

# Discovering the right place to check-in using web-based proximate selection

**Rui José**

Algoritmi Research Centre, University of Minho, Guimarães, Portugal

[rui@dsi.uminho.pt](mailto:rui@dsi.uminho.pt)

**Ana Inês Xavier**

Algoritmi Research Centre, University of Minho, Guimarães, Portugal

[a.ines.xavier@gmail.com](mailto:a.ines.xavier@gmail.com)

## Abstract

With information technology becoming increasingly embedded in our everyday physical world, there is a growing set of mobile applications that involve a connection with the digital representation of physical places. This association is normally initiated with a check-in procedure, through which a person asserts her presence at a particular place and determines the context for subsequent interactions. The common assumption is that a mobile application will be able to search the surrounding environment and present the user with the intended check-in target. However, in a world of ubiquitous place-based services, this assumption may no longer hold. A person in an urban environment would, at any moment, be surrounded by a large number of places, all of which could be regarded as possible interaction contexts for that person. In this work, we investigate the real-world challenges associated with wide-scale place selection and how the process can be affected by the place environment, by the position of the person in relation to the target place and by positioning errors. To study this reality, we used Google Places as a directory of georeferenced places. We conducted 14400 nearby place queries structured around different combinations of our three independent variables. The results suggest that the overall performance is poor, except for low density scenarios, and that this discovery process, albeit relevant, should always be combined with other place discovery approaches. The results also help to understand how this performance is affected by check-in positions and by the properties of the place environment.

## Keywords

Proximate selection, Place-based, Google Places, Mobile check-in

## 1 Introduction

Integration between physical and virtual environments is a major trend in information technologies. The digital world is becoming increasingly embedded in our everyday physical world, not just through technological infrastructures, but also through multiple digital services that are inherently blended with physical space. People now expect public and semi-public places to offer some type of digital representation. As they move around in their daily lives, and visit these places, they may occasionally want access to the respective digital services. For example, a user entering a restaurant may gain access to a menu, local quizzes, or a customer feedback survey. Likewise, a user may also want to connect with local digital devices, e.g. gaining mobile access to the interactive features of a public display. These associations between a physical environment and the respective digital entity are normally initiated through some type of check-in procedure, whereby a person explicitly asserts her presence at that particular context. This initiates a temporary situated session, which will then provide the context for subsequent interactions.

We refer to a place where a user wants to check-in as a target place. We assume that these places are symbolic locations, with a recognizable name and a physical existence, e.g. a store, a café or a hospital. Therefore, people are expected to have a shared understanding of what their current place is, and be able to select it from a list of places. Using the terminology proposed in [1], we consider that this check-in procedure may involve two distinct phases. The first is a discovery phase, which is essentially a proximate selection process [2] in which the system selects the most relevant candidates and presents them to the user, possibly ranked by some measure of relevance, e.g. distance. The second is the association phase, in which the user, or in some cases the system, selects the target place from that list and the association is established.

In this work, we are concerned with the real-world challenges and the main implications of large-scale place discovery. We are only considering the case where place discovery is based on queries made to a place directory. This model assumes that physical places have a digital representation with information about their physical location. It also assumes a place directory that supports location-based queries of nearby places. This should enable a mobile application to use information about its own location to obtain a list of nearby places, which can then be presented to the user for selection. A prevailing assumption about this type of place discovery is that when a person is in a particular place and makes a query for nearby places, that place would be the first in a list of results. The person would then be able to easily select it as the intended place for check-in. However, this assumption can be challenged by important technical issues associated with positioning services and the number of other nearby places [3]. In a world where place-based services become truly ubiquitous, a person in an urban environment would, at any moment, be surrounded by multiple candidate places corresponding to the many places around her, all of which could be regarded as possible interaction contexts for that person. Our research objective is to investigate to what extent a mobile application would actually be able to assist a person in the selection of a target place, when that person is in an environment that is densely populated with digital services. In particular, we aim to understand what are the main elements that may affect the reliability of proximate selection under those circumstances.

