
Conveyal Analysis Documentation

Release v4.13.0

Conveyal

Mar 26, 2020

1	Contents	3
1.1	Preparing a baseline transport network	3
1.1.1	Setting up a new region	3
1.1.2	Uploading a GTFS bundle	5
1.1.3	Creating a Project	6
1.2	Preparing opportunity datasets	6
1.2.1	Uploading an opportunity dataset	6
1.2.2	Managing opportunity datasets	7
1.2.3	LODES dataset import	7
1.3	Overview of network editing	7
1.3.1	Create a new modification	8
1.3.2	Create a new scenario	8
1.3.3	Activate a modification in a scenario	8
1.4	Modification types	8
1.4.1	Add trip pattern	9
1.4.2	Adjust dwell time	10
1.4.3	Adjust speed	10
1.4.4	Convert to frequency	10
1.4.5	Remove stops	11
1.4.6	Remove trips	11
1.4.7	Reroute	12
1.4.8	Custom modifications	12
1.5	Timetables	13
1.5.1	Exact Times	13
1.5.2	Copying Timetables	13
1.6	Phasing	14
1.7	Managing modifications	15
1.7.1	Toggling display of modifications	15
1.7.2	Exporting modifications	15
1.7.3	Importing modifications	15
1.7.4	Manually edit raw modifications (beta)	16
1.8	A simple walkthrough	16
1.9	Single point analysis	17
1.9.1	Isochrone map	17
1.9.2	Analysis panel	18
1.10	Regional analysis	19

1.10.1	Starting a regional analysis	19
1.10.2	Inspecting results	19
1.10.3	Downloading regional results	19
1.10.4	Comparing regional analyses	19
1.10.5	Measuring aggregate accessibility	20
1.11	Options and configuration	21
1.11.1	Advanced settings	22
1.11.2	Errors and warnings	22
1.12	Developers Guide	22
1.12.1	Installing your own copy	22
1.12.2	Updating this manual	23
1.13	Frequently Asked Questions	23
1.13.1	How does Conveyal Analysis calculate travel times?	23
1.13.2	When starting an analysis, why does the “initializing cluster” message persist for so long?	23
1.13.3	Why are travel times longer than expected?	24
1.13.4	Who can see and edit my projects?	24
1.13.5	How long should it take to start a single-origin analysis?	24
1.13.6	How long should each regional analysis take?	25
1.13.7	How do I model corridor upgrades where a new trunk is added and existing services are converted to feeder routes?	25
1.13.8	What new features is Conveyal working on?	25
1.14	Methodology	26
1.14.1	Spatial resolution	26
1.14.2	Accounting for variability	26
1.14.3	Time percentile	27
1.15	Glossary	28
	Index	29

This manual will help you use [Conveyal Analysis](#) to edit land-use and transportation scenarios, and evaluate them in terms of accessibility.

If you are using the version hosted by Conveyal at analysis.conveyal.com, log in with the Chrome browser using the instructions sent via e-mail, then select a region and project to use. Once you have selected a project, you can [edit](#) and [analyze](#) scenarios.

Most users of analysis.conveyal.com will already have a project prepared when they first log in. If you don't see a region or a project, you'll first need to [prepare a baseline network](#) or get in touch with your support team.

1.1 Preparing a baseline transport network

Accessibility calculations will be based on the *transport network* that is set up for your *region*. This includes scheduled transit services as well as the region's network of streets, sidewalks, bikeways, etc. Initially you will set up a *baseline network* representing the transit services described in the *GTFS bundle* for your region. Later on you will likely want to compare alternative *scenarios* by creating *modifications* to the baseline network.

1.1.1 Setting up a new region

The *region* is a *bounding box* defining the area to be used for accessibility calculations. It should generally cover the entire service area of the agency or agencies you will be working with. Opportunities and network components outside this area will generally not be considered.

The regions page, shown after you log in, provides a list of existing regions if any, and the option to create a new region. It is also accessible at any later point by clicking the globe icon (). From the regions page, set up a new region by clicking:

Start by using the search bar in the map to automatically locate your city or country by name. You can also move the bounding box by dragging its corners on the map or by manually entering coordinates in the sidebar. Enter a name for the region and an optional description.

You must also upload an appropriate [OpenStreetMap \(OSM\)](#) extract at this point, which will serve as the street layer of the transport network. This will be used for any walking or biking segments of a trip as well as for some transfers between stops and stations. Note that while several formats exist for OSM data, we require the [PBF format](#) because it is more compact and faster to process. Your extract should cover your entire service area or region, but not extend unnecessarily far beyond it as that may impact processing time.

The process of downloading and processing OSM data into an appropriate format is covered in the next section. For now, if all looks good with your region, you should be able to click the *Set up a new region* button. After a bit of time for uploading and processing the OSM data, you should be prompted to upload a GTFS bundle.

Preparing the OSM data

Downloading

Extracts from the global OSM database can be downloaded in [many different ways](#). Some popular services like those provided by [Geofabrik](#) or [Nextzen](#) provide easy downloads for selected cities and regions. Be sure to download data covering your entire region - the predefined areas used by these sites may or may not align well with your region. You'll often need to download an extract for a country or region larger than your true analysis area, then crop it to the size you need.

Cropping

Performing accessibility analysis with excessively large OSM data can lead to significant increases in computation time and complexity. We strongly recommend cropping the data if they cover an area significantly larger than your transportation network or opportunity data. Several command line tools are able to perform these cropping operations:

- [Osmosis](#) is a multi-platform Java tool that works on Windows, Linux, and MacOS.
- [OSMConvert](#) is a fast tool pre-built for Windows and Linux and available on MacOS and Linux as part of the `osmctools` package.
- [Osmium-Tool](#) is a personal favorite that is extremely fast but only straightforward to install on Linux and MacOS platforms.

Below are some example crop commands for these different tools. You'll need to replace `input.osm.pbf` with the name of your downloaded PBF file and change the coordinates of the area to clip to.

When creating a region, the panel will show an `osmconvert` command pre-filled with the current regional bounding box. If you have `osmconvert` installed locally, you can paste this command into to your local command line and modify the filenames to crop your OSM data to regional boundaries before upload.

Osmosis:

```
osmosis --read-pbf input.osm.pbf \  
  --bounding-box left=-79.63 bottom=43.61 right=-79.12 top=43.83 \  
  --write-pbf cropped.osm.pbf
```

OsmConvert:

```
osmconvert input.osm.pbf \  
  -b=-79.63,43.61,-79.12,43.83 --complete-ways -o=cropped.osm.pbf
```

Osmium-Tool:

```
osmium extract \  
  --strategy complete_ways --bbox -79.63,43.61,-79.12,43.83 \  
  input.osm.pbf -o cropped.osm.pbf
```

The latter two commands expect bounding boxes to be specified in the format `min_lon,min_lat,max_lon,max_lat`. We frequently find bounding boxes using the convenient [Klokantech bounding box tool](#). Selecting the "CSV" format in the lower left will give exactly the format expected by these tools. You can also adapt the bounding box values shown in the region setup panel of Analysis.

Note that files larger than 500MB may be rejected on upload. Please contact your support team if you genuinely need to upload a file of this size, or need assistance in cropping and filtering OSM data.

Filtering

OpenStreetMap contains a lot of data besides the streets, paths, and platforms we need for accessibility analysis. As of this writing more than half of the ways in OSM are buildings, and slightly less than a quarter are roads or paths. Filtering out unneeded data will reduce your file size and speed the upload and processing by Analysis. As in the previous section, sample commands are provided below that will remove any unnecessary tags and should dramatically reduce the output file size.

Osmosis:

```
osmosis \
  --read-pbf input.osm.pbf \
  --tf accept-ways highway=* public_transport=platform railway=platform park_ride=* \
  --tf accept-relations type=restriction \
  --used-node \
  --write-pbf filtered.osm.pbf
```

Osmium-Tool:

```
osmium tags-filter input.osm.pbf \
  w/highway w/public_transport=platform w/railway=platform w/park_ride r/
↪type=restriction \
  -o filtered.osm.pbf
```

1.1.2 Uploading a GTFS bundle

Start by gathering *GTFS feeds* for the transit agencies whose service will be included in your region. A GTFS feed is a set of CSV files inside a .zip archive, and a *GTFS bundle* is a set of one or more GTFS feeds.

If you just created a new region, you will be prompted to upload a GTFS bundle:

Otherwise, you can click the database icon () on the sidebar to access your GTFS Bundles at any time.

First you need to assign your new bundle a name. We recommend staying organized by basing the name on the time period and/or services included such as “TTC&GO - June 2019”. Next choose one or more .zip files to upload. You can select multiple GTFS feeds in the file dialogue by shift-clicking, control-clicking or command-clicking (depending on your browser/operating system). Finally, click the create button to confirm.

Again, note that files larger than 500MB may be rejected on upload. The largest GTFS feeds in regular use are below 400MB and most are much smaller than this. A larger file may indicate a problem. Please contact your support team if you genuinely need to upload a larger file.

If there is no GTFS available for your region you can, as a workaround, create a blank slate by using a valid GTFS feed from somewhere else in the world. After creating a project as described below, you can then attempt to [import modifications](#) from a shapefile representing service in your region.

