

MOHANBIR SAWHNEY, BIRJU SHAH, RYAN YU, EVGENY RUBTSOV, AND
PALLAVI GOODMAN

Uber: Applying Machine Learning to Improve the Customer Pickup Experience

In 2018, Birju Shah, group product manager of maps and sensors, Ryan Yu, senior product manager of pickup experience, and Evgeny Rubtsov, product analyst of maps at Uber Technologies, were working on the best way to measure and improve the quality of the pickup experience for riders and drivers. Ensuring that the pickup experience went flawlessly was a top priority at Uber. Flawed pickups could lead to rider and driver dissatisfaction, reduce driver productivity, and increase the frequency of canceled rides.

Picking up a rider sounded like a simple task. Pulling off a flawless pickup experience was very challenging in practice, however. Finding the precise location of the rider and navigating the driver to the best rendezvous location in an efficient manner was not easy, especially in crowded locations like airports and concert venues. GPS navigation signals could be flawed in urban areas with tall buildings. One-way streets and parking restrictions could also create problems for drivers. Further, drivers and riders across the world had different definitions of what constituted a good experience.

The Uber team members had structured the initiative as three steps. First, they needed to analyze the pickup experience so they could identify potential problems faced by riders, as well as drivers, at each step in the pickup experience. Next, they needed to create a model for pinpointing the best location for picking up a rider. Finally, they needed to develop a quantitative metric for the quality of the pickup experience. The eventual goal of the Uber team was to create an automated

©2020 by the Kellogg School of Management at Northwestern University. This case was prepared by Professor Mohanbir Sawhney, Birju Shah, Ryan Yu, Evgeny Rubtsov, and Pallavi Goodman. Cases are developed solely as the basis for class discussion. Cases are not intended to serve as endorsements, sources of primary data, or illustrations of effective or ineffective management. Some details might have been fictionalized for pedagogical purposes. To order copies or request permission to reproduce materials, call 800-545-7685 (or 617-783-7600 outside the United States or Canada) or e-mail custserv@hbsp.harvard.edu. No part of this publication may be reproduced, stored in a retrieval system, used in a spreadsheet, or transmitted in any form or by any means—electronic, mechanical, photocopying, recording, or otherwise—without the permission of Kellogg Case Publishing.

system that would use machine learning (ML) technologies to improve the pickup experience for all Uber riders and drivers.

History of Uber

Uber was founded in March 2009 in San Francisco by Travis Kalanick and Garrett Camp as a luxury car service through which users could request a ride via an application on their phone. Over the years, Uber transformed into a mobility platform, offering its users various ways to get from place to place, including various ridesharing services from personal cars (UberX, Uber Black for luxury cars) to shared carpool rides (UberPool), as well as electric bike and scooter services. Uber also entered the food delivery market with Uber Eats and was developing shared air transportation with Uber Elevate and investing in autonomous vehicles. Uber services could be used via the company's websites or mobile apps. By 2018, Uber had a worldwide presence and operated in more than 785 metropolitan areas around the globe. That year, the company reported top-line gross bookings of \$50 billion.

The Ridesharing Industry in 2018

Ridesharing had become a popular idea, but it was not a new phenomenon. Evidence of ridesharing could be found as far back as World War II, when acute gasoline shortages led to shared rides. The same idea reappeared during the oil and energy crisis in the 1970s. The arrival of high-end technology, GPS navigation devices, and smartphones gave birth to the modern ridesharing market, in which drivers who owned mostly cars teamed up with companies such as Lyft and Uber to provide rides to potential customers through dispatching platforms.

The ridesharing market had seen explosive growth around the world, fueled by aggressive expansion of global players like Uber, as well as local competitors like Didi Chuxing in China, Ola in India, Yandex in Russia, and Grab in Southeast Asia. In 2017, the size of the worldwide ridesharing market was estimated to be \$36 billion and was projected to grow to \$285 billion by 2030. In the United States, ridesharing companies transported 2.61 billion passengers in 2017, a 37 % increase from 1.90 billion in 2016.* Ridesharing trips were highly concentrated in large cities. The nine largest cities (Boston, Chicago, Los Angeles, Miami, New York, Philadelphia, San Francisco, Seattle, and Washington, DC) accounted for 70% of all ridesharing trips, totaling 5.7 billion miles. In December 2018, Uber's share of rideshare spending was 68.8%, whereas Lyft accounted for 28.8%. Ridesharing was rapidly becoming an integral part of everyday life in urban areas, particularly for younger riders who preferred ridesharing to owning a car.

Overview of Pickups

The quality of pickups was a crucial factor for the success of ridesharing companies. To keep riders and drivers happy, Uber needed to ensure that riders were picked up on time at the right

* "The New Automobility: Lyft, Uber and the Future of American Cities," Schaller Consulting, July 25, 2018. <http://www.schallerconsult.com/rideservices/automobility.pdf>.

place and that drivers spent their time efficiently and productively. The quality of pickups was also important to the cities in which Uber operated. Uber needed to collaborate with city and state administrators in the management of road construction, road closures, and traffic. By working with cities, Uber could ensure efficient use of roads and compliance with traffic and parking laws. As a start, Uber partnered with cities to provide anonymized data from more than two billion trips to help urban planning around the world through Uber Movement. Access to such data could inform decisions on adapting existing infrastructure, reducing congestion, and investing in efficient transportation in the future.

Challenges in Pickups

The quality of a pickup experience could be affected by several factors. To illustrate, consider the examples below:

Cultural Differences

In the United States, riders preferred to interact with drivers through the Uber application or through text messaging. Riders would only call drivers if a problem occurred with the pickup, such as difficulty finding the right rendezvous location or a delay in a driver's arrival. The contact rate for rides in the US was less than 20%. In contrast, riders in India liked to call drivers to give them precise directions or to make sure the driver was on the way. The contact rate between riders and drivers was more than 80% in India. Such cultural differences needed to be considered in designing a stress-free pickup experience. For instance, using the contact rate between drivers and riders as a signal of friction in the pickup experience would be appropriate for the US market but not for the Indian market.

