

AUTOMATICALLY DETERMINING OPTIMAL TRANSPORT SERVICE LOCATIONS FOR POINTS OF INTEREST FROM NOISY MULTIMODAL DATA

TECHNICAL FIELD

[0001] Various aspects of this disclosure generally relate to transport services, and more particularly, to the determination of optimal transport service locations.

BACKGROUND

[0002] An on-demand transport service provider (e.g., a driver) receives requests from clients for transportation of one or more persons or items from an initial location to a destination. Typically, the client designates a location for the driver to pick up the one or more persons or items and begin the trip; the driver then determines a location to drop off the persons or items at the destination to complete the trip. The manual determination of pickup and drop-off locations can often yield pickup and drop-off locations that are not optimized for actual road conditions, or may delay pick up time or total trip time due to poorly chosen pickup and/or drop-off locations. Manually finding suitable pickup or drop-off locations for every possible point of interest (POI) is tedious and un-scalable to curate, and unpleasant for the client and the driver.

SUMMARY

[0003] The following presents a simplified summary in order to provide a basic understanding of various aspects of the disclosed invention. This summary is not an extensive overview of all contemplated aspects, and is intended to neither identify key or critical elements of all aspects nor delineate the scope of any or all aspects. The sole purpose is to present some concepts of one or more aspects in a simplified form as a prelude to the more detailed description that is presented later.

[0004] A solution to automatically infer or discover transport service locations (e.g., pickup locations, drop-off locations, etc.) specific to a point of interest or area of interest (AOI) that are suitable or optimal for both the service provider and the requester is provided. These locations may be inferred from a multitude of signals including noisy

Global Positioning System (GPS), images, videos, audios, e-wallet transactions, telematics data (e.g., door opening, closing, seat belt usage from car sensors), and user interaction data from the provider and requester application programs. The audio, visual, textual, and location data may come from the driver application program and/or the passenger application program. Although transport service is illustrated as an example in which the solution may be applied, the solution may be extended and generalized to other services such as food delivery and logistics.

[0005] In one aspect of the disclosure, a method, a computer-readable medium, and an apparatus for transport service are provided. The apparatus may receive a plurality of transport service transactions associated with a point of interest entity. For each transport service transaction of the plurality of transport service transactions, the apparatus may determine a transport service location at which the transport service transaction is executed. The apparatus may cluster the transport service locations determined for the plurality of transport service transactions. The apparatus may determine one or more candidate transport service locations for the point of interest entity based on the clustering. The apparatus may provide the one or more candidate transport service locations to a client or a service provider associated with a transport service transaction that is to be executed at the point of interest entity.

[0006] To the accomplishment of the foregoing and related ends, the aspects disclosed include the features hereinafter fully described and particularly pointed out in the claims. The following description and the annexed drawings set forth in detail illustrate certain features of the aspects of the disclosure. These features are indicative, however, of but a few of the various ways in which the principles of various aspects may be employed, and this description is intended to include all such aspects and their equivalents.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1 is a diagram illustrating an example of transport service transaction.

[0008] FIG. 2 is a flowchart of a method of automatically determining optimal service locations for point of interests.

[0009] FIG. 3 is a diagram illustrating an example of a general framework of how other data sources may be used for inferring transport service locations.

[0010] FIG. 4 is a diagram illustrating an example of the clustering result for a POI at 1K samples using certain fixed clustering parameter.

[0011] FIG. 5 is a diagram illustrating an example of the clustering result for the POI described above in FIG. 4 at 20K samples (e.g., 1-month dataset) using the same fixed clustering parameter.

[0012] FIG. 6 is a chart illustrating an example of the cumulative distribution function of 5-NN distances (in km) for a POI.

[0013] FIG. 7 is a diagram illustrating an example of clustering result for the POI described above in FIGS. 4 and 5 using the methods described above to address the issue of traditional clustering algorithm.

[0014] FIG. 8 is a flowchart of a method of transport service.

[0015] FIG. 9 is a conceptual data flow diagram illustrating the data flow between different means/components in an exemplary apparatus.

[0016] FIG. 10 is a diagram illustrating an example of a hardware implementation for an apparatus employing a processing system.

DETAILED DESCRIPTION

[0017] The detailed description set forth below in connection with the appended drawings is intended as a description of various possible configurations and is not intended to represent the only configurations in which the concepts described herein may be practiced. The detailed description includes specific details for the purpose of providing a thorough understanding of various concepts. However, it will be apparent to those skilled in the art that these concepts may be practiced without these specific details. In some instances, well-known structures and components are shown in block diagram form in order to avoid obscuring such concepts.