Ensure any GTFS you upload follows requirements of the specification. Various [validation tools](#) are available. Common issues include:

- Missing required files
- Calendar dates that do not cover an intended date of analysis.
- Using values other than 0-7 in the route_type column of routes.txt

1.1.3 Creating a Project

Uploading and processing a GTFS bundle may take several minutes. Once processing is complete, you should be able to create a new *project* based on the bundle you uploaded. If you aren't on the projects page already, click the project icon () and then,

A project is essentially a wrapper around a bundle which associates it with any scenarios and modifications you may create later on. The purpose of projects is to allow multiple users to work simultaneously on the same baseline network without stepping on each other's toes. The section on *Managing modifications* describes how modifications can be shared between projects that are based on the same bundle. Once created, the bundle associated with a project cannot be changed.

Give the project a descriptive name, select a bundle to which the project will be associated, and click the create button to confirm. You are now ready to move on to [editing scenarios](#).

1.2 Preparing opportunity datasets

To measure access to spatially distributed opportunities (e.g. jobs, people, schools), you will need to ensure at least one **opportunity dataset** has been added to your project. To add or manage opportunity datasets, click the grid icon on the sidebar:

Opportunity datasets can be added either by uploading a CSV or shapefile representing the spatial distribution of opportunities, or by using the LODES data import tool (US only).

1.2.1 Uploading an opportunity dataset

A spatial dataset representing any opportunities of interest may be uploaded as either a CSV or shapefile, formatted to the specifications described below. To start an upload, click:

Enter a name for the opportunity dataset source, then select the appropriate file(s). Once this is done, click the upload button to start the upload. Opportunity datasets will be saved with the source name you entered and the attribute/column names of the uploaded files.

After processing is complete, you can refresh the page and see dot-density maps of your datasets, converted to the [analysis grid](#) used in Conveyal Analysis.

CSV

A [Comma Separated Value](#) (CSV) file can be used to represent opportunities as point features. The CSV should have columns for latitude, longitude, and one or more associated opportunity counts.

In the following example, a CSV file represents the location of restaurants along with the number of restaurants (1 per point) and the estimated number of employees at each. Such a dataset could be used to represent either opportunities for entertainment or for work.

```
lat,lon,restaurant_count,est_employees,name
39.111475,-84.633360,1,5,"La Rosa's Pizza"
39.105990,-84.556554,1,12,"Primavista"
39.103175,-84.509386,1,8,"Shanghai Mama's"
39.160658,-84.543407,1,11,"Boswell Alley"
39.100883,-84.512742,1,25,"McCormick & Schmick's"
...
```

Note that only *numeric* fields will be accepted. The `name` field in this example would be dropped. Be careful to give each column a meaningful name in the header row and remove any numeric columns which do not represent opportunity counts such as ID fields. Be sure to name the file with a `.csv` extension and use only commas as separators, i.e. no tab-separated or fixed-width text files.

You will be prompted for the names of the fields containing latitude and longitude values. In the example above you would enter `lat` and `lon` respectively. We currently only support CSV files which specify points in WGS 84 latitude/longitude coordinates.

Shapefiles

A shapefile may represent opportunities as either points or polygons. Opportunity counts associated with polygons will be treated as though they are uniformly distributed within the given area. For the easiest experience, any numeric fields not representing opportunities (ID fields, etc) should be removed before uploading. As with CSV, text fields will be dropped - this includes fields which may look like numbers (e.g. "1", "NA") but are actually stored as text. If a field is not showing up after upload, ensure that it was actually stored as a numeric value rather than text.

Shapefiles should not be zipped; select all of the files in the Shapefile when uploading (at the very least, `.shp`, `.shx`, `.dbf` and `.prj`). How you select multiple files depends on your browser and operating system, but generally will involve shift-clicking, control-clicking or command-clicking.

1.2.2 Managing opportunity datasets

Selecting an existing opportunity dataset from the dropdown menu will give you options to:

(e.g. all attributes from a shapefile)

1.2.3 LODES dataset import

For projects in the United States, employment data from from the *Longitudinal Employer-Household Dynamics (LEHD)* Program can be imported with a one-click import function.

The **LODES** (*LEHD Origin-Destination Employment Statistics*) dataset provides block-level data on the home and work locations of employed people living in the US. These job and worker counts are made available both as totals and as counts disaggregated by industry, earnings, education level, etc.

Conveyal will update the LODES one-click import function after the Census Bureau addresses certain data limitations for more recent years. The current releases for 2016 and 2017 exclude federal employees, and the 2017 release is missing data for Alaska and South Dakota. For more information, see the [LODES technical documentation](#).

1.3 Overview of network editing

After selecting a *project*, you will arrive at the modifications page, shown below. Each project has an expandable list of numbered *scenarios* followed by a list of *modifications* grouped by type. Each modification represents a single operation on the *baseline network* (for example adding a line, or adjusting the speed of an existing line) and can be activated in multiple scenarios. Scenarios are a way of packaging isolated modifications into meaningful groups representing changes to the broader network. For example, you could create a scenario representing each of several funding scenarios with one for service cuts, one for extended service and so on.

To prevent data conflicts, the system does not support multiple users or browser tabs editing the same project simultaneously. If you expect multiple people in your organization to be using Conveyal Analysis at the same time, you may want to create projects for each of them. Modifications can later be imported into a project from other projects based on the same *GTFS bundle* (See [Managing modifications](#)).

1.3.1 Create a new modification

To add a modification, first navigate to the modification page (the icon) where you will need to select a project if you haven't already. If your project is already selected you should see the following button below the project name.

This will open a window prompting you to select the **modification type** and enter a name for the modification. Choose a descriptive name - you will generally want to include at least the name of the route created or affected by the modification. After confirming these details, you will be taken to a page displaying options that vary by modification type.

1.3.2 Create a new scenario

From the modifications page, you can also create and edit *scenarios*. A scenario is used for packaging modifications into meaningful groups representing broader changes to the network. You may need to click the scenarios bar to expand it. Once expanded, you should see the baseline scenario and any others that you have created. Click the button to create a new scenario and give it a name e.g. “*Expanded BRT*” or “*Downtown Stop Consolidation*”. When editing modifications, you may opt to add them to one or more scenarios. The *baseline scenario* represents the *baseline network* and cannot be edited.

1.3.3 Activate a modification in a scenario

By default, each modification is active in all scenarios that exist when the modification is created. You can change which scenarios a modification is active in by using the checkboxes corresponding to scenarios at the bottom of the modification detail panel.

1.4 Modification types

When creating a new modification, you have several different options, each explained in detail below.

- *Add trip pattern*
- *Adjust Dwell Time*
- *Adjust Speed*
- *Convert to Frequency*
- *Remove Stops*
- *Remove Trips*
- *Reroute*
- *Custom* (not for general use)

Across all modification types the basic actions are consistent.

The *copy* command automatically creates a copy of the current modification with `(copy)` appended to the name. This copy can be found in the modification menu. The copy feature is useful for creating modifications that share common elements like adding a route with branches off a main trunk.

Clicking the small blue back arrow saves your changes and takes you back to the list of all modifications:

Changes are also automatically saved every 10 seconds.

1.4.1 Add trip pattern

The **add trip pattern** modification allows you to add new *trip patterns* to your transport scenario. A trip pattern is a set of stops visited in order by a transit vehicle. Often a route will consist of multiple trip patterns, e.g. one for each direction of travel. This modification also offers a bidirectional option to allow a single trip pattern to represent travel in both directions. This may be easier for modes with generally bidirectional stations like subways, ferries, or cable cars.

After creating a new modification you will see something like the view below.

You can set the mode (e.g. bus) and add a description for the modification at the top of the panel. To create an alignment, or to edit the alignment you've previously created, click

This will activate the map-based route editing mode. You can stop and save your work at any time with:

New route alignments are defined by an ordered set of stops and control points. The actual route taken between these can either follow existing streets or it can take a straight line between points. It is also possible to combine the two options in one route as when a bus runs on the street but then diverts into a planned busway which is not yet part of the street network. The distance between stops is used to estimate segment travel time on the new alignment so it is important to be as accurate with the alignment as possible.

By default when you begin editing a new route, you can click once on the map to place the first stop, then again to place the second stop, and so on. If you click on an existing stop (indicated by a small gray circle), the icon for the new stop will be black and the new transit service will stop at that existing stop. If you click in a place where there is not an existing stop, a new stop (in blue) will be created.

Once created, any stop or control point can be moved by dragging it to a new location. Clicking on a stop gives you the option to delete it, or convert it to a *control point* through which the route will pass without stopping. Similarly, control points can be converted back to stops or deleted by clicking on them.

You can also insert a stop into the middle of an existing alignment. Simply hover over the any part of the line and you should see a new stop appear below your cursor. Click to add it at that location.

On the panel to the left of the map there are a few options available while editing an alignment. These are described below.