Our research methodology is based on data collected from a real directory of georeferenced places, more specifically Google Places. We use nearby place search queries to analyse the effect of three independent variables in the place discovery process: the characteristics of the place environment; the position of the person in relation to the position of the target place; and the error introduced by positioning systems. We selected 40 target places around the world and we have made 14400 queries representing different combinations of these independent variables. The results suggest that the overall performance is poor, except for low density scenarios, and that this discovery process, albeit relevant, should always be combined with other place discovery approaches.

## 2 Related work

Check-in procedures are very common in location-based platforms. The term check-in has been made popular, mainly by Foursquare, as a form of location sharing. The popularity of Foursquare itself made it a valuable platform for research on mobility and check-in practices. Rallapaly et al. [4] use check-in data to create predictive user mobility models. Based on individual mobility data, the authors identify the main factors that affect user mobility and use them to support individual mobility prediction. They suggest that these predictions can yield significant benefit to content prefetching and shared ride recommendation. Cramer et al. [5] study performative uses and emerging social practices surrounding location-sharing of Foursquare check-ins. The authors identify the emergent, but sometimes conflicting practices that affect the decision to check-in or not to check-in.

The more general concept of bridging between physical and virtual worlds has also been explored from many different perspectives. In its essence, the problem of how to create an association between a mobile device and a nearby entity is essentially a service discovery problem. Zhu et al. [6] describe a taxonomy of service discovery techniques, according to which the approach we are studying can be considered a directory-based service discovery infrastructure. This model assumes the existence of a place registry, which includes, at the very least, their spatial location. Our interaction model can be described as a proximate selection approach, where a list of objects or places is presented to the user in a such a way that the items that are more relevant to the user's current context are emphasized or made easier to choose [2]. In our own previous work [1], we have studied the key processes involved in the interaction between mobile devices and large screen displays. Based on the experience gathered from the creation of 3 different mobile applications, we identified 8 high-level processes that shape this form of interaction, more specifically: discovery, association, presence management, exploration, media upload, interface migration, direct

control and media download. The issues addressed in this study can be described as matching the discovery and the association phases.

In one of the earliest works on bridging physical and virtual environments, Kindberg et al. [7] proposed that places and devices should have a systematic link with their virtual counterparts. In their approach, places would have a beacon or a tag, announcing to nearby devices the place URL. This would then enable anyone with a browser, who was physically visiting the place, to find out about the respective information and services. The authors argue that the Web is the most appropriate technological framework for this type of general-purpose infrastructure and describe a concrete architecture to support the processes of registration, physical discovery, id resolution and user interaction. We share most of the goals, concepts and general approaches that were described in this work. However, our focus is not on the specifics of the underlying technical architecture, but on the viability of a place selection process based on place discovery. Hosio et al. propose the social surroundings (SocS) [8] platform as an application that allows people in third-places to get a glimpse of social profiles of nearby people. The goal is to encourage and facilitate conversations and promote a sense of community. In this case, place provides the virtual context that mediates communication.

The  $\phi^2$  ('Phi-square') [3] project also explores the connection between physical locations and their virtual representations. The authors use barcode generators and scanners to support check-in procedures that connect users with digital services. They also describe how the check-in experience can be affected by location data, which can take time to acquire and have a lower precision indoors, and by the number of locations presented to a user in a list, which grows with the density of venues. We study these effects and present a quantitative analysis of their concrete impact in key scenarios. Klokmose et al. [9] propose a system for allowing proximate selection of exhibited artworks based on Wi-Fi signals. A web server looks up the proximity of requesting devices to a number of points-of-interest, and generates a response containing the closest artwork. Weißker et al. [10] also use a Wi-Fi based technique to enable ad-hoc connections to a web-based locative game with real-time interaction. Users can join the game by simply connecting to the respective Wi-Fi network and opening their Web browser. These and many other systems assume the existence of physical objects or networks, which are physically present at the target place and can be used to trigger or assist the selection process. In our work, we are specifically considering the case where the entire selection process is based on virtual information about places.