Modalities

Uber offered various ridesharing modalities (UberX, UberPool, and Uber Black, as well as scooters, bikes, etc.). Riders could choose among all these options, so Uber needed to make sure that riders were matched to the appropriate type of Uber service. For modalities like UberPool, which involved sharing a car with other riders, Uber would need to choose routes that made sense for all the riders in a car, as well as for the driver.

Venue Pickups

In many airports and event venues, Uber was allowed to pick up passengers only in designated areas. These rules needed to be communicated to both riders and drivers in an intuitive way when they opened their app. Being able to operate in such venues was critical to Uber's business, as venue pickups accounted for a significant percentage of gross bookings and annual trips.

Safety

Safety was a top priority for riders and drivers. To establish trust with its users, Uber needed to build safety into the trip experience and to be prepared to help when incidents arose. All riders, drivers, and cars were protected by insurance maintained by Uber, which encouraged riders and drivers to participate in creating safe environments. For instance, in Uber's safety tips, drivers were advised to check for cyclists and pedestrians on the road before pulling over for a pickup and to encourage their riders to sit in the back seat so that they could safely exit either side of the vehicle when they reached their destination.

Compliance

Cities were an important part of Uber's ecosystem, and the company wanted to help foster a compliant and safe environment. For example, compliance in the New York City area could be confusing. A New Jersey driver would be able to pick up a passenger in New Jersey and drop him or her off in New York City. However, without an NYC Taxi & Limousine Commission (TLC) license, the driver would not be able to pick up another rider in New York. Uber worked with drivers in NYC to help them obtain TLC licenses and provided assistance and guidance throughout the process, which included a drug test, a background check, and required classes.

Rider Expectations

Uber catered to a diverse set of users, so riders' expectations could vary significantly based on the rider persona and the context of the trip. For instance, a commuter who called an Uber to work every day might expect a consistent and reliable pickup from the same spot near her home each morning. In contrast, a suburban rider who used Uber occasionally might be more concerned about the choice and availability of drivers in remote suburbs. Drivers also wanted different things from a pickup. Some drivers might expect riders to be waiting when they arrived, whereas other drivers might not like the inconvenience of shared rides.

These examples are just a small illustration of the complexity that Uber needed to consider in understanding and optimizing pickups across its vast global network. The challenge for Uber was to deliver a quality pickup experience that could scale with the company's growth.

History of Pickups with Uber

2014–2016: Pickup-First Paradigm

From 2014 to late 2016, Uber's mobile application centered the rider experience around the pickup experience, positioning pickups as the first step of the flow to requesting a trip (see **Exhibit 1** for an illustration of the pickup-first paradigm). The onus of setting a pickup was almost entirely on the rider. Riders needed to take the key decision to set the location of the pickup. All other actions downstream (e.g., selecting the modality, setting the destination) depended on this initial action. If riders made a poor choice of pickup location, the Uber app did very little to

deal with downstream problems with the pickup. Essentially, the app merely matched riders and drivers, without offering assistance during pickup.

2016–Present: Destination-First Paradigm

In late 2016, Uber launched a new rider application redesigned from the ground up. In the redesigned app, riders were first asked where they wanted to go. They then selected a product and set their pickup. The app marked a radical shift in the request flow, flipping the paradigm from pickup-first to destination-first (see **Exhibit 2** for an illustration of the destination-first paradigm). The destination-first paradigm had an unintended consequence. The action of setting the pickup became a secondary action, so riders treated the pickup location as an afterthought. Uber found it more difficult to focus the riders' attention on the pickup location, which could increase the possibility of pickup problems.

Mapping the Pickup Experience

Shortly after the rider app redesign in 2016, Uber developed an end-to-end customer journey map for the customer pickup experience that consisted of six steps (see **Exhibit 3** for a visual illustration). These steps were defined as follows:

1. **Contextualize:** When a rider opened the application, Uber tried to deduce his or her **anchor location**, a best guess at where the rider was. The anchor location was determined either using the rider's GPS data and/or search data entered for their pickup.
2. **Assist:** Uber next tried to determine the best rendezvous location (i.e., pickup location) for the anchor location. For instance, the anchor location in Chicago's O'Hare International Airport might be the arrivals level in Terminal 1, but the rendezvous location for this anchor location would be the rideshare pickup area between Terminal 1 and Terminal 2 at the departures level. To find a good pickup spot, Uber used historical trips and outcomes to suggest a rendezvous location to the rider.
3. **Depict:** In this phase, Uber presented riders with suggestions for improving their pickup location. Uber could present suggested rendezvous location candidates to the rider for selection and confirmation. Uber could also show riders a screen to refine their pickup location if it needed more information about where they were. Then, with the user's input, Uber dispatched and passed along the rendezvous location to a driver.
4. **Navigate:** Uber navigated the driver to the rendezvous point. The routing algorithm guided the driver to the end-of-route (EOR) location, which was the closest point on a road segment to the confirmed rendezvous location. The Uber driver app provided a route-line, or suggested navigation path, that the driver could choose to follow. Drivers could also choose to use their own navigation provider to get to the EOR location.
5. **Track:** After dispatch, Uber enabled the rider to track the car and take action appropriately. The rider might receive push notifications about the car arriving or the driver waiting. In addition, riders could monitor and manage their pickup in several ways: They could see the