- **Auto-create stops at set spacing:** Specifies whether stops should be created automatically at a specified interval along the route. When using this option, you will likely want to define your route alignment with control points rather than stops, however if stops are used new stops will not be added unless the gap between existing stops is sufficiently large. The stop spacing can easily be changed later.
- **Bidirectional:** If this option is checked (the default value), vehicles will travel in both directions along the described geometry. If it is toggled off, vehicles will only travel in the direction the line is drawn. Directed alignments can be useful when there are *couplets* or other aspects of the route that don't follow the same alignment in both directions. Directed alignments are also necessary for *phasing*.
- **Follow streets:** If this option is unchecked (the default value), stops and control points will be connected by straight lines. If it is checked, the route will follow streets. This setting only applies to segments that are actively being edited and will not cause already drawn segments to follow the streets, allowing you to draw part of a route on street and part off-street. The length of each segment is used to estimate travel times between stops, so this feature is very useful for approximating the length of bus/streetcar routes.
- **Extend:** If this option is checked (the default), clicking on the map will extend the route forward from the end by adding a new stop in the direction the line was drawn. If you wish to edit segments you already created, by adding stops or control points, it is often convenient to uncheck this option so that stray clicks do not extend the route from one of its termini.
- **Extend from end:** When *Extend* is enabled, this option allows you to specify whether new stops will be added moving forward from the last stop of the route (the default) or backward from the first stop of the route. When this option is *unticked*, new stops will be added as extensions backwards from the beginning.

- **Snap to existing stops (beta):** Experimental feature that adds stops from the baseline GTFS bundle to the new route. This feature adds stops within 20 meters of the alignment, but not more than one stop approximately every 80 meters.

Once you have created an alignment, you'll need to specify when the route runs using a [simplified timetable](#). You can do this by [copying a previously created timetable](#), or by clicking:

1.4.2 Adjust dwell time

You may also want to adjust the dwell time along a route or at a particular stop, for example to model the effects of off-board fare collection, or the effects of increasing ridership at a particular stop. As with the remove-stops modification, you can select a feed, route and optionally patterns. You can then use the map to select the affected stops (if you skip this step, all stops will have their dwell times adjusted). You can then choose to either enter a new dwell time (in seconds), or scale the existing dwell times (for instance, entering 2 would double existing dwell times).

Unfortunately, the `stop_times.txt` files of many GTFS feeds use equal `arrival_time` and `departure_time` values. For such feeds that do not represent dwell time explicitly, this modification type may not be immediately applicable.

1.4.3 Adjust speed

This modification can be applied to multiple routes, but only one route will be shown on the map.

This modification does not automatically increase frequency in response to the more efficient routes. However, it can be paired with an adjust frequency modification to model such a response.

Sometimes you will want to adjust the speed on an entire route or just a *segment* of one. For instance you might want to model the addition of bus lanes or an application of transit signal priority. This modification allows you to scale the speed of existing scheduled *trips*, or segments of them, by a constant factor. It can be applied to whole routes (in which case more than one route may be selected) or to selected segments of a particular route or to just a subset of *trip patterns*.

You will need to select a GTFS feed, routes and optionally *trip patterns* if only one route is selected. All trip patterns will be selected by default. Segments can be selected by clicking *Select* from the toolbox shown in the figure below. This will allow you to begin drawing a polygon selection area on the map. Any segments within this area will be selected when the polygon is closed and selected segments will then render on the map in purple.

The *Select* option will begin a new selection and the *Add to* option will add to the current selection if any. *Remove from* allows you to select segments to remove from the current selection and *Clear* un-selects all segments.

Finally, enter a numeric value in the *Scale speed by* field — this is the factor to multiply the speed by. For instance, if you enter 1.3, the speed of vehicles will increase by 30% when traveling between stops. Note however that this modification does not affect dwell times; to model changes in dwell time, see the *Adjust dwell time* modification. It also does not take into account the possibility of increased frequency due to faster, more efficient routes. However, it can be paired with a *Convert to frequency* modification to model that scenario.

1.4.4 Convert to frequency

Often a scenario will include changes to the number of trips per hour on an existing route. We support this using the *convert to frequency* modification. It works by replacing the scheduled trips for one or more existing *trip patterns* with frequency based *Timetables*. You can opt either to

- delete all existing trips and replace them with your own timetables or
- replace trips only when the new timetables would be in effect.

The travel and dwell times on the replacement trips are based on those of an existing trip which you can select from the affected trip pattern. The new timetables can make use of *Phasing* to reproduce timed services on branching lines, etc.

First, create a new modification and select *Convert to frequency*. Give the modification a name; you will likely want to name it after the route you plan to modify.

Start by selecting a *GTFS feed* and then the route from that feed that you want to adjust. You should see all trip patterns for the selected route displayed on the map.

You can choose to remove all existing trips on the route (the default) and start from scratch with new timetables. Or you may choose to retain trips outside the time windows in which you specify frequencies which is useful when you are changing the frequency for only part of the day (e.g. increased weekend frequency) and want to retain the existing scheduled service at other times. This is controlled using the checkbox labeled “*Retain existing scheduled trips at times without new frequencies specified*”.

You then create any number of frequency entries using simple *timetables*. Within each timetable, you will need to select a trip pattern from your route, then a particular trip from that trip pattern to be used as a template. Travel times for the new timetable will be based on this template trip so it is important to choose one that is representative. For example, you may want to select one of the slower trips when specifying frequencies during congested peak-hour service.

Typically, you will need to create *at least* two new timetables, one for each direction of travel.

Note: Once converted to a frequency-based route with this modification, any of a route’s patterns not represented by a timetable are effectively removed. With “retain trips” set to the default value of false (unchecked), these patterns will be removed at all times of day. With “retain trips” set to true (checked), they will be removed when any frequency entry is active.

1.4.5 Remove stops

It is also possible to remove some of the stops from a route, while leaving the rest of the route untouched. To do this, create a new *remove stops* modification, and select a feed, route, and patterns as you did when removing trips. You can then use the map to select which stops to remove. Under “Selection,” click “new,” “add to,” or “remove from” to select new stops to remove, add to the existing selection, or remove from the existing selection. Stops that are selected for removal are listed at the bottom of the modification, as well as being highlighted in red on the map.

You can also specify that you wish to remove a certain amount of time at each removed stop (in addition to the dwell time specified in the original feed, which is automatically removed). This could be used to account for acceleration and deceleration, and can also be used when the original feed does not specify dwell time. This feature is useful when testing the effects of stop consolidation.

When removing the beginning of a route, the dwell times at each stop are removed, as is any time specified to be removed at each removed stop. Any remaining travel time is preserved as an offset from the start of the trip in the original GTFS to the start of trips at the new first stop. This is effectively as if the vehicles were leaving the original terminal at the same time but deadheading past all of the removed stops.

This modification does not take into account the possibility of increased frequency due to more efficient routes. However, it can be paired with a *Convert to frequency* modification to model that scenario.

1.4.6 Remove trips

Another common modification is to remove trips. The most common use is to remove entire routes, but it is also possible to remove only specific *trip patterns* on a particular route. In order to remove trips, create a new *remove trips* modification, and select a GTFS feed, route, and optionally the trip pattern of the trips to be removed. All trips on

this route/pattern combination will be removed and the route will be shown in red on the map. Note that the “active in variants” selector always specifies whether the modification is active. In this case it implies that the trips will be removed from the selected variants.

1.4.7 Reroute

This modification type can be used to represent detours, extensions, and curtailments. When creating a *reroute* modification, you first select a *GTFS feed*, route, and *trip patterns*. Once trip patterns are selected, you then select a stop at which the reroute alignment will start, or a stop at which the reroute alignment will end, or both, by clicking

then clicking an existing stop on the baseline pattern.

Once a start or end stop is specified, you can add and modify segments by clicking the *Edit route geometry* button, then clicking on the map. Editing behavior is similar to editing mode for *adding trip patterns* though with some options fixed depending on whether the route is being extended in either direction, or if the reroute is happening in the middle of the alignment.

- If only the “start of reroute/extension” is set, new stops and segments will be added extending *forward* from the selected stop on the baseline pattern (“Extend from end” turned on).
- If only the “end of reroute/extension” is set, new stops and segments will be added extending *backward* from the selected stop on the baseline pattern (“Extend from end” turned off).
- If both are set, the new segment between the stops cannot be extended, but it can be modified by clicking on it to add stops and control points (“Extend” turned off).

A reroute modification can apply to multiple trip patterns in a single direction as long as the patterns all contain the start and end stop in order; you will generally need to create one reroute modification for each direction of the route.

A few examples should help to illustrate how this modification works. Consider a baseline pattern passing through stops A -> B -> C -> D:

- To curtail this pattern at stop C, eliminating service to stop D, select C as the “start of reroute/extension.”
- To reroute this pattern from C to another stop X instead of D, select C as the “start of reroute/extension,” activate route editing, and click on stop X to add a segment from C -> X. Speeds and dwell times can be set on this added segment. Baseline speeds and dwell times between A and C are not affected.
- To extend this pattern backward, to originate at a stop Y, select A as the “end of reroute/extension,” activate route editing, and click on stop Y to add a segment from Y -> A. Speeds and dwell times can be set on this new segment. Baseline speeds and dwell times between A and D are not affected.
- To detour this pattern so that it serves a stop Z between B and C, select B as the “start of reroute/extension”, select C as the “end of reroute/extension”, activate route editing, click on the new segment to add a stop, and drag the added stop to Z. Speeds and dwell times can be set on this new segment. Baseline speeds and dwell times between A and B, and between C and D, are not affected.