[0018] Several aspects of providing transport service will now be presented with reference to various apparatus and methods. The apparatus and methods will be described in the following detailed description and illustrated in the accompanying drawings by various blocks, components, circuits, processes, algorithms, etc. (collectively referred to as “elements”). These elements may be implemented using electronic hardware, computer software, or any combination thereof. Whether such elements are implemented as hardware

or software depends upon the particular application and design constraints imposed on the overall system.

[0019] By way of example, an element, or any portion of an element, or any combination of elements may be implemented as a “processing system” that includes one or more processors. Examples of processors include microprocessors, microcontrollers, graphics processing units (GPUs), central processing units (CPUs), application processors, digital signal processors (DSPs), reduced instruction set computing (RISC) processors, systems on a chip (SoC), baseband processors, field programmable gate arrays (FPGAs), programmable logic devices (PLDs), state machines, gated logic, discrete hardware circuits, and other suitable hardware configured to perform the various functionality described throughout this disclosure. One or more processors in the processing system may execute software. Software shall be construed broadly to mean instructions, instruction sets, code, code segments, program code, programs, subprograms, software components, applications, software applications, software packages, routines, subroutines, objects, executables, threads of execution, procedures, functions, etc., whether referred to as software, firmware, middleware, microcode, hardware description language, or otherwise.

[0020] Accordingly, in one or more example embodiments, the functions described may be implemented in hardware, software, or any combination thereof. If implemented in software, the functions may be stored on or encoded as one or more instructions or code on a computer-readable medium. Computer-readable media includes computer storage media. Storage media may be any available media that can be accessed by a computer. By way of example, and not limitation, such computer-readable media may include a random-access memory (RAM), a read-only memory (ROM), an electrically erasable programmable ROM (EEPROM), optical disk storage, magnetic disk storage, other magnetic storage devices, combinations of the aforementioned types of computer-readable media, or any other medium that can be used to store computer executable code in the form of instructions or data structures that can be accessed by a computer.

[0021] FIG. 1 is a diagram 100 illustrating an example of a transport service transaction. In the example, a client may request a service provider 106 to pick up or drop off a rider 104 at a POI 102. In some embodiments, the client may be the rider 104. The POI 102 may occupy a large area. Therefore, a precise transport service location 108 may need to be determined in order to facilitate the transport service transaction.

[0022] Understanding precise service locations for pickup, drop-off, and delivery is crucial for ride-hailing and delivery services. While most mapping solutions have POI information, they rarely have detailed information about various access points, entrances and exits etc., which is vital to complete the first mile and last mile segments of the service successfully. Towards this goal, a machine learning based solution is provided to automatically infer the service locations (e.g., given by latitude and longitudinal coordinates on the map) by mining GPS trails of historical bookings.

[0023] Inferring the precise service locations helps to understand the exact locations of pickups and drop-offs for a POI, and hence the travel times corresponding to a trip originating or ending at the POI. Further, inferring the precise service locations helps to uncover new locations that are preferred or selected by passengers. Thus, inferring the precise service locations helps to better estimations of travel times, as well as improve both passenger and driver experience.

[0024] In this disclosure, an automated method to infer preferred service locations for POIs is provided. While traditional implementation uses GPS probes extensively, some embodiments of the solution may be easily extended for processing a multitude of signals including images, videos, voice, text, telematics, and human feedback. Whereas most traditional methods rely solely on GPS data for POI inference, a solution is provided to use a variety of multi-modal data coming from audio, video, text (e.g., passenger-driver chat), and telematics logs that can all be combined to provide a more precise inference of the service locations. In some embodiments, a method is provided to integrate human feedback when uncertainty is high in the inference of the service locations.

[0025] FIG. 2 is a flowchart 200 of a method of automatically determining optimal service locations for point of interests. In some embodiments, the method may be performed by an apparatus (e.g., the apparatus 902/902' shown in FIG. 9 or FIG. 10).

[0026] At 202, the apparatus may combine noisy GPS data with driver application program interaction data (e.g., when the driver notifies the system that the vehicle has arrived; the drivers are not enforced to perform this action as accurately as possible) to infer the point where each driver triggered a pickup/drop-off event for each POI entity associated with that pickup/drop-off event.