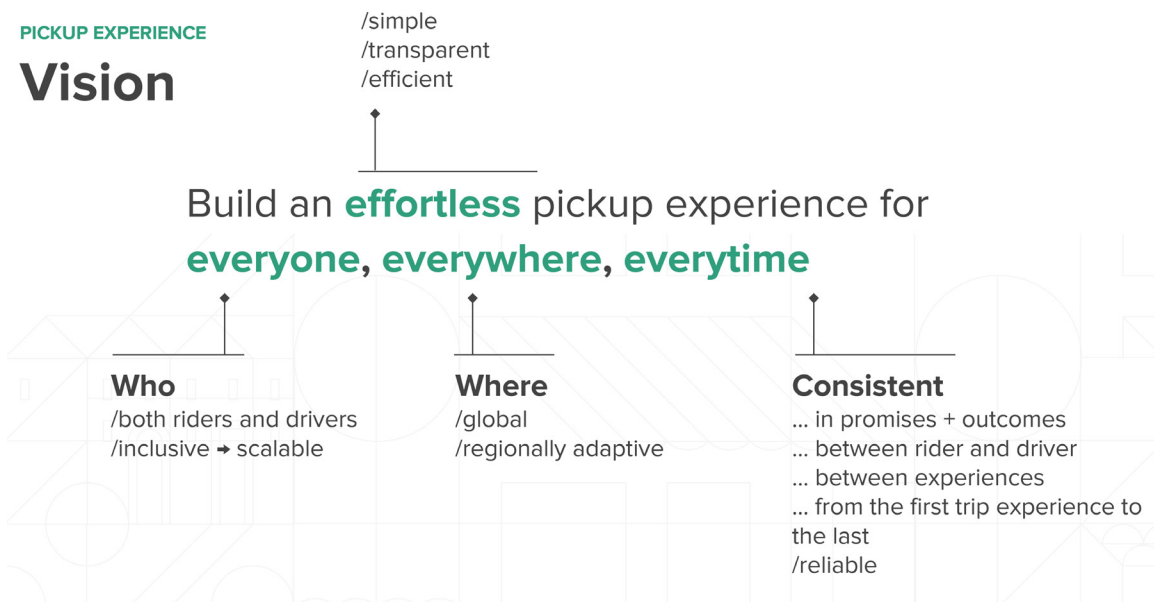
real-time GPS location of the car; get an estimate of the time that the driver would take to arrive; contact the driver; or change their pickup location post-dispatch.

6. **Meet:** In the final moments of the pickup, Uber did all it could to make sure the rider and driver could find each other. Most of this involved metadata the rider could see, such as rider/driver names and the car model and license plate number. In addition, Uber provided riders with tools to facilitate the finding and meeting process, including the ability to share their real-time GPS location and use their phone as a color spotlight to help their driver see them (see **Exhibit 4** for details of the “meet toolkit” for riders).

Analyzing Pain Points in the Pickup Experience

Uber’s vision was to build an effortless pickup experience for everyone, everywhere, every time. *Effortless* implied a simple, transparent, and efficient pickup. *Everyone* included riders, drivers, and city residents; Uber wanted to improve the experience and efficiency of traveling. *Everywhere* meant all markets where Uber operated around the world. *Every time* meant that Uber aimed to deliver a consistent experience riders and drivers could rely on. These goals were represented in the vision statement image below (**Figure 1**):

Figure 1: Uber’s Pickup Experience Vision



Source: Uber.

The Uber team created a wish list for a perfect pickup experience:

1. The rider requested a car at the *best pickup location possible*.
2. The best available driver was assigned to the rider.

3. Driver took the most effective route to the pickup location.
4. Rider knew exactly how to get to the pickup location.
5. The rider and driver arrived at the pickup location at the exact same time.
6. The rider got into the car as soon as the driver stopped, and the car was on its way immediately.

Pickup experiences in practice often were less than ideal. With the ideal pickup experience in view, the Uber team next started to analyze the potential problems that riders and drivers could experience during the pickup experience. These pain points could be something as simple as an Uber driver who stopped farther down the street from where a rider actually was because the rider's phone GPS placed the rider in a location significantly off from where they actually were. Drivers could reach the rendezvous point only to find that the rider was on the wrong side of a one-way street. Drivers could reach a pickup location on a busy city street to find that the rider had not yet arrived at the location, forcing the driver to abandon the pickup. Drivers could be stymied by the lack of parking at a pickup location, gated communities with restricted access, and temporary parking restrictions at crowded concert or sporting event venues. Riders could have trouble finding drivers at busy pickup locations like airports. The Uber team created a narrative that illustrated potential pain points at each phase of the pickup experience (see **Exhibit 5**).

Although this illustration of general pain points was useful, the Uber team realized that the pain points, as well as customer expectations, could vary substantially based on the customer persona. To get a better understanding of pain points for specific customer contexts, the team created a set of personas to represent different customer contexts:

- **Premium:** Riders who wanted the best quality of transportation that Uber could provide. They were typically professionals who were traveling for a business purpose and were able to expense their trip to their companies. Premium customers opted for Uber Select (premium rides in higher-end cars with leather seats) or Uber Black (luxury rides for six in high-end cars with professional drivers). They expected curb-to-curb service and a personalized experience, and they were willing to pay a hefty premium for a high-quality ride. Premium riders enjoyed pickups that included a little extra time to get to the car and features tailored to suit their needs, such as temperature control and luggage service.
- **High-Frequency Riders (HFR):** Riders who used Uber very frequently for short rides in a densely populated urban area like Manhattan or downtown Chicago. These riders could use Uber several times a day, and their rides were typically a few miles or less. They expected quick curbside pickups on busy streets, a challenging task for drivers dealing with dense traffic, one-way streets, and parking restrictions. HFRs qualified for notable rewards, such as discounts on future rides, better customer support, and priority airport pickups. For such frequent riders, Uber introduced Ride Pass, designed to let riders use Uber anywhere in a city or at any time for a low price. Ride Pass began with an introductory rate of \$14.99 monthly.
- **Commuters:** Riders who used Uber to get to and back from work on a regular basis. They had the same origin and the same destination for most trips. They expected a consistent pickup and dropoff experience and wanted to make sure that they reached work on

time without having to wait a long time for the Uber driver. For these commuters, Uber launched Uber Express Pool, which cost as much as 50% less than the standard shared UberPool ride. Users could select an area to wait at rather than a pickup point, then wait a maximum of two minutes in which Uber matched them with their rideshare buddies.

- **Social:** Riders who used Uber for social activities like going to restaurants, concerts, sporting events, or bars. These riders tended to use Uber during late evening hours and on weekends. They often traveled in groups, and they were sometimes inebriated. They expected drivers to pick them up at late hours and at busy venues, and they were often frustrated with “surge pricing” on busy weekend nights.
- **Travelers:** Riders who used Uber at an airport or train station as part of a long-distance travel experience. These riders included business travelers, families traveling on vacation, and international tourists visiting the United States. They expected timely and stress-free pickups at airports, and they were frequently frustrated with finding the right pickup location and the right Uber driver at congested rideshare pickup locations at busy airports. These riders also tended to have a lot of luggage, and they appreciated the driver’s assistance in loading and unloading it.