1.4.8 Custom modifications

The custom modification type allows us to try out new development features that are not yet generally supported. *You should not use this modification type without consulting with your support team; an improper configuration may produce errors during analysis.*

1.5 Timetables

Add trip pattern and *Convert to frequency* modifications require simple timetables. By default, a newly added timetable or frequency entry will specify trips running every 10 minutes from 7 AM to 10 PM, every day of the week.

Each timetable or frequency entry allows for the specification of days of service, span of service, and frequency. You can add as many timetables as you need to specify different frequencies or speeds at different times of days. Overlapping windows can be specified, but keep in mind that this means that trips on *both* entries will operate at the specified frequencies (e.g. if you have a ten-minute frequency and a 15-minute frequency overlapping, there will be one set of vehicles coming every ten minutes and another, independent set coming every 15).

For *Add trip pattern* modifications, speed and dwell time values can be set for each timetable, either at the segment level or as overall average values. The user interface also displays travel times derived from these values. While segment-level running-time values can be modified, the speed values are what Conveyal Analysis actually saves and uses for calculation. Recalculated travel time values may differ slightly from explicitly entered values, due to rounding of speed values or if segment lengths change. Before analyzing scenarios, we recommend re-opening modifications with timetables and double checking that values reflect desired travel times.

For *Convert to frequency* modifications, speed and dwell time values are copied from the template trip. Oftentimes, travel time will vary throughout the day due to varying traffic and passenger loads, so it makes sense to choose a template trip that is representative of the time window for which you are creating frequency service.

1.5.1 Exact Times

You can specify whether the timetable represents an assumed headway, or represents the exact schedule, using the *Times are exact* checkbox. The default setting, with this box unchecked, should be used when a trip pattern's frequency is known but an exact schedule has not yet been defined (e.g. for a new service still in the planning stages).

For example, consider a timetable specifying that a particular *trip pattern* runs every 15 minutes from 9 AM until 7 PM. With the checkbox unchecked, the routing engine knows that vehicles depart the first stop on the route every 15 minutes between 9 AM and 7 PM, but has no information as to exactly when that will happen. For example, vehicles might leave at 9:02, 9:17, 9:32, and so on, or they might leave at 9:10, 9:25, 9:40, etc.; many of these possibilities will be simulated in order to get a complete picture of how different possible schedules might perform. See *Methodology* for more details.

When the complete schedule is defined, it is appropriate to activate the *Times are exact* checkbox. In this case, a single fixed timetable will be created, with the first departure at the start time, and then additional departures with exactly the specified frequency until (but not including) the end time. For example, in the scenario given above, the vehicles would be scheduled to depart at exactly 9:00, 9:15, 9:30 until 6:45 (but not at 7:00 because the end time is not included).

If the schedule is not known, but it is known that the schedules of two lines will be interrelated (e.g. using timed transfers or pulsed schedules), the *Phasing* feature may be enabled.

1.5.2 Copying Timetables

Timetable entries can be copied between *Add trip pattern* and *Convert to frequency* modifications.

Some users may find it convenient to use a single template *add trip patterns* modification that specifies commonly used service windows and frequencies. For example, you could create a “Base Timetables” *modification* and deactivate it from all *scenarios*. You could then add multiple timetables to this template, for example:

- AM Peak, Frequent Network: 7 AM to 9 AM, every 5 minutes, every day
- Midday, Frequent Network: 9 AM to 3 PM, every 10 minutes, every day
- PM Peak, Frequent Network: 3 PM to 6 PM, every 5 minutes, every day

- Base Weekday: 6 AM to 11 PM, every 30 minutes, Monday through Friday
- Base Weekend: 7 AM to 11 PM, every 60 minutes, Saturday and Sunday

When proceeding to add new routes with *Add trip pattern*, the appropriate timetables could then be copied from this template by clicking:

1.6 Phasing

The phasing feature can be used to tie the timetables of two *Add trip pattern* or *Convert to frequency* modifications together. For example, if you are creating a *pulsed system* where buses run infrequently but are timed to meet at specific transfer points, you might want to specify that while you don't know the exact schedules of any of the lines, they will all meet at a central point at the same time. Alternately, you might have a *branching service* where two lines share part of their alignment but then branch. If each branch runs every 30 minutes, you might want to write the schedules such that they provide service every 15 minutes on the shared alignment, rather than both coming at the same time.

You can accomplish both of these things using the phasing settings. We implement this by allowing you to lock the departures on one timetable or frequency entry to those of another, saying that vehicles on one frequency entry depart a stop a certain number of minutes after vehicles on another frequency entry (generally on a different modification) depart a stop (the same one or a different, nearby one).

The image below shows the settings for a pulse between two lines in Atlanta.

This line on Memorial Drive is phased at the stop “Memorial Drive SE at Moreland Ave SE” from the stop “Moreland Ave SE at Memorial Drive SE” on the timetable “Weekday” in the “Moreland” modification, with a phase of three minutes. This means that vehicles running on this timetable (7 AM-10 PM every day) will depart the stop “Memorial Drive SE” will depart “Memorial Drive SE at Moreland Ave SE” three minutes after vehicles on the “Weekday” timetable of the Moreland route depart the stop “Memorial Drive SE at Moreland Ave SE.”

First, you select one of the stops on this line to phase at (the “phase at stop”). You then can select a modification (either an Add Trips modification or an Adjust Frequency modification) and frequency entry or timetable to phase from (the “phase from timetable”); the dropdown shows the name of the modification, as well as the name of the frequency entry and details about it. You can then choose the stop on the other modification where the phase should be applied (the “phase from stop”).

Finally, you can specify the phase in minutes, which is how long after the other vehicle departs that this one should depart (at the last stop on either modification, the arrival time will be used in lieu of the departure time).

If you wish to create a pulse, generally you would want to specify a dwell time at the stop as well, to accommodate bidirectional transfers. Adjust frequency modifications may be paired with an *adjust dwell time* modification to accommodate this. There is not currently a way to set the dwell time at a specific stop in an add trips modification, but this feature is on the short-term development roadmap. The analysis engine requires at least a minute between alighting from one vehicle and boarding the next to make a connection, plus any walking time if the boarding and alighting do not occur at the same stop.

Pulsing will be disabled if either of the modifications involved is bidirectional (because we won't know how to properly phase the directions together), or if either of the frequency entries involved are specified using “exact times” (because the times are already set and can't be adjusted to match the other entry). It is possible to phase two frequency entries which do not have the same headway, although this should be used with caution. If one headway is a factor of the other, they will just be lined up so that they occur together when possible (for instance, vehicles on frequency entries with headways of 15 minutes and 30 minutes set up to pulse will arrive together every 30 minutes, with an additional trip from the 15 minute headway occurring in between). However, if they are not factors of each other, the results will not be as desired. For instance, phasing frequency entries with headways of 15 and 20 minutes will result in trips being aligned as desired once per hour, with other trips misaligned (for example, vehicles on the frequency entry with 20-minute headway might arrive at 7:00, 7:20, 7:40 and 8:00, while vehicles on the frequency entry with 15-minute headway might arrive at 7:00, 7:15, 7:30, 7:45 and 8:00).

1.7 Managing modifications

Modifications can be grouped by *project* and *scenario*, and different projects and scenarios can be compared against each other in analysis mode, giving you flexibility on how to use them. Depending on your use cases, different approaches may make sense.

If one user will be responsible for analyses in your region, involving a relatively small number of modifications, we recommend doing your work in one project and assessing the impact of different combinations of modifications by creating and using scenarios within that project.

If multiple users will be involved in editing scenarios, or if you want to assess more than 10 different combinations of modifications, which would make the list of scenarios annoyingly long, we recommend dividing the modifications among different projects. For example, one team member could code rail scenarios in Project A, another team member could code bus scenarios in Project B. Modifications can be *imported* between projects that use the same *GTFS bundle*; in this example, modifications from the two projects could be combined in a third Project C.

1.7.1 Toggling display of modifications

In the list of modifications on the initial view in editing mode, clicking the title of a modification will open it and allow you to edit it. To control whether each modification is displayed on the map, click

Stops and segments representing modifications are displayed on the map, using different colors to indicate their state relative to the baseline GTFS:

- : Added trip pattern
- : Removed trip pattern
- : Changed timetable (e.g. modified frequency, speed, or dwell time)
- : Unchanged (alignment is unchanged but the *trip pattern* is effected somehow, e.g. *Reroute*)

Projects start with only a “Default” scenario (plus a locked *Baseline* in which no modifications can be active). You can create additional scenarios expanding the list of scenarios, clicking the create button, and entering a name.

When the Scenario list is expanded, options next to each scenario allow you to: the modifications active in the scenario the scenario the scenario (not available for baseline or default scenario)

1.7.2 Exporting modifications

To see options for exporting scenarios from the top of the editing panel, click

A panel will then be shown with multiple options to download or share the scenarios in your project:

- all scenario details
- alignments of add-trip modifications
- new stop locations created in the scenario
- a summary of all modifications in a scenario, for printing or reference. Keep in mind that some browsers may not print more than 30 pages or so of a long report.

1.7.3 Importing modifications

To import modifications from another project or a shapefile, click

From another project

Occasionally, you may want to copy all of the modifications from one project into another. This may be useful to make a copy of a project, or to combine modifications developed by different team members into a single project (for instance, one team member working on rail changes and another on bus changes).