Tacita [11] includes an explicit discovery process where the user is shown a map with the nearest public displays and their associated capabilities, as well as the user's own position. The goal is to allow people to choose a nearby display that may be suitable for accessing specific services. A similar feature is also available in Digifieds [12], which includes a mobile client that provides an interactive map with the locations of nearby Digifieds-enabled public displays, thus allowing people to more easily find a display where they can make a new post. In Instant Places [13], the nearby displays are also shown in the form of a list. In this case, the main goal is to facilitate the association, allowing people to recognize and select the display they already intend to use. These examples demonstrate another relevant scenario for proximate selection, i.e. pairing mobile applications with target devices in the public space.

This work expands a previous publication by the same authors, also addressing the same issue [14]. This previous publication was based on early results. The analysis was still superficial and only considered the distance to the ground position and the place environment. In this work, we analyse the full set of independent variables as well as their combination in regard to the distance to the position of the places.

### **3 Research design**

Our research methodology is based on data collected from place discovery queries made to Google Places. We identified Google Places as the ideal platform for our research because it has a global presence and pervasive usage. It also offers a rich API, which supports the discovery of places around a given position. We will use the expression target place to refer to the place in which the person intends to check-in. This place is assumed to be registered in Google Places and is associated with a specific geographic position, which we call the ground zero position of the target place. Ideally, when a person is in a place and makes a place discovery query, that same place should be the first in a list of nearby places. We hypothesize that in a real-world situation, this discovery process may be affected by three main variables: the characteristics of the place environment, the position of the person in relation to the ground zero position of the target place and the error introduced by positioning systems.

### 3.1 Place environment

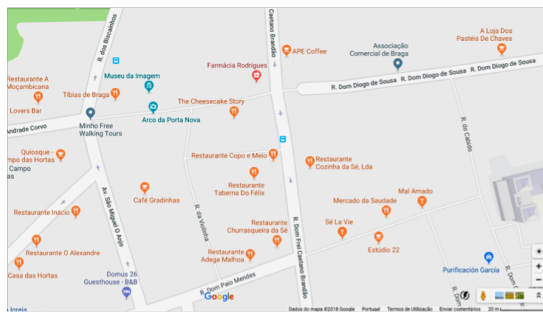
Place discovery can be strongly affected by the properties of the surrounding environment, particularly by the local place density and by the size of the target place. The more places there are, the more likely it will be that other places can be seen as being closer to the user. The larger the target place, the less likely it is that other places may be closer to its ground zero position.

To understand the effect of the place environment in the discovery process, we devised 4 base scenarios. They are all urban scenarios with a significant number of places, but they were carefully designed to represent realistic and relevant use cases in regard to different density levels and place sizes. Each scenario describes a concrete situation, where a person is expecting to check-in to a particular target place, of a particular size (small, large), in a physical environment characterised by a particular place density (low, high). The properties embedded in these scenarios, which are described in Table 1, should provide relevant insights on how the discovery process can be affected by the nature of different urban environments.

**Table 1.** Place scenarios used in this study.

Scenario	Condition	Description
1	High density Small target	The person is in an area with high concentration of services and multiple buildings. The target place is a restaurant located at the street level
2	Low density Small target	The person is in an area with low concentration of services, and only a few small and mostly residential buildings. The target place is a small business located at the street level
3	High density Large Target	The person is in an area with high concentration of services and multiple buildings. The target place is a large business located in the area, e.g. Hotel, Department store, shopping centre
4	Extreme density Small indoor target	The person is inside a shopping centre. The target place is a particular store within that shopping centre.

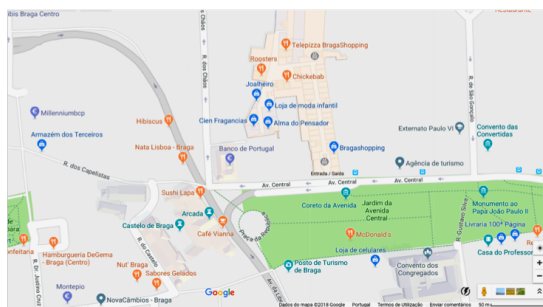
To have a representative sample of the diversity that can characterize physical and social settings, we selected 40 concrete places from different cities around the world, 10 for each of our place environment scenarios. These 40 places were selected by searching for appropriate locations in Google Maps. Figure 1 represents four concrete places used in this study, each corresponding to one of our four base scenarios.