The team wondered how to map the six stages in the pickup experience to the pain points for each segment. They would need to create a persona-specific list of pain points for the different phases of the pickup experience. They would also need to define the ideal pickup experience for each persona, as the outcome expectations varied dramatically across personas. The pain points and outcome expectations would be based on the persona profiles, customer interviews, and persona experience. This map of pain points and outcome expectations would help the Uber team to understand common themes in improving the pickup experience that could scale across customer segments and pickup locations.

Improving the Pickup Experience with Automation

With an understanding of the pain points in the pickup experience, the Uber team turned to the solution approaches to the pain points. They embarked on an initiative to improve the pickup experience using advances in artificial intelligence (AI) and machine learning (ML) technologies. Before an ML-based model could be built, the team needed to create a model to predict the location of the rider and they needed to define a quantitative metric for the quality of a pickup experience.

Predicting the Location of the Rider

Uber used a sophisticated process to identify and refine the pickup location of the rider:

1. **Sense Signals:** Uber collected relevant data about the rider’s anchor location, such as the GPS coordinates of the anchor location and manual inputs provided by the rider.
2. **Score the Anchor Location:** Uber scored signals on the rider’s location by assigning either a low or a high confidence to the anchor location. If the rider had defined the anchor location

through manual text input or by manually locating a pin, Uber assigned a high confidence score to the anchor location. If the rider had not entered anything, Uber used GPS data as the default source of the anchor location. If the anchor source was GPS data, the anchor confidence could be high or low, depending on the context of the pickup. If the confidence was high, Uber would move on in the pickup process. If the confidence was low, Uber would try to leverage other data to improve the confidence by using hypotheses related to the rider's profile and behavior and the characteristics of the pickup location. These hypotheses would be used to create a confidence score for the rider's location using a predictive model.

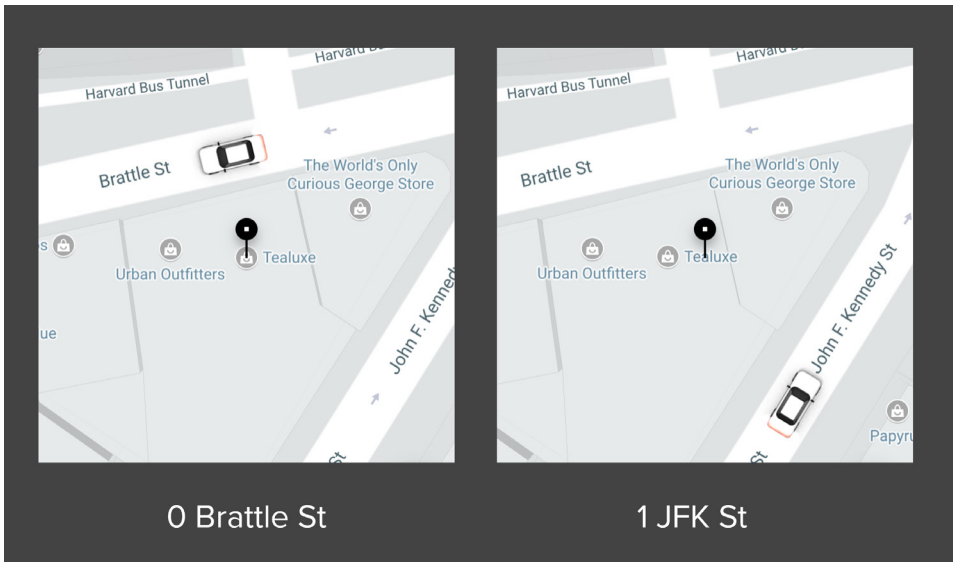
3. **Decide Pickup Location:** Once Uber had a score for the anchor location, it would determine what to show to the riders on the app to assist them with their pickup. For low anchor confidence pickup requests, Uber showed no suggested hot spots for pickups. Instead, it showed the rider their predicted location pin and a search bar to prompt the rider to give additional information on their pickup location. For high anchor confidence pickup requests, Uber showed one or more hot-spot suggestions nearby and let the rider choose one. These hot-spot suggestions were based on historical data accumulated from rides on the Uber platform or, for certain pickup venues like airports, curated hot spots for designated pickup areas. For instance, if the rider was at an airport terminal, Uber might have suggested Door 2 on the departures level, where rideshare pickups were allowed.
4. **Adapt Pickup Location:** Uber used the rider's input (either updated pickup location or selected hot-spot pickup) to move on with the pickup process or loop again with the new signal.
5. **Refine Pickup Heuristics:** Uber continued to evaluate the data from pickup experience to determine ways to optimize it. For instance, if a hot spot was rarely chosen for a specific location, often resulted in rides that were rated lower, or had led to more cancellations, then Uber might have stopped suggesting it. If Uber saw that riders always adjusted their pin to a certain spot or never adjusted it and were happy with their pickup, then Uber assumed that pickup location with high confidence going forward.

To improve the confidence score of the anchor locations when GPS data was the only signal available, the Uber team needed to come up with a model to predict the likely location of the rider. GPS coordinates were an imperfect signal for determining the best rendezvous location (i.e., where the rider should be picked up). First, the GPS location could be erroneous. The GPS system worked well in clear areas under open skies, achieving accuracy to within a 5-meter (16-foot) radius. However, its positioning accuracy could be degraded by tall buildings or walls, which could reflect the GPS signals. In densely populated and highly built-up urban areas, GPS location estimates could have a margin of error of 50 meters (164 feet) or more, which could place the driver on the wrong street. GPS coordinates often could not pinpoint at which entrance of a large building the rider was located. Buildings like train stations had several entrances and exits within the same city block, which opened on different streets. If drivers arrived at the wrong entrance, they could be forced to loop around one-way streets to get to the correct entrance. Drivers could also arrive at the pickup location only to find that the rider was at the other end of a divided street. In busy venues like airports and sporting arenas, drivers and riders could find it difficult to find

each other in a mass of drivers and riders (see **Figure 2** for illustrations of the challenges in using GPS signals to predict the rider's location):

Figure 2: Examples of Challenges in Using GPS Signal to Predict a Rider's Location

