5R package, this tool offers dramatic time savings in the calculation of OD matrices compared to the other software packages. The R5R tool scored highest in our ranking of the technical features of each tool and offers built-in support for calculating accessibility using a variety of impedance functions. But like OTP, R5R relies on OSM for streets data.

From these results, it appears each tool offers its own strengths and weaknesses based on computational time as well as their alignment with the technical criteria we consider as the state-of-the-art in applied accessibility analysis. Although we have tried to standardize inputs, the idiosyncrasies of the individual tools in terms of their underlying data structures and analysis parameters means these tests are not perfectly comparable and analysts should consider these differences when selecting a tool. Nevertheless, ArcGIS and R5R appear to be the most similar in terms of their focus on minimizing travel time, which is traditionally the foundation of accessibility analysis. But while their results are the most consistent across the tools analyzed, R5R offers the most rapid calculation of travel times.

With regards to implications for the practice of accessibility analysis, it is now easier than ever to move beyond cumulative opportunity measures and employ alternative impedance functions as well as consider multimodal travel, generalized costs, temporal variation, and traveler perceptions. Fast calculation times also enable accessibility analyses at finer spatial scales. The next generation of tools can also enhance practice by making it easier for practitioners to account for competition.

However, our comparative analysis of the different tools has revealed that each produces different travel time results for the same origin-destination pair, causing accessibility results to differ according to the tool used. Furthermore, while travel times and accessibilities are generally highly correlated across tools, there is evidence of some statistical tendencies and spatial clustering in the travel time differences that can have a significant effect on calculated accessibility scores. This may not be surprising given the different data structures and routing algorithms utilized by each software package. But beyond the effects of spatial dependence from modifiable areal units and frame dependence from modifiable temporal units in accessibility calculation (Pereira, 2019), these results highlight the potential for algorithmic dependence in calculated accessibility results.

Consequently, software developers must be transparent and rigorous in communicating the methods and assumptions employed in their tools. Practitioners and researchers engaged in accessibility analysis should interrogate and document the tools used as part of standard reporting procedures. With the increasing emphasis on accessibility-based planning and the associated implications of accessibility analyses for informing policy, such results warrant further investigation in the academic community and suggest that analysts and planners should approach accessibility calculation and the interpretation of results with care.

## Acknowledgements

## References

Allen, J. (2020). *OpenTripPlanner analysis*. Retrieved from https://github.com/SAUSy-Lab/OpenTrip-Planner_analysis

Anselin, L. (1995). Local indicators of spatial association—LISA. *Geographical Analysis, 27*(2), 93–115.

ARC. (2019). *Performance, the Atlanta region's plan*. Atlanta, GA: Atlanta Regional Commission.

Bast, H., Delling, D., Goldberg, A., Müller-Hannemann, M., Pajor, T., Sanders, P., . . . Werneck, R. F. (2016). Route planning in transportation networks. In L. Kliemann, & P. Sanders (Eds.), *Algorithm engineering* (pp. 19-80). Berlin: Springer.

Blanchard, S. D., & Waddell, P. (2017). Urban access: Generalized methodology for measuring regional accessibility with an integrated pedestrian and transit network. *Transportation Research Record, 2653*(1), 35–44.

Blumenberg, E., & Pierce, G. (2014). A driving factor in mobility? Transportation's role in connecting subsidized housing and employment outcomes in the Moving to Opportunity (MTO) program. *Journal of the American Planning Association, 80*, 52¬–66.

Boeing, G. (2017). OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems, 65*, 126–139.

Boisjoly, G., & El-Geneidy, A. M. (2017). How to get there? A critical assessment of accessibility objectives and indicators in metropolitan transportation plans. *Transport Policy, 55*, 38–50.

Conway, M. W., & Stewart, A. F. (2019). Getting Charlie off the MTA: A multiobjective optimization method to account for cost constraints in public transit accessibility metrics. *International Journal of Geographical Information Science, 33*(9), 1759–1787.

Conway, M. W., Burd, A., & van der Linden, M. (2017). Evidence-based transit and land use sketch planning using interactive accessibility methods on combined schedule and headway-based networks. *Transportation Research Record, 2653*(1), 45–53.

Conway, M. W., Byrd, A., & Van Eggermond, M. (2018). Accounting for uncertainty and variation in accessibility metrics for public transport sketch planning. *Journal of Transport and Land Use, 11*(1), 541–558.

Cooley, D., & Barcelos, P. (2020, November). *Package 'googleway.'* Retrieved from https://cran.r-project.org/web/packages/googleway/googleway.pdf

Crowley, D. F., Shalaby, A. S., & Zarei, H. (2009). Access walking distance, transit use, and transit-oriented development in North York City Center, Toronto, Canada. *Transportation Research Record, 2110*, 96–105.