To do so, select a project in the upload/import panel and click

If you choose a project associated with the same GTFS bundle, all modifications will be imported; when there are multiple scenarios, the scenarios in the project being imported will be mapped directly to the scenarios in the receiving project (i.e. modifications in the first scenario will remain in the first scenario in the new project).

If you choose a project associated with a different GTFS bundle, only add-trip modifications will be imported.

From shapefiles

In general, it is best to create all modifications directly in Conveyal Analysis as it allows full control over all aspects of transit network design. However on occasion it may be desirable to import modifications from a [shapefile](#). If you have a shapefile containing route geometries, you can upload it to Conveyal Analysis and have it turned into a set of *Add trip pattern* modifications. Your shapefile will need attributes (columns) for each line's:

- name (e.g. route id)
- approximate headway in minutes
- approximate speed in kmph

The shapefile should contain only linear features. Points on the lines will be converted into control points in the modification and by default stops will be spaced uniformly along the line. If think you may want to edit the alignment later in Conveyal Analysis, it might be helpful to simplify complex geometries before uploading them.

To upload a shapefile from the upload/import panel (), click

After selecting and uploading a zipped shapefile, you should see the following options. You'll need to use the drop-down menus to identify the attributes from your uploaded shapefile that will be used to generate the modifications.

Finally, as shapefiles only contain the route geometry and not the stop locations, stops can be created automatically at a specified spacing. Stops may be explicitly added or moved after import.

1.7.4 Manually edit raw modifications (beta)

By clicking on “Customize modification” in the details pane for any modification, advanced users can copy, paste, and edit its raw JSON data.

1.8 A simple walkthrough

To create a simple *modification*, follow this quick-start guide:

1. If you are not at the initial view in editing mode, with your project name at the top of the side panel, click the icon on the sidebar.
2. Select a *project* or add a new one
3. Click:
4. Leave *Add trip pattern* selected as the **Modification Type**, type “New Route Example” as the **Modification Name**. Confirm these options by clicking on the green button, which will create a modification and open the modification details panel for it.

5. In the modification details panel, click
6. On the map, click to add stops for the new route. For more details on editing route alignments are see [Add trip pattern](#).
7. In the modifications details panel, click:
8. Optionally, edit the default timetable parameters (e.g. set 5-minute headways for weekdays between 7 and 9 AM) and add additional timetables. For more on this see [Timetables](#).
9. To save your changes and return to the main list of Modifications, click the icon at the top of the modification details panel.
10. Add more modifications or proceed to [analyze](#) your scenario.

1.9 Single point analysis

The main analysis page is for generating *isochrones* (travel time contours) from selected origins. To enter analysis mode, click the icon on the sidebar. To start an analysis, ensure a *project* and *scenario* are selected.

To retrieve results for the origin marker shown on the map, you can either move the marker to a new location or click the *fetch results* button at the top.

This will initialize a compute cluster which may take a minute to start up. If this is your first time performing an analysis with a given *bundle*, it may take some time to build the network. This only needs to be done once for each bundle. For more information, see [When starting an analysis, why does the “initializing cluster” message persist for so long?](#)

Once the compute cluster is initialized, you should see an isochrone displayed in blue around your point on the map. If you have selected an opportunity dataset, you will also see a chart showing cumulative accessibility results at selected time and percentile thresholds. You may also select a comparison project and scenario, which will be shown in red. Many other configuration parameters are described in [Options and configuration](#).

1.9.1 Isochrone map

After the server returns results, the map will show a blue *isochrone*. This represents the area reachable from the origin marker within a given travel time cutoff, to a given degree of certainty.

The *time cutoff* slider controls the travel time threshold between a range of one minute and two hours. The slider for *travel time percentile* controls the portion of departures within the time window that have to meet the travel time threshold. Reducing the travel time should smoothly decrease the size of the isochrone, as would increasing the travel time percentile. The default values are 60 minutes and 50th percentile. This would mean that the default isochrone boundary is drawn where exactly half of trips in the selected departure window would take exactly one hour.

The modifications displayed on the map are controlled in editing mode (See: [Toggling display of modifications](#)).

To change the origin of the analysis, drag the marker to a new location. Clicking on the map will display a box and whisker plot of the distribution of travel times from the origin to that location. For example, in the image below, the travel time varies between about 30 and 50 minutes depending on when one departs.

If multiple scenarios are being compared, the isochrone for the first scenario remains blue while the isochrone for the second is red. Thus, areas reachable under both scenarios are purple, areas reachable only under the first scenario are blue, and areas reachable only under the second scenario are red.

If an *opportunity dataset* is selected in the drop-down menu in the settings panel, the map will show gray dots representing the density of opportunities. For instance, if your selected opportunity data are jobs, there will be tightly packed dots in areas of dense employment, and less tightly packed dots elsewhere. One dot represents one or multiple

opportunities, and the scale may differ between zoom levels and opportunity datasets. For example, if at a given zoom level, one dot represents 4 jobs, at that same zoom level one dot might represent only two residents.

1.9.2 Analysis panel

The left panel has controls for the analysis and displays the access to opportunities afforded by the *scenario*. At the top of the panel, available scenarios and opportunity data layers are listed in drop-down menus. For example, you might be interested in how a given scenario provides access to jobs, access to schools, or some other variable of interest represented in an *opportunity dataset* you uploaded. Additional scenarios can be selected for comparison. A *baseline* scenario with no modifications (i.e. the unmodified *GTFS bundle* you uploaded) is automatically available for every project.

Charts of accessibility results

Directly below the comparison controls are readouts of the accessibility (number of opportunities reachable) from the chosen origin under the scenario and (if applicable) any comparison scenario.

The main display of accessibility results is the stacked percentile plot. The right portion of the plot shows the distribution of cumulative accessibility, i.e. the number of opportunities reachable given varying travel time cutoffs. The graph is not a single line, because there is variation in transit travel time depending on when a user of the transport system leaves their origin. Rather, the graph shows the number of opportunities given 95th, 75th, 50th, 25th, and 5th percentile travel time. The bottom of the shaded area is the number of opportunities which are almost always reachable, while the top is the number of opportunities that are reachable only in the best cases (e.g. when someone leaves their house at the perfect time and has no waiting time). The darkened line is the number of opportunities that are reachable at least half the time (i.e. have a median travel time of less than the travel time cutoff). For a more detailed explanation, see the [methodology](#) page.

When the cumulative plot is steep, areas with especially high opportunity densities (e.g. typical downtown areas for jobs) are reachable. Note that the Y axis is a square-root scale, so that the cumulative plot would be a straight line if both the opportunities and travel speeds radiating in all directions from an origin were uniform.

The currently-selected travel time cutoff is indicated by the vertical line on the plot.

To the left of the Y axis labels is a box-and-whisker plot. This shows the same information as the cumulative plot, but only for the currently selected travel time cutoff. The lowest whisker shows the number of opportunities accessible given 95th percentile travel time, the box shows the number of opportunities accessible given 75th, 50th and 25th percentile travel time, and the top whisker shows the number of opportunities reachable given 5th percentile travel time.

When multiple scenarios are selected, the charts will be slightly different, because they will include information for both scenarios.

Two box plots will be displayed, in red and blue, to the left of the axis. The blue box plot is for the first scenario, while the red one is for the second scenario. Above the chart, there is a selector that allows you to select whether to view the cumulative curves for the first scenario, the second scenario, or both (in which case the plots will be simplified and only the bands between the 75th and 25th percentile travel times will be shown, for simplicity).

Downloading

There are multiple options for downloading single-point analyses:

- saves the isochrone currently shown on the map. The downloaded file can be converted to other formats using a tool like [mapshaper](#). Note that these vector isochrones are interpolations of the [underlying analysis grid](#). They can be useful for visualizing results in GIS, but additional steps may be needed to prepare them for geoprocessing.

- saves the underlying travel time surface, a raster of travel times (in minutes) from the selected origin to the rest of the region. This raster has five bands corresponding to [time percentiles](#) of 5, 25, 50, 75, and 95. For geoprocessing, we often suggest using band 3 (the 50th percentile travel times) of this raw output.

1.10 Regional analysis

The *Single point analysis* interface also allows creating a *regional analysis*, which involves repeating an accessibility calculation for every location in a regular grid (See *Spatial resolution*).

1.10.1 Starting a regional analysis

To start a regional analysis, first set the appropriate parameters using the controls for a *Single point analysis*, and confirm that the *isochrones* and accessibility plots are as expected.

You can also choose geographic bounds for your regional analysis in the Advanced settings. By default, the entire region is analyzed, but for efficiency it is also possible to analyze a smaller area. You can set the bounds of the analysis by dragging the pins on the map, or by selecting an existing regional analysis and using the same bounds. If you plan to compare two regional analyses, make sure they have the same bounds and routing engine version.

When you have configured all of these options, click at the top of the panel and enter a name. Note that this button is disabled unless isochrones are displayed; checking single-origin isochrone results is a verification step that helps avoid heavy computation for analyses with invalid settings.

After a few seconds, you will see your regional analysis appear in the list with a progress bar. Since Conveyal Analysis is computing accessibility from every origin in the region, it can take a [few minutes](#) for each regional analysis to complete. Once a regional analysis is complete, it can be viewed by selecting it from the drop-down menu, which will take you to the [regional analysis view](#).