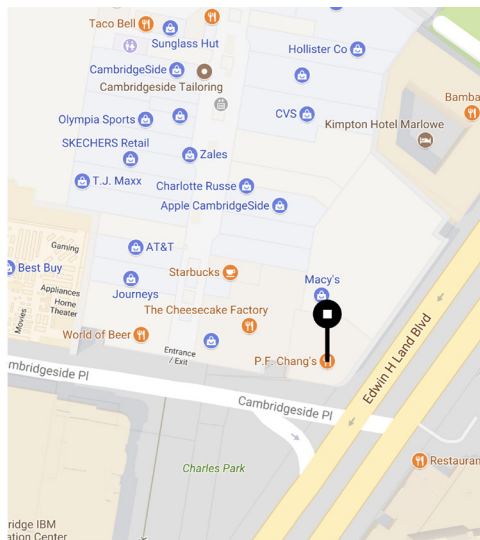
Multiple Interpretations of GPS: Wrong Street



Multiple Interpretations of GPS: Ambiguous Location

Which address is right?

- 📍 P.F. Chang's
- 📍 CambridgeSide Galleria
- 📍 100 Cambridgeside PI
- 📍 100 Cambridgeside PI Ste C101



Source: Uber.

Cópia autorizada para uso exclusivo no International Week AESE-SORBONNE 2026 (2026-04-23)

The Uber team needed to leverage historical and telemetric data to predict the rider's location. They could use the rider's historical behavior, characteristics of the location, time of day, and date of week, as well as the rider's profile, to build a predictive model for the rider's location. The team needed to come up with a set of hypotheses that would help them to build such a model. For instance, Uber could use the rider's trip history to come up with the following hypothesis:

Hypothesis: *If a rider was recently dropped off at a location (e.g., restaurant, event venue), it was likely that he/she was still there if they were requesting a ride nearby an hour later.*

These hypotheses would need to be categorized into hypotheses related to the rider's profile, the pickup location's characteristics, and the rider's intended destination. This initial set of hypotheses could be tested to see how useful they were in the scoring model for the rider location.

Defining a Pickup Quality Metric

The next step in automating the pickup experience was to define a quantitative metric for the quality of a pickup. The pickup quality metric would help Uber to define an "ideal" pickup that could serve as an "objective function" to be optimized by a machine learning model.

Defining a pickup quality metric was a challenging task. The quality of a pickup experience could be measured in many ways. Did the driver loop around multiple times unnecessarily while the rider was already at the curb? Did the driver stop at the wrong location? Did the rider walk to the wrong location? Did the rider have to do more walking than necessary to get to the pickup location? How long did the rider or driver wait at the pickup location? Did the rider have to contact the driver during the pickup experience? Did the rider or driver cancel the ride? Did the rider rate the ride poorly? These were all things someone at Uber could potentially deduce in isolation, but how would Uber ladder these up into a unified metric that could quantitatively represent that information?

Although the quality of a pickup would be ultimately subjective to the rider, Uber wanted a metric that could quantitatively represent the pickup quality. Uber had access to three types of signals that could be used to create a pickup quality metric. It could use *active signals* collected from customers, such as customer complaints, user research, and feedback cards. It also had access to *passive signals* such as GPS data and user activity data to identify when a rider or driver was confused or uncertain. In addition, it could leverage *third-party signals*, such as traffic patterns, congestion, and parking restrictions, to predict potential problems with pickups. The pickup quality metric would ideally combine sensor data, user activity data, and business metrics like contact and cancellation rates. Uber had been using cancellation rates and pickup location error (distance between in-app rendezvous point and actual point where the ride started) as its primary success metrics for a pickup experience. However, these were *proxy metrics* that suggested friction in the pickup experience. Further, these metrics were limited in assessing quality, as they did not capture good experiences.

The Uber team needed to determine what signals they should use and the scoring process they could employ to determine the quality of the pickup experience. A good pickup quality metric would have the following attributes:

- Represented the overall quality of a pickup (not just a single dimension of quality)
- Could be universally applied to every trip across markets and modalities
- Could distinguish between “good” and “great” pickups (i.e., does more than just flag the bad ones)
- Could be usefully visualized geographically
- Could be usefully compared across cities/countries
- Could be usefully compared across similar segments (i.e., pickups at JFK airport versus SFO airport)

The Uber team had put together a potential list of attributes (also called features) that could be used to create a pickup quality metric (see **Exhibit 6**). The team would need to select a subset of three to four attributes that would best capture the quality of a pickup experience. They would then need to assign weights to these attributes so that they could develop a single “pickup quality score.” For instance, the pickup quality metric could be defined as follows:

$$\text{Pickup happiness} = \text{Minimize} [(50\% * \# \text{ of steps to rendezvous}) + (25\% * \# \text{ of ETA Changes}) + (25\% * \text{time spent in search})]$$

The team wondered which of the potential attributes best reflected the overall pickup quality. They were also unsure how to assign initial weights to the selected attributes. The team understood that the pickup quality metric could be adapted and refined based on experience, but it was important for the team to start with a good guess for a robust pickup metric.

Automating Pickups at Scale

To improve pickup satisfaction on a consistent basis for millions of tips, the Uber team realized it needed to develop an automated model for predicting and improving the quality of the pickup experience. Automation of the pickup experience required the creation of a machine learning model for the quality of the pickup experience. Building an ML model involved several steps and required three sets of capabilities: business domain expertise, data science expertise, and user experience (UX) expertise. Developing an ML model took place over seven steps:

1. Define the business problem: The Uber team needed to begin by asking what the problem was and why it was important. How would the solution to the problem benefit customers and the company? How much business value could be generated by solving the problem? In this case, the problem to be solved was to improve the quality of the pickup experience to increase rider and driver satisfaction.
2. Assess if ML is appropriate: Not all business problems lent themselves to an ML solution. Viable ML models needed massive volumes of training data from different sources. The data

needed to be timely, accurate, and easily accessible. The cost of collecting the data had to be reasonable. The time taken to collect the data and make decisions needed to be acceptable, as some decisions had to be made in real time (for instance, letting a rider know during the trip that the pickup location needed to be changed based on a predictive model for estimating the time of arrival). The Uber team had to evaluate the volume, variety, cost, accuracy, and timeliness of the data needed to build the ML model.