[0027] Relying on only GPS probe from driver application program interaction to infer optimal service locations for a POI may not be always reliable, due to variations

in when and where the drivers decide to indicate that they have picked up a passenger (an action which is unenforceable), inherent noise and uncertainty of GPS signals. Hence, in some embodiments, the driver interaction data and/or GPS signal may be augmented with additional signals including, but not limited to, images, videos, speech, telematics and inertial motion sensors such as accelerometers, gyroscopes, etc.

[0028] FIG. 3 is a diagram 300 illustrating an example of a general framework of how other data sources may be used for inferring transport service locations. In the example, those other data sources may include environment 302, driver and passenger mobile phone 304, and vehicle 306. The environment 302 may provide (at 308) audios, videos, and images. The driver and passenger mobile phone 304 may provide (at 310) application program interaction data, GPS data, inertial measurement unit (IMU) sensor data, and text (e.g., passenger-driver chat). The vehicle 306 may provide (at 312) doors and seat belts telematics data. Data from different sources may be combined at 320 to form a signal resolution. At 322, final output of the inferred transport service locations may be obtained.

[0029] In order to address security concerns and resolve post-ride disputes, in-facing cameras, and microphones have been used to record data and perform post-mortem analysis. However, this data may be more actively used to serve other purposes. In some embodiments, an in-facing dashboard camera equipped with a face detection and/or recognition system may be used to detect faces inside the car to detect and record when passengers on-board and disembark the vehicle. When the algorithm detects a new face inside the cabin, this can indicate the beginning of a new ride. Similarly, absence of a previously detected face can indicate the end of the ride. The time instants and the GPS signals corresponding to these events may be combined to infer and refine the optimal service locations and the duration of the trip.

[0030] Typically, a ride hailing journey starts with driver and passenger greeting each other, the driver verifying the passenger's name and destination, and the passenger agreeing to the details. This also co-occurs with door closures that also have a unique audio signature. In some embodiments, by analyzing these signals, an audio signature may be derived to create canonical forms of such audio signals (similar to wake words used in personal assistants). These signals, in combination with the GPS data, may be used to identify the exact service locations.

[0031] Over the years, telematics sensors in cars have improved tremendously both in accuracy and coverage of activities that happen with the car. Telematics sensors may include an on-board GPS (typically more accurate than phone GPS receivers), sensors attached to rear doors, acceleration, brakes, child safety lock, wipers, etc. In a ride-hailing scenario, a driver may serve a sequence of rides continuously, in which case, open and closure of rear and trunk doors may indicate passenger on boarding or disembarking the vehicle. In some embodiments, the timestamps of these activities (as indicated by telematics sensor signals), in conjunction with the recorded GPS signals, may be used to infer the service locations more accurately.

[0032] In some embodiments, images and video of the pickup locations may be matched with an existing database of images of the city to narrow down the exact location within the POI. This is useful mainly when GPS signals are noisy and choppy, especially in indoor pickup areas.

[0033] Passengers often interact with the map to select a pin location for either pickup or drop-off. These “manual pins” are used typically when passengers do not find a pickup point nearby. In some embodiments, when (unsupervised machine learning) clustering algorithms are applied over manual pins aggregated over time, patterns of usage may be revealed, which in turn can be used to create new POIs. This may further identify potential blind spots in the POI database that can be fixed by more thorough mapping.

[0034] Passengers often use text conversations within the ride-hailing application program to provide additional information about the POIs. For example, a passenger might say, “waiting at the pillar marked 3”. In some embodiments, mining the chat logs in association with the POI and booking may reveal additional attributes about a POI and potentially help to create new POIs.

[0035] Referring back to FIG. 2, at 204, the apparatus may filter for outliers in the data due to excessively noisy GPS and/or incorrect driver interaction.

[0036] At 206, the apparatus may cluster together such pickup/drop-off points for each associated POI entity using a clustering algorithm, where the parameters of the algorithm are dynamically adjusted depending on the structure of the point dataset for each POI.

[0037] In some embodiments, the clustering algorithm may include several parameters, e.g., the maximum radius of the neighborhood from a point, and the desired minimum cluster size. Typically, the distribution of pickup signals (GPS positions corresponding to the driver's button press) at each POI can vary widely. Whereas popular locations have high density and large variance of occurrences, less known locations (e.g., residential POIs) tend to have sparse concentrated occurrences. Hence, fixed value for clustering parameters may not work in all cases. A conservative value for the clustering parameters may tend to create more than necessary number of clusters for the same POIs while relaxing these parameters may merge the necessary clusters, and furthermore, fail to remove outliers effectively thus impacting the final results negatively.