(a) Sample place for scenario 1



(b) Sample place for scenario 2



(c) Sample place for scenario 3



(d) Sample place for scenario 4

Figure 1: Sample places for each of the four scenarios in this study

These scenarios are meant to represent clearly distinct environments in regard to place density and size. They are not meant to correspond to any categorization of place environments. This has made the whole process of selecting target places much easier. We focused only on environments that were clearly aligned with one of the scenarios, thus avoiding any borderline cases in which classification into one of the 4 scenarios could be ambiguous.

### 3.2 Positions at the place

The geographic information associated with places is often just a pair of coordinates. This corresponds to a single point in space, to which we will refer as the ground zero position for that place. However, places are not normally points in space, but rather geographic areas, with very diverse sizes and shapes, in which people can go into and move within. Therefore, in the real-world, the position of a person at a place is not likely to match the ground zero position of that place. Taking into account that people can check-in from multiple positions around the respective ground zero position, we introduced an independent variable that represents a sample of different check-in positions. More specifically, we created a 50x50m 9-point square grid, centred at the ground zero position (point 5), as illustrated in Figure 2.




Figure 2 A grid representing 9 different check-in points for each place

We refer to these 9 points as grid points. For each place, we have made an equal number of queries from each of these 9 grid points. The distance of the grid points to the ground zero position varies between 0 m (for point 5 at the centre of the grid) to 25 m (for points 2,4,6 and 8 at the middle of the grid boundaries) and approximately 35.4 meters (for points 1,3,7 and 9 at the grid corners). This grid is only meant to represent a simplification of the many diverse and complex realities associated with specific places, which we could not directly address in this study. It serves its purpose as an approximation to the most common place scenarios, mitigating the effects of the very diverse spatial properties of real places.

### 3.3 Positioning errors

Location technologies have significantly evolved in recent years and are always getting more accurate and precise. Still, any location technology involves estimation errors, which necessarily affect the position from which the place queries will be made. Considering the uncertainty caused by common positioning errors in mobile phones, we created a random error generator with a normal distribution and four different accuracy/precision levels. These errors were added to the positions used as input to the nearby place discovery queries. All the error levels have the same mean, which is zero, and a reference error corresponding to  $2\sigma$ , i.e. twice the standard deviation of the respective distribution. This means that in 95,4% of the cases, the error will be between 0 and the reference error distances represented in Table 2.

Table 2: Positioning errors introduced into query positions

Accuracy Level	0	1	2	3
$2\sigma$ (Reference error)	0 m	3 m	10 m	25 m

The error level 0 is the ideal case where the positioning system would not produce any positioning error. This is the only case, where the 10 queries per grid point were all made from exactly the same position. For all the other cases, a specific error was generated for each query, resulting in 10 slightly different query positions for each of these queries made from a single grid point.

### 3.4 Data collection process

The data for this study was obtained from place discovery queries made to the Google Places API. The places database used by Google Places is shared with Google Maps and contains many millions of records of all types of places, mainly small businesses and points of interest. For each of the 40 places in our study, we manually collected information about the respective position and unique place ID. The position of a place in Google Places is our ground zero position for that place, and is the point from which all other positions will be generated to represent the combined effect of various grid positions and positioning errors. The place ID is also needed to unambiguously identify a target place in a list of query results.

We developed and incrementally refined a custom tool to programmatically collect data according to our research variables. Figure 3 depicts the essence of the tasks supported by that query tool.