3. Gather and label the data: The next step involved asking what sources Uber would use to collect the data. Could it gather the data using feedback from users and riders? Could it collect relevant data from third-party sources that provided information on traffic, weather, and road conditions? What were the trade-offs in collecting pickup quality data directly from riders versus inferring pickup quality from proxy variables?
4. Preprocess the data: The first three steps in developing the ML model were driven primarily by business analysts who knew the business domain and the business problem. Once the problem and the data sources had been defined, the data science team led the development of the actual model. The first step in this phase was to process the raw data into useful variables for the ML model. Raw data could be incomplete and could contain errors. Raw data needed to be formatted to fit the requirements of the ML model. How would Uber improve the data quality to correct erroneous or missing data? How would Uber define the features for the ML model and define the objective function once it had “cleaned” the data?
5. Choose the ML approach and algorithm: The data science team next needed to select an ML approach and algorithm for the model. ML approaches included two types of learning: supervised learning and unsupervised learning. Supervised learning began with a ground truth, which meant that Uber knew what the “right” or “good” output values were for its model. Unsupervised learning had no labeled outputs, so the ML model had to infer the patterns and structure present within the data. Which approach should Uber adopt for the pickup experience model? After an ML approach was selected, the Uber team would need to decide on the best algorithms. Uber could use many different algorithms, such as logistic regression, support vector machines (SVMs), and neural networks. They would need to consider the pros and cons of the different algorithms and decide if they should use a combination of algorithms (a technique called “ensemble methods”) to improve the predictive performance of the model.
6. Engineer features: ML models used a set of variables called “features” to predict the value of a dependent variable called the “objective function.” Building a good ML model required good features. Feature engineering was the process of creating new features, discarding those that were irrelevant, and selecting the set of features that will maximize the accuracy of the model. Feature engineering was different from data cleansing because it involved manipulating the variables to create a more effective set of variables. To conduct feature engineering, how could the Uber team test the predictor variables and ensure that they contributed to pickup quality in a meaningful way? How would they decide possible features of the model and how those features would work in the model?
7. Make UX decisions: The final step of creating an ML model involved making decisions and designing the user experience. This included triggering UX messages to customers, employees,

and partners. What messages should Uber trigger to riders and drivers based on the model's predictions? When and how should these messages be presented to riders and drivers?

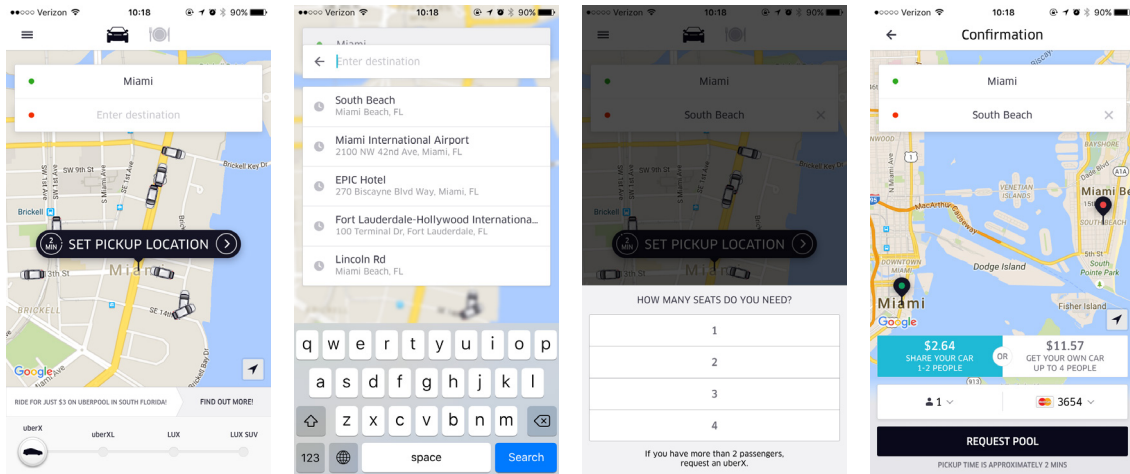
The Uber team would take the lead on the first three steps. The data scientists would direct the next three steps, and the UX design team would implement the final step.

For a summary of these steps to building an ML model, see **Exhibit 7**.

Conclusion

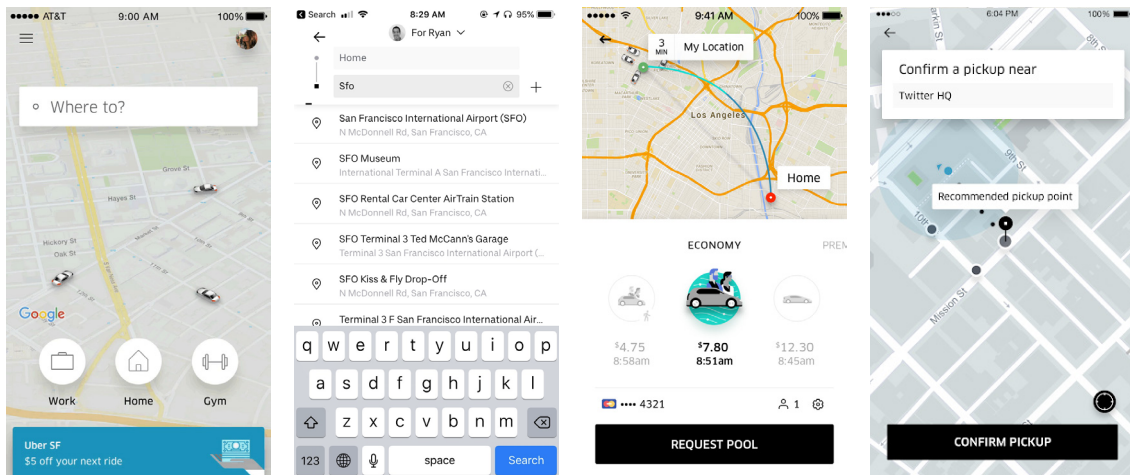
Uber viewed improving the pickup experience as a top priority. All other user interactions flowed from this first step, so it was essential to get the pickup experience right. Uber would continue to work on automating real-time actions based on signals and finding the best way to represent the pickup quality as a useful metric. But did the Uber team approach the pickup context the right way? Could the team members use simple heuristics to measure and improve the pickup experience, or should they invest in machine learning to achieve results at scale? Would there be a world where pickup recommendations truly could be automated, considering all the subjectivity and segmentation involved, or would it always have human involvement and curation? The answers to these questions would go a long way in improving the retention and loyalty of riders and drivers.