[0038] FIG. 4 is a diagram 400 illustrating an example of the clustering result for a POI at 1K samples using certain fixed clustering parameter. As shown, the clustering yields two clusters 402 and 406.

[0039] FIG. 5 is a diagram 500 illustrating an example of the clustering result for the POI described above in FIG. 4 at 20K samples (e.g., 1-month dataset) using the same fixed clustering parameter. As shown, the clustering yields only one cluster 502, because the clustering algorithm joined all the points together.

[0040] In some embodiments, in order to address the issue described above regarding the clustering algorithm, a two-stage process to set and adjust the parameters is provided. During the first step, a clustering parameter (e.g., the maximum radius of the neighborhood from a point) may be fixed to an initial value in all experiments, and another parameter (e.g., the desired minimum cluster size) may be set based on the density of the points observed for the POI. If the clustering algorithm fails (i.e., if it does not return any valid clusters), the clustering parameters may be dynamically adjusted (e.g., relaxed gradually) in order to adjust to the observations. In some embodiments, only a few consecutive attempts may be allowed for this adjustment and "no cluster found" will be reported if the clustering fails even after the maximum number of attempts is exhausted.

[0041] In some embodiments, in order to address the issue described above regarding the clustering algorithm, the values of the clustering parameters may be selected automatically. A clustering parameter may need to be re-tuned for different

dataset sizes, which is dependent on the popularity of the POI being clustered and the period of the snapshot of data used. In some embodiments, by using heuristics to approximate point densities, e.g. k -nearest neighbors algorithm (k -NN) distances, suitable clustering parameters for each POI may be decided automatically. This may be implemented in the clustering pipeline. As datasets increase in size for certain POIs, button presses may start to 'map' out the street network. This means that clusters that were once distinct may be 'bridged' by these spurious data points.

[0042] To this end, in some embodiments, for each point in a POI, the average distance to its k nearest neighbours may be measured. Then the clustering parameter may be set to be at some percentile p of this k -NN distance distribution. For example, in some embodiments, $k = 3, 4$, or 5 etc. and $p = 90\%$, 95% , or 99% etc. Setting p higher allows more 'inclusive' (and hence fewer) clusters. In some embodiments, to avoid excessively large clusters on sparser datasets, a maximum cap on a clustering parameter may be set.

[0043] FIG. 6 is a chart 600 illustrating an example of the cumulative distribution function (CDF) of 5-NN distances (in km) for a POI. As shown, around 80% of the points are within 1m of its neighbours. Further, 99% of the points are within 14m of its neighbours.

[0044] FIG. 7 is a diagram 700 illustrating an example of clustering result for the POI described above in FIGS. 4 and 5 using the methods described above to address the issue of traditional clustering algorithm. As shown, five clusters 702, 704, 706, 708, 710 are identified.

[0045] Referring back to FIG. 2, at 208, the apparatus may infer an aggregation or representative point for each cluster as a suitable pickup/drop-off point for that POI. In some embodiments, the aggregation or representative point for a cluster may be the densest point estimated for the cluster. In some embodiments, the inferred points may be stored either for further analysis, human verification, or to directly update the current database of POIs for suitable points.

[0046] Some embodiments may be executed regularly, taking input data in batches (e.g., data accumulated over a week or month), processing them in batches and generating or updating the output optimal service locations. However, executing in this manner inevitably leads to delays in changes. The process requires waiting for a

week/month for data to accumulate before re-processing. In addition, the process requires a computationally intensive and/or long duration to process because the data is processed as a batch.

[0047] Some embodiments may perform ‘online/incremental learning’ (instead of ‘batch learning’), thus solving these problems. Effectively, the learning algorithms may be refactored to update optimal service locations by only receiving a small batch of ‘new data points’, combined with ‘existing memory’ of all the ‘old data’ that the system has learnt from so far. In the extreme case, updates may be calculated with each single new data point that is received.

[0048] This allows the system to react faster and on the fly to changing conditions (e.g. an unforeseen closure in an existing road or relocation of an entrance) in a computationally efficient manner. In certain other cases, this also allows to flag anomalies - for instance, a service location bearing the name of a popular destination but moved significantly due to some manual (inserts/deletes) errors. For example, the clustering algorithm may be modified to allow for incremental changes received on the fly.

[0049] Referring back to FIG. 2, at 210, the apparatus may store and feedback human verification data of such points into the model to assist in refining the clustering algorithm. Some embodiments may have no human input and may be purely ‘mechanically’ executed. However, human expertise is still clearly valuable, especially in highly contextual, temporal or behavioral aspects. Some embodiments may have a team of human experts dedicated to curating service location data, though the methods described above help automate most of the process and are meant to augment the ability and speed of the human team in carrying out their work.