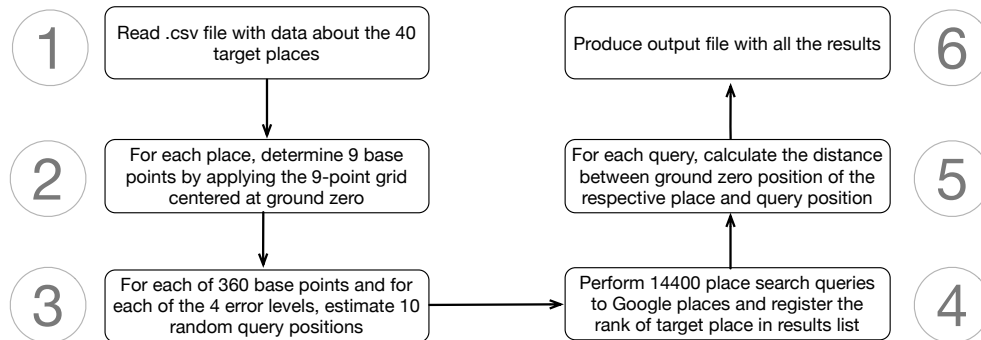


Figure 3: Overview of the programmatic process

The tool takes the list of 40 places with their IDs and ground zero positions. These 40 places are then combined with the 9 based points for each of those places and the 10 random errors for each of 4 error levels to generate 14400 unique Google Places queries, as summarised in Table 3.

Table 3: Structure of the queries by independent variable

#environment scenarios	4
#places per scenario	10
#grid points per place	9
#error levels per grid point at place	4
#queries per error level and grid point	10
#total place discovery queries	14400

The Google Places API supports many alternatives queries for nearby places. To guarantee that our query results were ranked only by distance, we used Nearby Search queries and carefully set the parameters to make sure that only distance was considered as a relevance criterion. The result of each query is a list of 20 nearby places ordered by their distance to the position used as the query parameter. We analysed each result list in search for the ID of the target place associated with that query. If the place was found, the rank position within that list was registered. Otherwise, the query was marked as a missed target. Additionally, we also registered the absolute distance between the target ground zero position and the position from which the query was actually made due to the combined effect of the grid points and positioning errors. To measure this distance, we used an implementation of the Cosine-Haversine formula [15].

## 4 Results

The analysis of the results is organized according to the effects of the 3 independent variables of this study: the position at the place, the level of positioning error and the place environment. We also include an additional section focused on the distance between the query position and the ground zero position of the target place. This distance is determined by the combined effect of two of our independent variables: the 9 grid points around the place and the four random positioning error levels.

### 4.1 Positioning errors

Regarding positioning errors, we consider how different error levels affected the rank position. Table 4 shows the rank results obtained with each of the four error levels considered in our study.

Table 4: Results per error level

Error Level (Reference error)	0 (0m)	1 (3m)	2 (10m)	3 (25m)
Study average positioning error	0m	1m	6m	29m
Average query distance to ground zero	33m	33m	34m	49m
Average discovery rank	12/20	12/20	12/20	14/20

The average positioning error represents the average of the random error that was effectively introduced during the study queries for each of those error levels. These are relative values, which measure the distance between a query position and the grid point from which the query should have been made. The Average query distance to ground zero represents the combined effect of check-in being made from different locations at the place and the positioning errors introduced by the positioning system. This is the absolute distance between the ground zero position of the target place and the concrete position used in each of the queries.

With this data, we can isolate the effect of positioning errors in the Average discovery rank. The results show that positioning errors seem to have very little effect in the system overall performance. Level 0, 1 and 2 all have the same average rank, which is 12, and level 3, shows only a very small increment to 14. Performance is not very good, but is stable, suggesting that positioning errors do not have a major effect. This can be interpreted as a consequence of the existence of multiple check-in points at the place, which in a sense are also a source of error. In some cases, the positioning error may potentiate that error, but in other cases, it may actually help to mitigate it. The only exception might be point 5, where the grid point matches the ground zero position and therefore, positioning errors are the only source of error.

### 4.2 Grid positions

We can also consider how the 9 grid positions, which represent a sample of a person possible positions at the place, affect the rank position. Table 5 shows the average results obtained with each of the 9 grid points across all the 40 places in this study.