Exhibit I: 2014–2016: Pickup-First Paradigm



Source: Uber


Exhibit 2: 2016–Present: Destination-First Paradigm



Source: Uber.

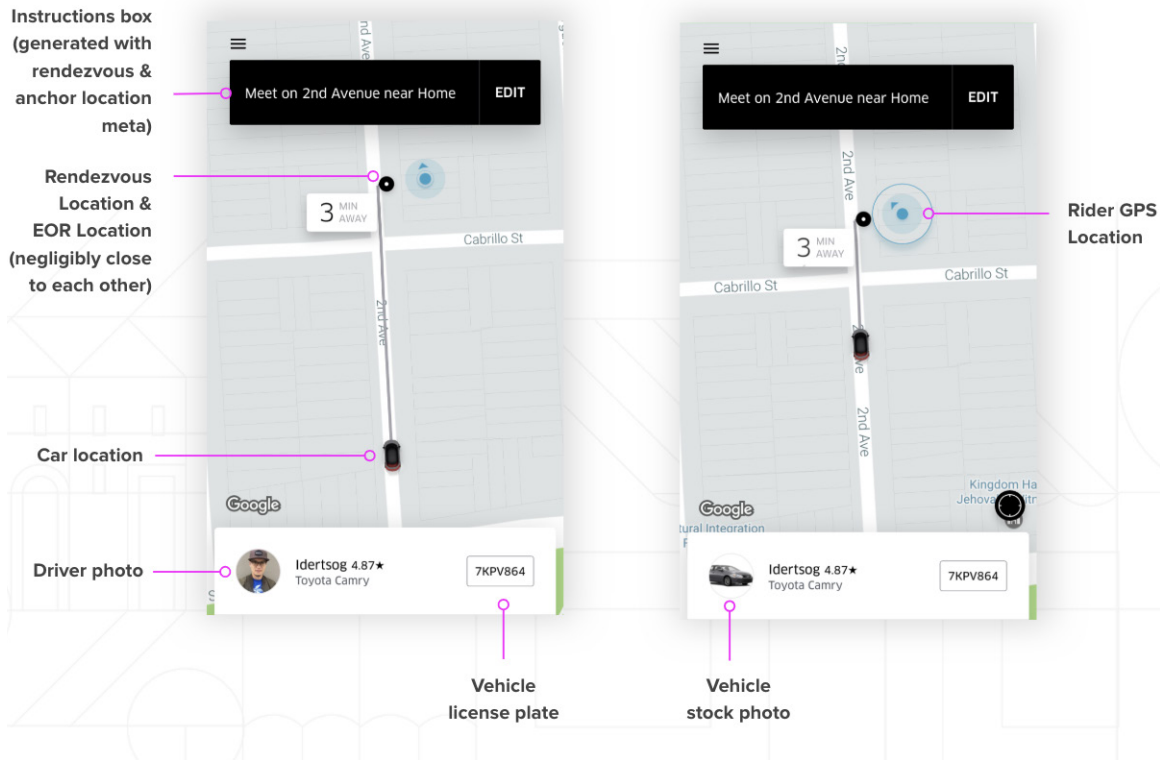
Cópia autorizada para uso exclusivo no International Week AESE-SORBONNE 2026 (2026-04-23)

Exhibit 3: Steps in the Customer Journey for a Pickup Experience



Source: Uber.

Exhibit 4: Meet Toolkit



Source: Uber.

Cópia autorizada para uso exclusivo no International Week AESE-SORBONNE 2026 (2026-04-23)

Exhibit 5: Illustrative Pain Points during a Pickup Experience

Contextualize

Rider

You've had a long day at work, and your feet are tired from wearing heels, so you call an Uber to take you home, rushing through the request process. You see the car pull up, but the driver goes past you and stops fifty feet down the road, where the pickup location seemed to be on their app. As you walk the fifty steps it takes to meet your car, you feel the pain in your feet more acutely. In this situation, it turns out that your GPS contextualized your pickup location slightly off, deducing an anchor location with low confidence that created a rendezvous spot halfway down the block from your actual location.

Driver

You have arrived at the location and don't see anyone nearby looking for an Uber. You double-check the address, and sure enough, the building you're directly in front of has the same address. You send a message and wait a minute before you finally see your rider walking slowly to the vehicle. You're annoyed that you've just lost two minutes waiting at the right spot for someone who should have been there already. Nevertheless, you politely greet the rider because this happens often, and you don't want to get a low rating.

Assist

Rider

You're running late to work, so you call an Uber instead of taking your usual subway ride. You see your GPS dot shows you at home in your backyard, where you have just fed your dog, so you confirm your pickup. As you track your car, you realize that it has been routed to the parallel street on the opposite side of your block. You quickly message your driver to clarify where you are and luckily make it to work on time for your meeting. In this scenario, Uber used your GPS location and anchored you to the other street, to which you happened to be closer to from your backyard but to which you had no access. Because you didn't see the pickup hot-spot candidates on the confirmation screen that you usually see when you're calling an Uber in the city, you assumed the pickup would be right.

Driver

As you drive around to the other side of the block, you wait at a stop sign longer than you feel you should to make a left turn. As you approach the right street, the traffic light turns red and you wait again to make another left. When you finally arrive, you're more annoyed than usual at the two minutes you've wasted to drive around the block.

Depict

Rider

You have just finished shopping for groceries and are standing outside the entrance holding heavy bags of laundry detergent. Because your arms always get sore when you walk thirty minutes home carrying groceries, you decide to call an Uber, entering the address for the pickup location. Your driver sees the address and navigates there, waiting for you on the street, and calls asking where you are. You tell your driver where you are and are thankful when they happily drive into the parking lot, up to the entrance to pick you up. In this scenario, your pickup location was correctly passed along to the driver but depicted you on the street rather than the storefront.