[0050] Normally, such a working dynamic requires the automated system to ‘blindly’ process and output results, by which the human team would evaluate and/or correct if necessary. In some embodiments, improvements to this workflow involve either ‘ranking’ the results for evaluation in some manner (so that the human team prioritizes the most ‘important corrections’), and or feeding back the results in some way to improve the system, usually via the team that creates or maintains the algorithms/systems.

[0051] A non-trivial improvement to this machine-assisted human curation paradigm involves ‘active learning’. Essentially, the system of some embodiments is active in choosing specific clustering tasks it has attempted to solve and ‘asking’ the human experts whether its solution was ‘correct’, with the goal of improving itself on subsequent tasks.

[0052] In some embodiments, an active learning system may perform these steps:

1. Given a baseline model and set of parameters, find the optimal service locations for some of the data. It need not be all existing data available, a subset will do. For brevity, these output optimal points may be referred to as ‘proposed solutions’.
2. Evaluate which proposed solutions should be shown to the human expert for feedback. This can be done by:
 - a. First, measure either of these metrics:
 - Uncertainty of the correctness of a proposed solution, e.g., the algorithm judges that its inference is likely to be erroneous and would need a human-expert to correct it.
 - Informativeness of the potential human-expert answer, e.g., the algorithm judges that the human's answer to this proposed solution has the highest potential to improve the model's overall correctness in as few iterations as possible.
 - The importance of this data point, e.g., it is important to get this location right because it is a high-value POI.
 - A suitable balance of the above.
 - b. Then, show the solutions which rank highly on the metric(s) in (2a).
3. The prioritized solutions in (2b) are shown to human experts (e.g., via a computer dashboard or other suitable human-computer interface). Human experts then provide the ‘correct answer’ as ground truth (e.g., where the optimal service location really is as judged by a human expert).
4. The human experts’ judgments are fed back to the system, and the model parameters are automatically tweaked to reduce its own error on human-provided ground-truth answers.

5. Repeat from Step 1 with the newly tweaked model and parameters, and on either a completely different subset of data, or some of the data used earlier.

[0053] Curated service location quality therefore improves at a faster rate and with fewer human interactions than if examples were shown at random for human feedback. This is because human experts are only required to expend efforts on providing more informative answers where they are needed most.

[0054] Although the disclosure focusses on ride hailing, one of ordinary skill in the art would recognize that the technology described in this disclosure may be generalized to other applications such as collection and delivery of items such as food or merchandise items. In these cases, the same methodology may be used to infer the service locations, i.e., locations of pickup, locations of parking, and locations of item delivery, etc.

[0055] In some embodiments, the optimal transport service locations inferred for a POI may be provided to an automated vehicle (also referred to as a self-driving vehicle). The automated vehicle may provide transport/delivery service at the POI based on the optimal transport service locations to achieve the various benefits described in this disclosure.

[0056] FIG. 8 is a flowchart 800 of a method of transport service. In some embodiments, the method may be performed by an apparatus (e.g., the apparatus 902/902' shown in FIG. 9 or FIG. 10). In some embodiments, operations performed by the method may correspond to operations described above with reference to FIG. 2.

[0057] At 802, the apparatus may receive a plurality of transport service transactions associated with a point of interest entity. In some embodiments, a transport service transaction may include picking up or dropping off one or more persons or one or more items by a service provider at a request of a client.

[0058] At 804, for each transport service transaction of the plurality of transport service transactions, the apparatus may determine a transport service location at which the transport service transaction is executed. In some embodiments, a transport service location may be one of a pickup location, a drop-off location, or a delivery location.

[0059] In some embodiments, the transport service location may be determined based on GPS data and provider application software interaction data. In some embodiments, the transport service location may be determined based on at least one of

an image of the transport service location, a video of the transport service location, telematics data, audio signals, dashboard camera output, client interactions, or chat logs.

[0060] At 806, the apparatus may cluster the transport service locations determined for the plurality of transport service transactions. In some embodiments, to cluster the transport service locations, the apparatus may iteratively relax clustering parameters. In such embodiments, the clustering parameters may include the maximum neighborhood radius and the minimum number of points required to form an effective cluster.