Cui, M., & Levinson, D. (2020). Primal and dual access. *Geographical Analysis, 52*(3), 452–474.

El-Geneidy, A., Levinson, D., Diab, E., Boisjoly, G., Verbich, D., & Loong, C. (2016). The cost of equity: Assessing transit accessibility and social disparity using total travel cost. *Transportation Research Part A: Policy and Practice, 91*(302–316).

Farber, S., & Allen, J. (2019). *The Ontario Line: Socioeconomic distribution of travel time and accessibility benefits*. Toronto: Metrolinx.

Farber, S., Ritter, B., & Fu, L. (2016). Space–time mismatch between transit service and observed travel patterns in the Wasatch Front, Utah: A social equity perspective. *Travel Behavior and Society, 4*, 40–48.

Ford, A. C., Barr, S. L., Dawson, R. J., & James, P. (2015). Transport accessibility analysis using GIS: Assessing sustainable transport in London. *ISPRS International Journal of Geo-Information, 4*(1), 124–149.

Fransen, K., Farber, S., Deruyter, G., & De Maeyer, P. (2018). A spatio-temporal accessibility measure

for modelling activity participation in discretionary activities. *Travel Behavior and Society, 10*, 10–20.

Geurs, K. T., & van Wee, B. (2004). Accessibility evaluation of land-use and transport strategies: Review and research directions. *Journal of Transport Geography, 12*(2), 127–140.

Giles-Corti, B., Broomhall, M. H., Knuiman, M., Collins, C., Douglas, K., Ng, K., . . . Donovan, R. J. (2005). Increasing walking: How important is distance to, attractiveness, and size of public open space? *American Journal of Preventative Medicine, 28,* 169–176.

Handy, S. (2008). Regional transportation planning in the US: An examination of changes in technical aspects of the planning process in response to changing goals. Transport Policy, 15, 113–116.

Higgins, C. D. (2019). Accessibility toolbox for R and ArcGIS. *Transport Findings.* https://doi.org/10.32866/8416

Higgins, C. D., DeJohn, A., Farber, S., Palm, M., Vaughan, J., Widener, M., . . . Miller, E. J. (2020). *Transportation accessibility advice* (Final report). Toronto: University of Toronto Transportation Research Institute.

Hu, Y., & Downs, J. (2019). Measuring and visualizing place-based space-time job accessibility. *Journal of Transport Geography, 74*, 278–288.

Huber, S., & Rust, C. (2016). Calculate travel time and distance with OpenStreetMap data using the Open Source Routing Machine (OSRM). *The Stata Journal, 16*(2), 416–423.

INRO. (2019). Introducing EMME scenes. Retrieved from https://info.inrosoftware.com/blog/introducing-emme-scenes

Iseki, H., & Taylor, B. D. (2009). Not all transfers are created equal: Toward a framework relating transfer connectivity to travel Behavior. *Transport Reviews, 29*, 777–800.

Kelly, C., Hulme, C., Farragher, T., & Clarke, G. (2016). Are differences in travel time or distance to healthcare for adults in global north countries associated with an impact on health outcomes? A systematic review. *BMJ Open, 6*(11), e013059.

Kwan, M. P. (1998). Space-time and integral measures of individual accessibility: A comparative analysis using a point-based framework. *Geographical Analysis, 30*(3), 191–216.

Lovelace, R. (2021). Open source tools for geographic analysis in transport planning. *Journal of Geographical Systems, 23*, 547–578.

Martens, K., & Golub, A. (2018). A fair distribution of accessibility: Interpreting civil rights regulations for regional transportation plans. *Journal of Planning Education and Research, 41*(4), 425–444.

Mayaud, J. R., Tran, M., Pereira, R. H., & Nuttall, R. (2019). Future access to essential services in a growing smart city: The case of Surrey, British Columbia. *Computers, Environment and Urban Systems, 73*, 1–15.

Merlin, L. A., & Hu, L. (2017). Does competition matter in measures of job accessibility? Explaining employment in Los Angeles. *Journal of Transport Geography, 64,* 77–88.

Metro. (2018). 2018 regional transportation plan. Portland, OR: Oregon Metro.

Miller, E. J. (2018). Accessibility: Measurement and application in transportation planning. *Transport Reviews, 38,* 551–555.

Miller, E. J., Vaughan, J., King, D., & Austin, M. (2015). *Implementation of a "next generation" activity-based travel demand model: The Toronto case.* Presentation at the Travel Demand Modelling and Traffic Simulation Session of the 2015 Conference of the Transportation Association of Canada, Charlottetown, PEI.