1.10.2 Inspecting results

Upon selecting a completed regional analysis, you will see a screen like the following:

The map shows the number of opportunities reachable from each location within the travel time cutoff specified when creating the regional analysis. Using the download button, you can save regional analysis results in a [GeoTIFF](#) raster format. These files can then be opened in a GIS to conduct additional analyses or create custom maps. Downloading results also allows you to see the raw [grid cells](#) used for analysis, rather than the smoother interpolated results shown in your browser.

1.10.3 Downloading regional results

Using the download button, you can save regional analysis results in a raster format (GeoTIFF). These saved files can be opened in GIS to conduct additional analyses or create custom maps. In QGIS, you will likely want to [style](#) the accessibility layer with a singleband pseudocolor scheme. If you prefer to work with the results as a regular grid of vector points, you may find the “Raster values to points” tool in the QGIS processing toolbox helpful. Downloading results also allows you to see the raw [grid cells](#) used for analysis, rather than the smoother interpolated results shown in your browser.

1.10.4 Comparing regional analyses

You can also compare two regional analyses from different projects in the same region. The map will show the differences in accessibility between the two analyses, with blue areas showing increased accessibility, and red areas

showing decreased accessibility, relative to the comparison analysis. Again, you can download raw results for the two analyses being compared for further styling and analysis in GIS.

1.10.5 Measuring aggregate accessibility

These regional analyses present a wealth of information, and maps of regional accessibility are frequently the best way to communicate the accessibility impacts of a transit plan. However, in some cases there is a need to summarize accessibility in a single aggregate metric. Conveyal Analysis allows aggregating the results of a regional analysis to *aggregation areas* (e.g. neighborhood, city council district, transit agency service area, or region). For a selected aggregation area, Conveyal Analysis can report the distribution of accessibility in that area, weighted by a value (e.g. population or households) at each grid cell origin.

Uploading aggregation areas

In order to accomplish this aggregation, you first need to choose a region and code it as a polygon Shapefile (unprojected WGS84, CRS 4326). The choice of region can have a significant impact on the final metric. If a very large region is chosen, where much of the region does not have transit service, there will be a large number of people with very low job access via transit, deflating the aggregate numbers. Conversely, if a small region is chosen, segments of the population that should be served by transit will be excluded from the aggregate metrics. Conveyal Analysis will not accept aggregation areas larger than 2 square degrees.

To create aggregation areas when viewing completed regional analysis, click

You will then see options to upload and process a shapefile:

Shapefiles used as aggregation areas must fit within a bounding box smaller than 250,000 sq. km, must not have any feature with area exceeding 2 square degrees, and should preferably use unprojected WGS84 coordinates. Be sure to select all components of the Shapefile (i.e. .shp, .dbf, .shx, .prj, etc.) when uploading. It may take some time to upload the files and process them, depending on their size and complexity.

If **Union** is selected, the union of the uploaded shapefile's features will be used as a single combined aggregation area.

If **Union** is un-selected, each feature will be used as a separate aggregation area, named with the value of the "attribute" specified in the upload options. For example, if you want to report a project's change in accessibility for each city council district, and you have a shapefile of with each district's name in a column "CD_Id," you would un-select "union" and type "CD_Id" as the attribute.

Viewing aggregate metrics

The choice of aggregation area can have a significant impact on the final metric. If a very large area is chosen, where much of the region does not have transit service, there will be a large number of people with very low job access via transit, deflating the aggregate numbers. Conversely, if a small area is chosen, segments of the population that should be served by transit will be excluded from the aggregate metrics.

To view aggregate metrics, select an uploaded aggregation area and what you want to weight by. For example, if you select "Workers total," your histogram will represent how many workers experience different levels of accessibility. If instead you select "Workers with earnings \$1250 per month or less," you will see how accessibility is distributed among workers with low earnings. Any opportunity dataset you have uploaded is available for use as weights.

Once you have an aggregation area and weight selected, you will see a display similar to the one below, showing a histogram of how many of the units you weighted by experience a particular level of access.

The plot is a histogram of the number jobs accessible within the full opportunity dataset (not just Brooklyn) for workers residing in Brooklyn (since this is a regional analysis of access to jobs, aggregated to Brooklyn, and weighted by workers). We can see that the distribution is bimodal. There are a large number of workers with relatively lower accessibility values (the left peak), most likely residing in outer Brooklyn further from the subway and from job

centers . There are also a number of workers with a higher level of accessibility, probably those residing in downtown Brooklyn and near Manhattan. This plot is conceptually similar to academic work on accessibility by population percentile (e.g. [Grengs et al. 2010](#)).

Below the histogram is a readout of the metric mentioned above: how many opportunities are accessible to a certain quantile of the population; in this case, 90% of workers residing in Brooklyn have access to more than 267,000 jobs. You can adjust the percentile by dragging the slider. The area of the histogram above the cutoff is highlighted. By setting it below 100%, you are effectively recognizing that some percentage of a given region may not be feasible to serve with transit, so accessibility metrics should not be penalized for the lack of access in those areas.

You can also view these aggregate metrics when comparing two regional analyses.

In the example below, the baseline scenario is shown in blue and an alternative scenario is shown in red, with a darker area where the two overlap. The horizontal axis is a scale of number of jobs accessible within 45 minutes. The height of each histogram bin represents the number of households who have the corresponding level of job access. The farther to the right the red distribution is, the higher the job access gains across households.

In the baseline, the 10th percentile household (marked by the leftmost vertical line) has access to fewer than 50 thousand jobs, while in the alternative scenario, the 10th percentile household (the second vertical line) has access to more than 141 thousand jobs. The large increase at the lower end of the distribution suggests the alternative scenario could help advance more equitable access to jobs.

Finally, Conveyal Analysis displays weighted mean accessibility, which represents the average accessibility experienced by all residents (or whatever unit you have weighted by) in the aggregation area. However, since it uses the mean, it is strongly affected by outliers, and may not be representative of the full range of accessibility experienced. In the Brooklyn example above, it falls between the two peaks, and does not reflect that many people have accessibility below that. Since aggregate accessibility frequently has a long right tail, with many residents with low to medium accessibility and few with very high accessibility, the mean is often higher than the accessibility experienced by a majority of the population. For these reasons, we do not recommend the use of this metric.

1.11 Options and configuration

Below the accessibility charts, different parameters for the analysis can be set:

The first panel allows the creation and use of **bookmarks**, which store particular analysis settings (e.g. origin location, type of opportunity, departure date and time, travel time cutoff, etc.). Once you have a set of settings you would like save, click . Once that is done, you can select a bookmark from the dropdown box to automatically fill in all of the settings from that bookmark. Bookmarks are shared by all projects in a region.

Next are selectors for **access modes** and **transit modes**; you can choose to perform your analysis with or without transit, and using walking, biking or driving. For instance, in the image above, a combination of walking and transit has been chosen. Note that traffic congestion is not taken into account in driving time estimates, though this may be a feature of a future release when more detailed datasets are available.

Next are the **date**, **from time**, and **to time**, which define the time period analyzed (i.e. the opportunities accessible by someone leaving the chosen origin point on the chosen date between the chosen times). The first time you open a project, these will default to the current date and 7:00 to 9:00. To avoid inadvertently introducing differences in results due to differences in service on different days, we recommend choosing a single date and using it for the duration of a project. You should check that the date chosen is sufficiently representative in the GTFS feeds you are using (e.g. a non-holiday weekday).

You will also need to select a **Routing engine** version, which should default to the highest available version of [Conveyal R5](#).

The final option is the **Percentile of travel time**. In single-point analyses, this is rounded to one of five pre-defined values (5, 25, 50, 75, and 95). For more information, see [methodology](#).

1.11.1 Advanced settings

Maximum transfers is an upper limit on the number of transfers that will be considered when finding optimal trips.

Egress mode defaults to walking. Non-walk egress modes may require a lengthy initial computation step; contact your support team if you need assistance.

Walk speed, Bike speed, Max walk time, and Max bike time apply to each access, transfer, and egress leg of trips when public transport is enabled. Any requested values will be applied for access and direct (walk- or bike-only) legs. For transfer and egress legs, the requested values will be applied, up to a distance limit of 2 km for walk and 5 km for bicycle. At reasonable speeds, the default value for Max walk time allows trips requiring a 20 minute walk to initial boarding stops, 20 minute walks at each transfer, and a 20 minute walk from alighting stops to final destinations; it also allows trips requiring a 20 minute walk from origin to destination, when walking directly is faster than using public transport.

In certain cases, these limits may lead to counterintuitive results. For example, consider a destination that is about a 30-minute walk north of a given origin; an east-west bus route 15 minutes north of the origin would let travelers “circumvent” a 20-minute max walk time by walking 15 minutes to the bus, riding one stop, then walking 15 minutes to the destination.

When public transport is disabled, these limits are not applied and the travel time slider below the accessibility chart effectively controls the maximum time for the requested direct mode.

If your scenario includes frequency-based routes (either in the baseline GTFS or in modifications with **exact times** not selected), **simulated schedules** controls the number of schedules simulated for sampling. The sampling process introduces random uncertainty, so you may see results change slightly when you repeatedly request accessibility results. When comparing regional analyses that include frequency-based routes, you may see small unexpected increases or decreases attributable to this random noise. Final results will be more accurate when **simulated schedules** is set to higher values, but computation will take longer. For quick, interactive analysis, we recommend setting it to 200, whereas, for final analysis, we recommend setting it to 1000. For more information, see [methodology](#).