[0061] In some embodiments, to cluster the transport service locations, for each transport service location of the transport service locations, the apparatus may measure the average distance to k nearest neighbors of the transport service location. The apparatus may automatically adjust the maximum neighborhood radius of the clustering to be at a percentile of the k nearest neighbors distance distribution. In some embodiments, k may be 3, 4, or 5 etc. and the percentile may be 90%, 95%, 99% etc.

[0062] At 808, the apparatus may determine one or more candidate transport service locations for the point of interest entity based on the clustering.

[0063] At 810, the apparatus may provide the one or more candidate transport service locations to a client or a service provider associated with a transport service transaction that is to be executed at the point of interest entity.

[0064] In some embodiments, the apparatus may further select a proposed solution to seek human feedback. The apparatus may refine the one or more candidate transport service locations based on the human feedback.

[0065] In some embodiments, the apparatus may receive a second plurality of transport service transactions associated with the point of interest entity. The second plurality of transport service transactions may be the most recent transactions. The apparatus may update the one or more candidate transport service locations based on the second plurality of transport service transactions.

[0066] FIG. 9 is a conceptual data flow diagram 900 illustrating the data flow between different means/components in an exemplary apparatus 902. The apparatus 902 may be a computing device or a system including multiple computing devices.

[0067] The apparatus 902 may include a location determination component 904 that determines transport service locations for a POI based on multimodal data. In one

embodiment, the location determination component 904 may perform the operations described above with reference to 202 in FIG. 2 or 804 in FIG. 8.

[0068] The apparatus 902 may include a clustering component 906 that clusters the transport service locations provided by the location determination component 904. In one embodiment, the clustering component 906 may perform the operations described above with reference to 206 in FIG. 2 or 806 in FIG. 8.

[0069] The apparatus 902 may include a location optimization component 908 that determines candidate transport service locations for the POI based on the clusters obtained by the clustering component 906. In one embodiment, the location optimization component 908 may perform the operations described above with reference to 208 in FIG. 2 or 808 in FIG. 8.

[0070] The apparatus 902 may include additional components that perform each of the blocks of the algorithm in the aforementioned flowcharts of FIGS. 2 and 8. As such, each block in the aforementioned flowcharts of FIGS. 2 and 8 may be performed by a component and the apparatus may include one or more of those components. The components may be one or more hardware components specifically configured to carry out the stated processes/algorithm, implemented by a processor configured to perform the stated processes/algorithm, stored within a computer-readable medium for implementation by a processor, or some combination thereof.

[0071] FIG. 10 is a diagram 1000 illustrating an example of a hardware implementation for an apparatus 902' employing a processing system 1014. In some embodiments, the apparatus 902' may be the apparatus 902 described above with reference to FIG. 9. The apparatus 902' may include one or more computing devices. The processing system 1014 may be implemented with a bus architecture, represented generally by the bus 1024. The bus 1024 may include any number of interconnecting buses and bridges depending on the specific application of the processing system 1014 and the overall design constraints. The bus 1024 links together various circuits including one or more processors and/or hardware components, represented by the processor 1004, the components 904, 906, 908, and the computer-readable medium / memory 1006. The bus 1024 may also link various other circuits such as timing sources, peripherals, voltage regulators, and power management circuits, which are well known in the art, and therefore, will not be described any further.

[0072] The processing system 1014 includes a processor 1004 coupled to a computer-readable medium / memory 1006. The processor 1004 is responsible for general processing, including the execution of software stored on the computer-readable medium / memory 1006. The software, when executed by the processor 1004, causes the processing system 1014 to perform the various functions described *supra* for any particular apparatus. The computer-readable medium / memory 1006 may also be used for storing data that is manipulated by the processor 1004 when executing software. The processing system 1014 further includes at least one of the components 904, 906, 908. The components may be software components running in the processor 1004, resident/stored in the computer readable medium / memory 1006, one or more hardware components coupled to the processor 1004, or some combination thereof.

[0073] In the following, various aspects of this disclosure will be illustrated:

[0074] Example 1 is a method, apparatus, or computer program for transport service. The apparatus may: receive a plurality of transport service transactions associated with a point of interest entity; for each transport service transaction of the plurality of transport service transactions, determine a transport service location at which the transport service transaction is executed; cluster the transport service locations determined for the plurality of transport service transactions; determine one or more candidate transport service locations for the point of interest entity based on the clustering; and provide the one or more candidate transport service locations to a client or a service provider associated with a transport service transaction that is to be executed at the point of interest entity.

[0075] In Example 2, the subject matter of Example 1 may optionally include that a transport service location may be one of a pickup location, a drop-off location, or a delivery location, where a transport service transaction may include picking up or dropping off one or more persons or one or more items by a service provider at a request of a client.