Morgan, M., Young, M., Lovelace, R., & Hama, L. (2019). OpenTripPlanner for R. *Journal of Open Source Software, 4*(44), 1–2.

North Central Texas Council of Governments. (2018). *Mobility 2045: The Metropolitan Transportation Plan for North Central Texas.* Arlington, TX.

Openshaw, S. (1984). *The modifiable areal unit problem.* Norwich: Geo Books.

OTP. (2020). Comparing OTP2 and OTP1. *OpenTripPlanner.* Retrieved from http://docs.opentrip-planner.org/en/latest/Version-Comparison/#comparing-otp2-and-otp1

OTP. (2020). *Routing bibliography.* Retrieved from https://docs.opentripplanner.org/en/latest/Bibliography/

Owen, A., & Levinson, D. (2014). *Access across America: Transit 2014 methodology* (Final report No. 14-12). Minneapolis: University of Minnesota, Accessibility Observatory.

Páez, A., Higgins, C. D., & Vivona, S. F. (2019). Demand and level of service inflation in floating catchment area (FCA) methods. *Plos one, 14*(6), e0218773.

Páez, A., Scott, D. M., & Morency, C. (2012). Measuring accessibility: Positive and normative implementations of various accessibility indicators. *Journal of Transport Geography, 25*(C), 141–153.

Padgham, M., Stepniak, M., & Kapp, A. (2021, June). *gtfsrouter.* Retrieved from https://cran.r-project.org/web/packages/gtfsrouter/gtfsrouter.pdf

Pereira, R. H. (2019). Future accessibility impacts of transport policy scenarios: Equity and sensitivity to travel time thresholds for bus rapid transit expansion in Rio de Janeiro. *Journal of Transport Geography, 74*, 321–332.

Pereira, R. H., Grégoire, L., Wessel, N., & Martins, J. (2019). *Tutorial with reproducible example to estimate a travel time matrix using OpenTripPlanner and Python.* Retrieved from https://github.com/rafapereirabr/otp-travel-time-matrix. doi:10.5281/zenodo.3242134

Pereira, R. H., Saraiva, M., Herszenhut, D., Braga, C. K., & Conway, M. W. (2021). r5r: Rapid realistic routing on multimodal transport networks with R5 in R. *Findings.* https://doi.org/10.32866/001c.21262.

Proffitt, D. G., Bartholomew, K., Ewing, R., & Miller, H. J. (2019). Accessibility planning in American metropolitan areas: Are we there yet? *Urban Studies, 56,* 167–192.

Saxon, J., Koschinsky, J., Acosta, K., Anguiano, V., Anselin, L., & Rey, S. (2021). An open software environment to make spatial access metrics more accessible. *Journal of Computational Social Science,* 1–20. https://doi.org/10.1007/s42001-021-00126-8

SCAG. (2016). *Environmental justice appendix, 2016-2040 regional transportation plan.* Los Angeles, CA: Southern California Association of Governments.

Shen, Q. (1998). Location characteristics of inner-city neighborhoods and employment accessibility of low-wage worker. *Environment and Planning B: Planning and Design, 25,* 345–365.

Siddiq, F., & Taylor, B. D. (2021). Tools of the trade? Assessing the progress of accessibility measures for planning practice. *Journal of the American Planning Association.* https://doi.org/10.1080/01944363.2021.1899036

Stewart, A. (2020, July). *Comparison of R5 and OTP. Github: Conveyal R5.* Retrieved from https://github.com/conveyal/r5/issues/575#issuecomment-655237929

TfL. (2006). *Transport 2025.* London, UK: Transport for London.

Tilhaun, N., & Li, M. (2015). Walking access to transit stations: Evaluating barriers with stated preference. Transportation Research Record, 1534, 16–23.

TTC. (2017). *Service standards and decision rules for planning transit service.* Toronto: Toronto Transit Commission.

Unterfinger, M., & Possenriede, D. (2021, April). *Package 'hereR'.* Retrieved from https://cran.r-project.org/web/packages/hereR/hereR.pdf

Wardman, M. (2004). Public transport values of time. *Transport Policy, 11*, 363–377.

Wessel, N., & Farber, S. (2019). On the accuracy of schedule-based GTFS for measuring accessibility. *Journal of Transport and Land Use, 12*(1), 475–500.

Xi, Y. L., Miller, E. J., & Saxe, S. (2018). Exploring the impact of different cut-off times on isochrone measurements of accessibility. *Transportation Research Record, 2672,* 113–124.