Driver

You don't mind losing a few seconds to drive up to the grocery store entrance to pick up the rider, but you dislike making phone calls on the road. To avoid using your phone while driving, you pull over closer to the sidewalk and hope it doesn't bother nearby drivers on the road too much. You feel bad when you see that cars still need to merge into the next lane to drive around you. When your rider picks up the call, you're relieved to know where to go so you don't need to stall longer on the road.

Navigate

Rider

You have just had your first doctor appointment in the new city you're living in and call an Uber to head home. As your car approaches, you realize that it's on the opposite side of the street. You see the car waiting for you, and you step forward to cross the street but then stop after noticing all the traffic. You eventually walk down to the corner of the block to cross, and the driver pulls forward to meet you. You apologize to the driver for the wait as you get into the car.

Driver

You get matched to a new rider and quickly open Google Maps to navigate to the pickup spot. When you arrive, you notice that your rider is on the opposite side of the street. You look around to see if you can do a U-turn to meet them, but the street has a double-yellow line and too many cars coming. The rider motions that they'll walk down the street, and you nod and drive forward to meet them. In this scenario, the end-of-route pickup location was matched to the closest road segment, so when you used Google Maps to navigate there, you were routed the quickest way to get to that point on the street, not necessarily to the right side of the street. It would have been faster if you had been routed in the opposite direction than have the rider walk down the block.

Track

Rider

After several hours of interesting talks, you leave the conference center with some new acquaintances in your field and are all looking forward to the happy hour. Because you flew in from out of town and don't know the city well, you decide to call an Uber. You get matched to a car a couple blocks away, and the ETA is four minutes. As a few minutes pass, you notice the car icon has barely moved and is alternating directions on your map. You cancel and call another Uber. In this scenario, the area's congestion may have weakened your driver's GPS location data as it was routed to your app for you to track, or the ETA might not have accounted for the added congestion due to the conference.

Driver

You have lost several minutes navigating to the pickup location for a ride that was later canceled and are annoyed with the app and rider for wasting your time without additional compensation. You soon get another match nearby but worry that with the current traffic conditions, the rider might also cancel. You resolve that if you get another cancellation, you're going to stop driving for the day and get something to eat instead.

Meet

Rider

The last act of a music festival has just ended, and you head out of the venue to call an Uber home. The app lets you know where you need to go to get picked up, and you make your way over. You see your car has arrived and begin getting notifications from your driver. You try to find the car, but it's a black car like most of the other cars waiting in the area, and it's impossible to scan all the license plate numbers as hundreds of people are walking across the parking lot, as well. You see many others on their phone talking to their driver or walking down the line of cars trying to look at license plates. It takes a few minutes, but you finally find the right car and get home safely awhile later. In this scenario, the designated rideshare pickup area was so congested that it was difficult to find and meet your driver.

Driver

As you pull into the venue, you begin to feel anxious seeing all the other Ubers and people walking around. You get to the designated pickup spot and wait for the rider. As you look around the sea of people, you realize you don't have much information to find your rider except their name and a low-quality picture. You look over at the other cars nearby and share some sympathetic looks with other drivers also waiting for their passengers. Someone opens your car door, thinking you're their driver, but you check the name and let them know they've made a mistake. It takes some time for you to meet the rider you have been matched with, and although you're a little peeved it took so long, you can also empathize because of the congestion at the venue. You wonder if the surge pricing on this ride is worth the additional time and anxiety.

Source: Uber.

Exhibit 6: Potential Features to Define a Pickup Quality Metric

Environment Related

- Congestion along a potential hot spot or pickup (real-time pruning)
- Event-based pruning around a hot spot (e.g., detecting lane closure, construction, sidewalk construction)

Behavioral

- Sensor Inferences
- User Generated Feedback Cards
- Support Tickets opened
- UX Research

Passive (Sensors)

- Location scoring during uncertainty
- Driver loops around rendezvous
- Driver ETA jumps within last 300 meters of pickup
- ETA versus ATA (estimated time of arrival versus actual time of arrival)
- Number of steps walked by riders
- Number of heading changes for driver
- Number of heading changes for rider
- Driver idle time (seconds)

User Signal

- Cancellation Rate (by driver or rider)
- Rider/driver Contact Rate
- PLE

Location Related

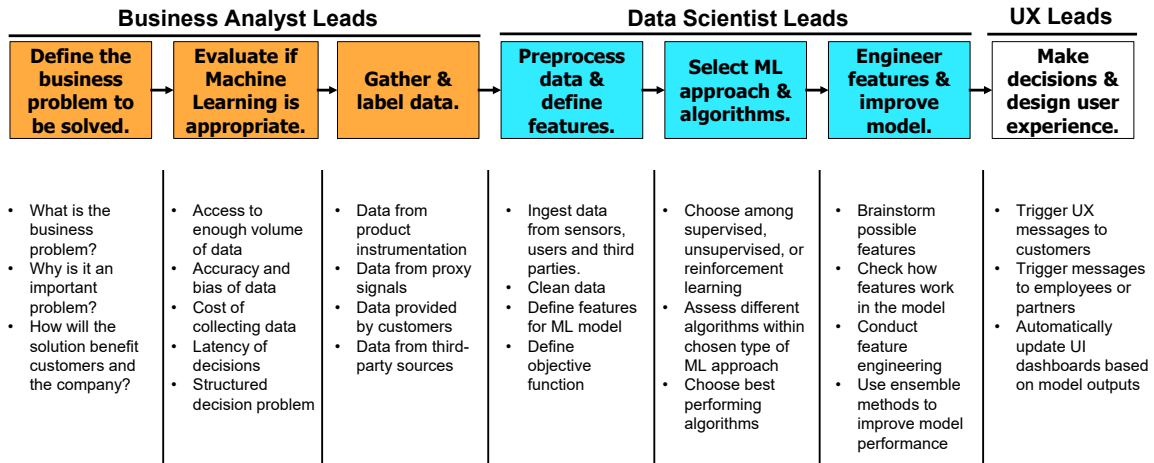
- X-axis, z-axis confidence swings

Mobile Events / Behavioral Indicators

- Number of taps until request
- Time spent in pin edit

Source: Uber.

Exhibit 7: Steps in Creating a Machine Learning Model



Source: Created by the authors.