Table 5: Results per grid point

Grid Point	5	6	4	8	2	1	7	3	9
Average positioning error	9m	9m	9m	9m	9m	9m	9m	9m	10m
Average discovery rank	4/20	12/20	12/20	12/20	12/20	15/20	15/20	15/20	15/20

The results are ordered by the average discovery rank (last row) associated with each grid point. They show that grid points have a very strong effect in the overall performance of the system. We can easily identify 3 main groups. A first group is composed by the grid corners (points 1,3,7 and 9). As corners, these are the most distant points (35m) and, not surprisingly, the worst performers with an average rank of 15. The second group is composed by the points located on the grid side lines (points 2, 4, 6 and 8). These are still distant from the ground zero position, but slightly closer (25 m) and their average rank is 12. Finally, point 5 is a special case because it corresponds to the central grid point, and consequently the best possible discovery grid point, as confirmed by the average rank of 4.


To further explore the discovery behaviour under the ideal situation, we specifically analysed the results obtained with the queries made from the central grid point (point 5) and without any positioning error (level 0). For 37 of our 40 target places the results were as expected, i.e. queries made directly from the ground zero position of the place would produce an average rank of 1. However, we identified 3 cases where the average rank was 2. Despite being a small difference, this is a somewhat unexpected result that required further explanation. When looking at this data more closely, we identified 3 places where even with distance 0 (point 5 and error level 0), the target place was second or even third in the results list. After a more detailed investigation, we found that these situations were either caused by the existence of distinct places with the same ground zero position or by the existence of multiple representations of the same place, with the same name but distinct ids and slightly different ground zero positions. These abnormal behaviours are part of the challenges that affect large-scale real-world place discovery.

While not all places correspond to 50x50 meters square area, we believe that this grid provides a convenient approximation to the fact that most check-ins will occur at some distance from the ground zero position of the target place. Regardless of the specific shape or usage patterns of a particular place, the grid provides a very generic distribution of queries made from different distances to that ground-zero position. At least for places that are not much larger than the 50x50 meters, this approach should be representative of the most common cases, while avoiding complex place models involving the size, shape, and organization of places, as well as the usage patterns and their effect in creating specific check-in hotspots.

#### 4.3 Combined effect of grid points and positioning errors

A key dimension of the analysis is the performance of the place queries in relation to the distance between the ground zero position and the position from which those queries were actually made. This distance is determined by the combined effect of two independent variables: the 9-point grid representing samples of people locations at the place; and the effect of positioning errors. This combined analysis is particularly relevant for assessing the viability of this place discovery process because it addresses the real-world situations, where both variables will always be present and they will both contribute to the overall performance of the process.

For each target place, and for each of its 9 grid points, we made 10 queries for each of the four error levels. The combined effect of these two independent variables is a set of 360 queries per place, all made from slightly different positions around the ground zero position of the target place, as represented in Figure 1.



**Fig. 4.** The query grid with its 9 grid points and the 360 query positions use in one of the places in this study.

In this analysis, we only considered the combined effect of grid points and positioning errors, expressed as the distance between the position used for each query and the ground-zero position of the place. Table 4 summarises the key data for this analysis.

**Table 6.** Frequency distribution of query distances and corresponding results

Distance	(0,10]	(10,20]	(20,30]	(30,40]	(40,50]	(50,60]	(60,70]	(70,80]	(80,90]
Query Frequency	868	492	2881	2941	5040	731	343	216	177
Average Rank	1.7	5,0	8.0	8.5	8.6	9.1	9.4	9.3	12,2
Missed Target Ratio	0.7%	5.7%	26.3%	31.3%	47.0%	61.1%	68.2%	68.5%	79.1%



This table represents, for various distance ranges: the frequency distribution of the queries made under those circumstances; the corresponding average position rank (only for queries in which the target place was included in the result list); and the percentage of queries in which the target place was not included in the result list. Another important result that is not explicit in this table is the total number of cases where the target place was not found (missed target situations). From the 14400 queries made in this study, 5323 (36%) queries did not include the target place in their list of results. We will analyse further the implications of these numbers, but for now the key idea is that the inclusion or not of those missed target situations must be considered when analysing the results.

Figure 5 presents a histogram of the rank results, where the missed targets were not considered.