If your scenario does not include frequency-based routes, there is no need to simulate schedules, so the requested number of simulated schedules is ignored. In other words, when departure times are explicitly specified for all trips in a scenario, only that single fully specified set of exact departure/arrival times needs to be tested, which speeds computation and eliminates random noise from sampling.

1.11.2 Errors and warnings

Occasionally, analysis will fail because there is an error in a scenario. When this occurs, error messages will be displayed detailing the issues, as shown below. One simply needs to return to the modification editor and correct the errors in the relevant modifications.

In other cases, the scenario may generate a warning, for instance if you remove more time from a segment when speeding it up than the length of that segment. This is not necessarily an error, but may require attention.

1.12 Developers Guide

1.12.1 Installing your own copy

Instructions on configuring and running a local copy of Conveyal Analysis are available [here](#).

1.12.2 Updating this manual

If parts of this manual are unclear or confusing, we'd welcome your feedback. At the top of each page is a link to "Edit on GitHub." Clicking that link, when logged into GitHub, will allow you to edit the markdown files for this manual and propose changes by opening a pull request.

1.13 Frequently Asked Questions

1.13.1 How does Conveyal Analysis calculate travel times?

Conveyal Analysis calculates door-to-door total travel time, without any subjective impedance factors. Travel times are always computed by finding actual paths through a full street and (where applicable) public transport ("transit", in American English) network. By default, Conveyal Analysis uses the centers of [high-resolution grid cells](#) as destinations.

For analyses with transit disabled, the total travel time from an origin to a destination includes time spent:

1. Walking from the selected origin to the nearest point on the street network
2. Traveling to the point on the street network nearest to the destination, using the selected direct mode (walk, bicycle, or car)
3. Walking from the street network to the destination

For analyses with transit enabled, the total travel time from an origin to a destination includes time spent:

1. Walking from the selected origin to the nearest point on the street network
2. Reaching nearby transit stops via the street network, using the selected access mode (walk, bicycle, or car)
3. Waiting to board transit (initially and at transfers)
4. Riding in transit vehicles
5. Walking via the street network between transit stops for transfers
6. Traveling from transit stops to the point on the street network nearest to the destination, using the selected egress mode
7. Walking from that point on the street network the directly to the center of the grid cell.

If it is faster to reach a destination directly, without using transit, the routing engine will replace the second step above with direct travel along the street network to the point on the street network nearest the to the destination, then skip to the final step. Public transport routing uses actual schedules from the imported GTFS feeds. For every origin, Conveyal Analysis finds a large number of travel times to every destination grid cell. Different travel times may be found for each departure time (each minute in the specified time window). If new transit routes have been added with frequencies but without specific departure times, many different possible schedules are also tested producing hundreds or thousands of possible travel times (according to the "simulated schedules" advanced parameter). All these different travel times imply a statistical distribution, from which a certain percentile of travel time, set in the user interface, can be read.

1.13.2 When starting an analysis, why does the "initializing cluster" message persist for so long?

First, the main Analysis server must request and initialize a computation cluster from Amazon Web Services. For scalability, we start a "worker" server for each *transport network* being analyzed. This means that even if you are

already successfully fetching analysis results for a project, a new server will be needed if you switch to a project associated with a different GTFS bundle or region, or if you change the routing engine.

To complete regional analyses quickly, we can clone hundreds of servers within a cluster for a transport network. We haven't yet built the user interface to launch additional servers from within Conveyal Analysis, so for the moment, get in touch with your support team if you need to speed up analyses.

1.13.3 Why are travel times longer than expected?

Because Conveyal Analysis includes waiting time, which can vary greatly depending on when one departs, total travel times may be longer than you might initially expect. Varying the travel time percentile can help you determine what role service frequency plays in waiting times (and therefore total travel times) for your network. The fifth percentile usually reflects trips with close-to-minimal waiting times (e.g. when passengers consult schedules and adjust the start of their trips to minimize overall travel times).

Another potential explanation for unexpectedly long travel times is that origins, destinations, and stops are all linked to the nearest edge in a network's street layer. This linking behavior can lead to unexpected results, especially for stops, origins, or destinations in areas with sparse street or pedestrian networks. Consider these examples:

- You place the origin directly on a transit stop with fast, frequent service, but travel times using the service show an unexpected delay. Because total travel time includes walking from the origin to the nearest street or pedestrian path in the baseline street layer, then from the street layer (back) to the transit stop, there may be an unexpectedly long access time if the stop is far from a nearby street.
- You notice a very long transfer time between two nearby stops. If these stops are linked to different street edges, and the street edges are only connected via a long, circuitous path, the walking time for the transfer will be unexpectedly long.

To avoid such issues, we suggest placing stops close to streets or pedestrian paths, and snapping to stops that exist in the baseline transit layer when feasible.

Uploaded OpenStreetMap data with poorly connected subnetworks may also lead to issues. If you notice strangely non-contiguous isochrones, they may reflect destinations being linked to such a subnetwork (e.g. a network of walking paths in a park that only connects to the rest of the street network at a few locations).

1.13.4 Who can see and edit my projects?

All authorized users within an organization have access to that organization's regions, projects, and scenarios. Multiple users should not edit or analyze the same project concurrently. If multiple users try to edit the same modification simultaneously, or if you have Conveyal Analysis open in multiple browser tabs, you may see "data out of date" errors.

1.13.5 How long should it take to start a single-origin analysis?

Three steps take place when starting a new single-origin analysis.

First, the main Analysis server must request and initialize a compute cluster from Amazon Web Services. For scalability, we start a "worker" server for each [transport network](#) being analyzed. This means that even if you are already successfully fetching analysis results for a project, a new server will be needed if you switch to a project associated with a different GTFS bundle or region, or if you change the routing engine version.

Second, the worker server needs to set up a transport network. It first checks whether the required transport network is already built and available for download from Amazon S3. If it is, the server downloads it and proceeds to step 3. If a pre-built network cannot be downloaded, the server downloads the required input files (OSM extract and GTFS bundle). It then builds the transport network by combining the road and transit layers, which can be a lengthy process

(on the order of 10 minutes for large regions, and an hour for very large regions with dense networks). To avoid having to repeat this step, the server will upload the built network to S3 for future use.

Third, once a transport network is downloaded or built, certain caching operations related to egress mode (e.g. walking or biking) need to be completed.

Recent routing engine versions provide status updates on these steps, including:

- *Starting routing server. Expect status updates within a few minutes.*
 - Trigger: analysis using a project that has not been analyzed within the last hour
 - Expected duration: 2-3 minutes
- *Building network...*
 - Trigger: analysis using a new GTFS bundle, scenario, or routing engine version
 - Expected duration: 5-50 minutes, depending on size of transit network
- *Linking grids to the street network and finding distances...*
 - Trigger: analysis using an updated egress mode or scenario
 - Expected duration: up to a few minutes, depending on extent of region
- *Performing analysis...*
 - Expected duration: a few seconds

By default, servers automatically shut down after one hour of inactivity. If you to override this default shutdown behavior for critical analyses, get in touch with your support team.

1.13.6 How long should each regional analysis take?

The time required to complete a regional analysis is a function of the size of the region, number of origins, and size and complexity of the transit network. Generally, computation is slower for routes specified as add-trip modifications than for routes specified in the baseline GTFS bundle. If frequency-based routes are present in the scenario's add-trip modifications or the base GTFS, compute time will also depend on the number of simulated schedules (Monte Carlo draws).

Additional servers start automatically once a results for few origins in a regional analysis have completed successfully. The number of additional servers is a function of the number of origins, with a default cap set to approximately 250. If you need a higher cap to complete a large batch of analyses quickly, contact your support team.

1.13.7 How do I model corridor upgrades where a new trunk is added and existing services are converted to feeder routes?

The new trunk service can be represented with an [add-trip modification](#). Existing services can be curtailed using [reroute modifications](#). You will need separate modifications for each direction of each curtailed route, setting the start/end stop to the stop at which the route will be curtailed as a feeder to the new trunk.

1.13.8 What new features is Conveyal working on?

Conveyal releases new features and improvements for the cloud-hosted version of Conveyal Analysis every few months. We regularly discuss with customers to understand how they use the system and identify features they would find useful. We create a roadmap containing features that we can realistically implement, considering long-term reliability and maintainability. If a customer wishes to pay for development effort or use development/support hours bundled in their contract to cover additional software development, features can be prioritized.

1.14 Methodology

This is a summary of the accessibility indicators used in Conveyal Analysis. For more details, please see Conway, Matthew Wigginton, Andrew Byrd, and Marco van der Linden (2017). “[Evidence-Based Transit and Land Use Sketch Planning Using Interactive Accessibility Methods on Combined Schedule and Headway-Based Networks](#)”

1.14.1 Spatial resolution

As a shared geographic foundation for regions, opportunity datasets, and accessibility indicators, Conveyal Analysis uses a regular grid of cells. Each cell measures approximately 300 meters by 300 meters, depending on the latitude (with larger cells at the equator and smaller cells near the poles). These grid cells are not spaced at a round number of meters. Instead they are one-unit steps in the same global projection used for map display in our web interface. This helps accelerate map display, while eliminating the complexity of switching to different local projections. Large regions may contain more than a million such cells, reflecting a much finer resolution than typical regional travel demand models. Conveyal implemented this approach to counter the effects of the “modifiable areal unit problem” and in response to systematic bias we observed when using the center points of polygons defined by roads, we have made an intentional methodological choice to perform all measurement on regular grids whose geometry is independent of roads and land parcels.