[0076] In Example 3, the subject matter of Example 1 or 2 may optionally include that the transport service location may be determined based on GPS data and provider application software interaction data.

[0077] In Example 4, the subject matter of any one of Examples 1 to 3 may optionally include that the transport service location may be determined based on at

least one of an image of the transport service location, a video of the transport service location, telematics data, audio signals, dashboard camera output, client interactions, or chat logs.

[0078] In Example 5, the subject matter of any one of Examples 1 to 4 may optionally include that the apparatus may: select a proposed solution to seek human feedback; and refine the one or more candidate transport service locations based on the human feedback.

[0079] In Example 6, the subject matter of any one of Examples 1 to 5 may optionally include that, to cluster the transport service locations, the apparatus may dynamically adjust clustering parameters.

[0080] In Example 7, the subject matter of Example 6 may optionally include that the clustering parameters may include the maximum neighborhood radius and the minimum number of points required to form a cluster.

[0081] In Example 8, the subject matter of any one of Examples 1 to 7 may optionally include that, to cluster the transport service locations, the apparatus may: for each transport service location of the transport service locations, approximate the point density distribution of the transport service location; and automatically adjust the maximum neighborhood radius of the clustering based on the approximate point density distribution.

[0082] In Example 9, the subject matter of Example 8 may optionally include that k may be an integer between 2 and 10, and the percentile may be between 90% and 99.9%.

[0083] In Example 10, the subject matter of any one of Examples 1 to 9 may optionally include that the apparatus may: receive a second plurality of transport service transactions associated with the point of interest entity; and update the one or more candidate transport service locations based on the second plurality of transport service transactions.

[0084] A person skilled in the art will appreciate that the terminology used herein is for the purpose of describing various embodiments only and is not intended to be limiting of the present invention. As used herein, the singular forms "a", "an" and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms "comprises" and/or "comprising,"

when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

[0085] It is understood that the specific order or hierarchy of blocks in the processes / flowcharts disclosed is an illustration of exemplary approaches. Based upon design preferences, it is understood that the specific order or hierarchy of blocks in the processes / flowcharts may be rearranged. Further, some blocks may be combined or omitted. The accompanying method claims present elements of the various blocks in a sample order, and are not meant to be limited to the specific order or hierarchy presented.

[0086] The previous description is provided to enable any person skilled in the art to practice the various aspects described herein. Various modifications to these aspects will be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other aspects. Thus, the claims are not intended to be limited to the aspects shown herein, but is to be accorded the full scope consistent with the language claims, wherein reference to an element in the singular is not intended to mean “one and only one” unless specifically so stated, but rather “one or more.” The word “exemplary” is used herein to mean “serving as an example, instance, or illustration.” Any aspect described herein as “exemplary” is not necessarily to be construed as preferred or advantageous over other aspects. Unless specifically stated otherwise, the term “some” refers to one or more. Combinations such as “at least one of A, B, or C,” “one or more of A, B, or C,” “at least one of A, B, and C,” “one or more of A, B, and C,” and “A, B, C, or any combination thereof” include any combination of A, B, and/or C, and may include multiples of A, multiples of B, or multiples of C. Specifically, combinations such as “at least one of A, B, or C,” “one or more of A, B, or C,” “at least one of A, B, and C,” “one or more of A, B, and C,” and “A, B, C, or any combination thereof” may be A only, B only, C only, A and B, A and C, B and C, or A and B and C, where any such combinations may contain one or more member or members of A, B, or C. All structural and functional equivalents to the elements of the various aspects described throughout this disclosure that are known or later come to be known to those of ordinary skill in the art are expressly incorporated herein by reference and are intended to be encompassed by the claims. Moreover, nothing disclosed herein

is intended to be dedicated to the public regardless of whether such disclosure is explicitly recited in the claims. The words “module,” “mechanism,” “element,” “device,” and the like may not be a substitute for the word “means.” As such, no claim element is to be construed as a means plus function unless the element is expressly recited using the phrase “means for.”