Figure 5: Frequency distribution for rank positions (without missed target results)

This diagram shows that rank 1 is far from being the expected result. In fact, the peak at rank 1 is biased by the large number of queries that were made directly from ground zero positions. For each place, we have made 10 queries from point 5 and without any positioning error (error level 0). This means 400 queries for which the normal expectation would always be rank 1.

A deeper analysis should also consider the effect of missed target situations and particularly the relationship between query distances and rank results. These two variables are represented in Figure 6, which shows the number of queries made for the different distance ranges and the corresponding number of missed target queries.




Figure 6: Queries made per distance range with hits and missed targets

The distribution of the queries over different distance ranges is a direct consequence of our research design and is not grounded on any empirical data. However, the key observation is the number of missed targets and how that number is affected by distance. This relationship can be analysed more explicitly in Figure 7, which shows the relation between query distance and the percentage of missed targets.




Figure 7: Percentage distribution of missed targets in relation to query distance

For distances between 20 and 30 m, which we can classify as very common, the ratio of missed target places is already above 25%. Consequently, in only 1 attempt among 4, a person would actually find the specific place that she intends to check-in in the result list. For queries with distances above 50 metres, the ratio becomes higher than 50%, and missing the target place becomes the most likely outcome.

#### 4.4 Environment scenarios

Regarding the effect of the target place size and local place density in the discovery process, our research design considered 40 places representing 4 base scenarios with different density levels and place sizes, as previously shown in Table 1. The result of combining these with the other independent variables was a total of 3600 observations per scenario. The respective results are summarised in Table 7.

Table 7. Position ranking per scenario.

Scenario	Condition	Average rank A	Average rank B	#missed target	%missed target
1	High density Small target	10	16	2063	57%
2	Low density Small target	6	7	339	9%
3	High density Large Target	7	13	1510	42%
4	Extreme density Small indoor target	8	13	1411	39%

The average rank A corresponds to the average position of the target place in the lists of results (only for queries in which the target place was included in the result list). The average rank B considers all the queries, including those in which the target place was not in the results list. A rank value of 21 was used to represent those cases. Data about missed target indicates the number and percentage of queries made under that scenario and in which the target place was not part of the list of nearby places.

The results show that the characteristics of the environment can strongly influence the proximate selection process. The best performing scenario was scenario 2, corresponding to the expectation that performance would be better when there is a lower concentration of services. All other scenarios have somewhat similar results. It is clear, however, that the worst performing scenario is scenario 1, which combines high density of places with small target places.

## 5 Conclusions

In this study, we have made a systematic analysis of some of the key effects that may influence the proximate selection of places from location-based place directories. Overall, the results suggest that the performance of the whole process is very poor. From 14400 observations, 5323 cases were missed targets, i.e. the target place was not even in the list of 20 nearby places returned by the query. This means that, in 37% of the cases, the selection of the correct target place would not be possible at all, regardless of how much effort the user could put into going through all the alternatives presented. Places in low density scenarios are the only ones where the missed target ratio (9%) is low enough for the process to be seen as viable.

These results seem to challenge the prevailing assumption that when a search for nearby places is made from a particular place, that place would be the first in a list of results. When we consider our initial research objective of investigating to what extent a mobile application would actually be able to assist a person in the selection of a target place, the first insight is clearly that, under such circumstances, automated proximate selection would be unfeasible. Having the target place as the first element in the list of results was much more of an exception than a common reality.

A second key insight is that even the ability to show to the user a list where the target place is included can be a major challenge. From this perspective, the main criterion is not the position rank, but the missed target ratio. The results show that discovery performance can be too low for an effective operation, with a third of the queries failing to include the target place in the list of nearby places.