Opportunity datasets from uploaded point or polygon files, as well as LODES imports, are resampled into this standardized regular grid. For example, if a large traffic analysis zone in an uploaded shapefile has 10 thousand jobs, those jobs are dispersed throughout the grid cells intersecting that zone, according to the areal proportion of the intersection of each cell with the zone.

1.14.2 Accounting for variability

Travel times by transit are often highly variable, depending on when travelers start their journeys and the timetables of routes they use. To account for this variability, we calculate the distribution of travel times experienced in a given departure window, then select specific values from this distribution according to a parameter we call the time percentile. This percentile ranges from 0 to 100 and represents how reliably a destination is reachable for travelers starting journeys in the selected departure window.

In general, this approach allows us to more appropriately characterize travel times between origin-destination pairs served by many alternative routes (e.g. a corridor with many overlapping bus services) than blanket half-headway assumptions. Note, however, that it assumes travelers have perfect knowledge of schedules in the network and act accordingly to minimize their travel time. In some cases, this assumption may not be appropriate. Travelers without access to complete schedule information may be inclined to board the first vehicle that arrives, even if others would allow them to arrive at their destination sooner.

We use the median or other percentile travel time, rather than the mean, because it can handle travel times that are infinite or undefined (for example, because the only route to a particular destination uses a bus that stops running before the end of the time window). For more details, please see Conway, Matthew Wigginton, Andrew Byrd, and Marco van der Linden (2017). “[Evidence-Based Transit and Land Use Sketch Planning Using Interactive Accessibility Methods on Combined Schedule and Headway-Based Networks](#).”

Others calculate the accessibility at each departure time and taking a mean of accessibility as is done by the University of Minnesota [Accessibility Observatory](#). Conveyal Analysis does not use this approach because the travel time to each opportunity should be treated independently. Consider a situation of a small town situated between two major cities each with 500,000 jobs, with hourly train service to each city. Suppose that we’re interested in the number of jobs reachable in one hour, given departure during a time window of 8:00 to 9:00 AM. Further suppose that, due to how the train schedules are written, it is possible to commute to all jobs in the first city in under an hour if you leave between 8:00 and 8:15, and all the jobs in the second city are accessible if you leave between 8:30 and 8:45, and at other times no jobs are accessible. This corresponds to hourly trains to each city, leaving at 8:15 in one direction and 8:45 in the other, and needing 45 minutes to travel to the respective city. Using the average of accessibility, this location would

have an accessibility of 250,000, because $\frac{1}{4}$ of the time 500,000 jobs are accessible in the first city, $\frac{1}{4}$ of the time 500,000 jobs are accessible in the second city, and $\frac{1}{2}$ of the time no jobs are accessible. Thus the accessibility value would be 250,000 even though there is no job that can be reached in an average travel time of less than 60 minutes.

1.14.3 Time percentile

Assuming service operates with perfect schedule adherence, the time percentile can be interpreted as how inclined travelers are to adjust their departure time based on transit schedules. Low time percentile values imply travelers are flexible and can start their journeys at times that minimize their waiting and overall travel time, which is appropriate for passengers accustomed to relying on schedules (e.g. commuter rail riders). Higher time percentile values are appropriate for passengers who expect “turn-up-and-go” service (e.g. users in dense areas with frequent service).

For most analyses, the default value of 50 is acceptable, though you may also want to consider 25th or 75th percentile travel times, depending on whether users are likely to use schedules or start trips randomly.

We assume passengers have perfect information about all schedules in the system and that they choose routes to minimize their total travel time. We also assume that passengers can board all vehicles, so we do not account for denied boardings due to vehicles or facilities that are inaccessible to people using wheelchairs, crowding, etc. Some routes may have frequencies/headways defined, but not explicit timetables, which can introduce uncertainty into results. Our routing engine (called R5) mitigates this uncertainty by simulating feasible timetables for all headway-based routes. For example, if the Red Line is set to depart Airport Station every ten minutes, without explicit timetables, one simulated timetable could include departures at 7:00, 7:10, 7:20... , and another simulated timetable could include departures at 7:02, 7:12, 7:22... . The number of simulated schedules parameter in the analysis settings determines how many simulated timetables are used, rounding up to an integer number of schedules per minute in the departure window.

Example

Consider travelers with the following hypothetical journey:

- Walk 10 minutes from their origins to a stop served by a single route
- Wait at the stop
- Board the first vehicle of the route that arrives at the stop, and ride it for 30 minutes
- Alight directly at their destinations

These travelers have at least 40 minutes of travel time, plus a waiting time that varies depending on when they begin their journey with respect to the route’s timetable.

Say these travelers want to begin their journeys sometime between 7:00 and 8:40 AM (a departure window of 100 minutes), and that they will only consider their destination reachable if their door-to-door travel time is less than 45 minutes. If the route departs the stop once per hour, at 7:00 and 8:00, there will only be five minutes of the departure window for which the destination is reachable. Travelers can start their journeys between 7:45 and 7:50, arriving at the station between 7:55 and 8:00 to board the 8:00 departure. The destination will only be considered reachable at time percentiles up to 5% (5 of the 100 possible departure minutes in the departure window). That level of reliability may be acceptable for travelers who are willing to structure their travel around schedules for services with low frequency.

If instead, the route departs the stop every 10 minutes, starting at 7:00, there will be fifty minutes of the departure window for which the destination is reachable (beginning trips between 7:05 and 7:10, or between 7:15 and 7:20, or between 7:25 and 7:30, etc.). The destination will be considered reachable at time percentiles up to 50%.

Finally, if the route departs every 5 minutes, the destination will be reachable within the 45-minute cutoff 100% of the time. Even if travelers just miss a vehicle, the next one will depart in 5 minutes, and they will arrive at their destination without exceeding the time cutoff.

1.15 Glossary

Several terms have a special or specific meaning in the context of Conveyal Analysis. While we hope that their meaning is usually clear enough from the context, this glossary is provided for reference.

baseline network The baseline network is the transport network (i.e. street network plus a GTFS bundle) without any associated modifications. It is a useful point of comparison during analysis and a starting point for most modifications.

GTFS

GTFS feed [General Transit Feed Specification](#), a format for transit network and schedule data. Most transit agencies produce GTFS feeds to power customer-facing trip planning applications, but they are also useful for analysis.

GTFS bundle A GTFS bundle is a grouped set of one or more GTFS packages associated with your region. A region can have multiple bundles and a bundle could include a single agency's GTFS or those of several adjacent or overlapping agencies. The bundle is used to represent the baseline transit service and as a starting point for some types of modifications.

isochrone An [isochrone](#) is a boundary around a given origin point defined such that travel from the origin to any point on the boundary takes an equal amount of time. Conveyal Analysis calculates travel time over a departure window, so the travel time used is actually a selected percentile in a distribution of travel times.

modification A modification is an alteration of some kind made to scheduled transit services. e.g. the removal of a line, a rerouting, a new service etc. These must be applied in scenarios to be used in analysis.

opportunity dataset An opportunity dataset is a spatial dataset with one or more numeric fields representing a count of opportunities (destinations) at particular locations in your *region*. See [Uploading an opportunity dataset](#).

project A project is means of associating scenarios and modifications with a particular GTFS bundle. A project is associated with only one bundle which cannot be changed after it is created.

region The region is the rectangular area used for accessibility analysis. It should geographically contain any scheduled transit services and network modifications.

routing engine [Conveyal R5](#) is the core computational software behind Conveyal Analysis. It performs one-to-many searches on multimodal (transit/bike/walk/car) transport networks.

scenario A scenario is a set of modifications packaged together. For example you might create a scenario, called "proposed service cuts" in which several modifications reduce service separately on lines A, B and C.

segment A portion of a [trip pattern](#) consisting of a transit vehicle traveling between two consecutive stops.

transport network A routable network with transit and road layers, built by a specific version of R5. The transit layer is built from a GTFS bundle and the road layer is built from the OpenStreetMap network associated with a region.

trip A single vehicle trip, representing a single vehicle visiting the stops on a particular trip pattern at a particular set of times.

trip pattern A trip pattern is a sequence of stops visited by a transit vehicle. A trip pattern can have many trips, and a route consists of a one or more trip patterns. A basic two-way route might consist of two trip patterns - e.g. one for inbound service, one outbound. Each might make stops on different sides of the street and/or in a different order.

B

baseline network, [28](#)

G

GTFS, [28](#)

GTFS bundle, [28](#)

GTFS feed, [28](#)

I

isochrone, [28](#)

M

modification, [28](#)

O

opportunity dataset, [28](#)

P

project, [28](#)

R

region, [28](#)

routing engine, [28](#)

S

scenario, [28](#)

segment, [28](#)

T

transport network, [28](#)

trip, [28](#)

trip pattern, [28](#)