CLAIMS

WHAT IS CLAIMED IS:

1. A method of transport service, the method comprising:
receiving a plurality of transport service transactions associated with a point of interest entity;
for each transport service transaction of the plurality of transport service transactions, determining a transport service location at the point of interest entity, the transport service transaction being executed at the transport service location;
clustering the transport service locations at the point of interest entity determined for the plurality of transport service transactions;
determining one or more candidate transport service locations at the point of interest entity based on the clustering; and
providing the one or more candidate transport service locations at the point of interest entity to a client or a service provider associated with a transport service transaction that is to be executed at the point of interest entity.
2. The method of claim 1, wherein a transport service location is one of a pickup location, a drop-off location, or a delivery location, wherein a transport service transaction comprises picking up or dropping off one or more persons or one or more items by a service provider at a request of a client.
3. The method of claim 1 or 2, wherein the transport service location is determined based on GPS data and provider application software interaction data.
4. The method of any one of claims 1-3, wherein the transport service location is determined based on at least one of an image of the transport service location, a video of the transport service location, telematics data, audio signals, dashboard camera output, client interactions, or chat logs.
5. The method of any one of claims 1-4, further comprising:
selecting a proposed solution to seek human feedback; and

refining the one or more candidate transport service locations based on the human feedback.

6. The method of any one of claims 1-5, wherein the clustering of the transport service locations comprises dynamically adjusting clustering parameters.

7. The method of claim 6, wherein the clustering parameters comprise a maximum neighborhood radius and a minimum number of points required to form a cluster.

8. The method of any one of claims 1-7, wherein the clustering of the transport service locations comprises:

for each transport service location of the transport service locations, approximating a point density distribution of the transport service location; and

automatically adjusting a maximum neighborhood radius of the clustering based on the approximate point density distribution.

9. The method of any one of claims 1-8, further comprising:

receiving a second plurality of transport service transactions associated with the point of interest entity; and

updating the one or more candidate transport service locations based on the second plurality of transport service transactions.

10. An apparatus for transport service, the apparatus comprising:

at least one memory; and

at least one processor coupled to the at least one memory and configured to:

receive a plurality of transport service transactions associated with a point of interest entity;

for each transport service transaction of the plurality of transport service transactions, determine a transport service location at the point of interest entity, the transport service transaction being executed at the transport service location;

cluster the transport service locations at the point of interest entity determined for the plurality of transport service transactions;

determine one or more candidate transport service locations at the point of interest entity based on the clustering; and

provide the one or more candidate transport service locations at the point of interest entity to a client or a service provider associated with a transport service transaction that is to be executed at the point of interest entity.

11. The apparatus of claim 10, wherein a transport service location is one of a pickup location, a drop-off location, or a delivery location, wherein a transport service transaction comprises picking up or dropping off one or more persons or one or more items by a service provider at a request of a client.

12. The apparatus of claim 10 or 11, wherein the transport service location is determined based on GPS data and provider application software interaction data.

13. The apparatus of any one of claims 10-12, wherein the transport service location is determined based on at least one of an image of the transport service location, a video of the transport service location, telematics data, audio signals, dashboard camera output, client interactions, or chat logs.

14. The apparatus of any one of claims 10-13, wherein the at least one processor is further configured to:

select a proposed solution to seek human feedback; and

refine the one or more candidate transport service locations based on the human feedback.

15. The apparatus of any one of claims 10-14, wherein, to cluster the transport service locations, the at least one processor is configured to dynamically adjust clustering parameters.

16. The apparatus of claim 15, wherein the clustering parameters comprise a maximum neighborhood radius and a minimum number of points required to form a cluster.

17. The apparatus of any one of claims 10-16, wherein, to cluster the transport service locations, the at least one processor is configured to:

for each transport service location of the transport service locations, approximate a point density distribution of the transport service location; and

automatically adjust a maximum neighborhood radius of the clustering based on the approximate point density distribution.

18. The apparatus of any one of claims 10-17, wherein the at least one processor is further configured to:

receive a second plurality of transport service transactions associated with the point of interest entity; and

update the one or more candidate transport service locations based on the second plurality of transport service transactions.

19. A non-transitory computer-readable medium storing computer executable code, comprising instructions for:

receiving a plurality of transport service transactions associated with a point of interest entity;

for each transport service transaction of the plurality of transport service transactions, determining a transport service location at the point of interest entity, the transport service transaction being executed at the point of interest entity;

clustering the transport service locations at the point of interest entity determined for the plurality of transport service transactions;

determining one or more candidate transport service locations at the point of interest entity based on the clustering; and

providing the one or more candidate transport service locations at the point of interest entity to a client or a service provider associated with a transport service transaction that is to be executed at the point of interest entity.

20. The non-transitory computer-readable medium of claim 19, wherein a transport service location is one of a pickup location, a drop-off location, or a delivery location,

wherein a transport service transaction comprises picking up or dropping off one or more persons or one or more items by a service provider at a request of a client.