We believe our experimental setting to be representative of a broad range of common everyday situations and a valuable contribution to understand this problem domain, but we also acknowledge two fundamental limitations. The first is that the study is only based on information obtained directly from Google Places. Therefore, we did not consider any additional data about the target places and the possible effects of place size, place shape, mobility patterns, spatial containment relationships or overlapping places caused by ambiguous boundaries or multi-floor buildings. Instead, we always assumed the same grid size as an approximation to what might be very diverse place settings. While this may have been reasonable for smaller target places, it may affect the generalization of these results to larger places, such as a shopping centre, a stadium or a University campus. The other key limitation is not considering the role of many physical proximity techniques, which have been extensively explored for proximate selection and association, e.g. Wi-Fi detection, beacons or QR-codes. This is not because of any judgment about the relevance of those techniques. These technologies can offer valuable alternatives for place discovery and may strongly influence place discovery patterns. The option not to address them in this study was merely an issue of focus. We wanted to specifically address the unique role that place-based queries can play as the fall-back solution and as a key expectation that most check-in procedures can have anywhere. As shown by the results of our study, a combination of different techniques would be essential to combine large scale availability with fine-grained discovery.

As future work, we expect to extend the findings of this research to integration approaches in which nearby discovery queries are combined with other proximity-based technologies. The key goal should be to understand their complementary roles and what might be the best way to combine their mutual strengths.

## References

1. José, R., Cardoso, J. C. S., Alt, F., Clinch, S., & Davies, N. (2013). Mobile Applications for Open Display Networks - Common Design Considerations. *Proceedings of the 2nd International Symposium on Pervasive Displays (PerDis '13)*, 97–102. doi:10.1145/2491568.2491590
2. Dey, A. K. (2010). Context-Aware Computing. In J. Krumm (Ed.), *Ubiquitous Computing Fundamentals*. CRC Press.
3. Büttner, S., Cramer, H., Rost, M., Belloni, N., & Holmquist, L. E. (2010).  $\phi^2$ : Exploring Physical Check-Ins for Location-Based Services. *Society*, 395–396.
4. Rallapalli, S., Dong, W., Lee, G. M., Chen, Y. C., & Qiu, L. (2013). Analysis and applications of smartphone user mobility. *Proceedings - IEEE INFOCOM*, (December 2011), 3465–3470. doi:10.1109/INFOCOM.2013.6567182
5. Cramer, H., Rost, M., & Holmquist, L. E. (2011). Performing a Check - in : Emerging Practices , Norms and Conflicts in Location - Sharing Using Foursquare. In *Proceedings of the 13th*

- International Conference on Human Computer Interaction with Mobile Devices and Services* (pp. 57–66). ACM.
6. Zhu, F., Mutka, M., & Ni, L. (2005). Service Discovery in Pervasive Computing Environments. *IEEE Pervasive Computing*, 4(4), 81–90.
  7. Kindberg, T., Barton, J., Morgan, J., Becker, G., Caswell, D., Debaty, P., ... Spasojevic, M. (2002). People , Places , Things : Web Presence for the Real World People , Places , Things : Web Presence for the Real World. *Mobile Networks and Applications*, 7(5), 365–376.
  8. Hosio, S., Kukka, H., & Riekkki, J. (2010). Social surroundings: Bridging the virtual and physical divide. *IEEE Multimedia*, 17(2), 26–33. doi:10.1109/MMUL.2010.3
  9. Klokmose, C. N., Korn, M., & Blunck, H. (2014). WiFi proximity detection in mobile web applications. *Proceedings of the 2014 ACM SIGCHI symposium on Engineering interactive computing systems - EICS '14*, 123–128. doi:10.1145/2607023.2610281
  10. Weißker, T., Berst, A., Hartmann, J., & Echtler, F. (2016). The Massive Mobile Multiuser Framework : Enabling Ad-hoc Realtime Interaction on Public Displays with Mobile Devices. *Proceedings of the 5th ACM International Symposium on Pervasive Displays - PerDis '16*, 168–174. doi:10.1145/2914920.2915004
  11. Kubitzka, T., Clinch, S., Davies, N., & Langheinrich, M. (2013). Using mobile devices to personalize pervasive displays. *ACM SIGMOBILE Mobile Computing and Communications Review*, 16(4), 26–27. doi:10.1145/2436196.2436211
  12. Alt, F., Kubitzka, T., Bial, D., Zaidan, F., Ortel, M., Zurmaar, B., ... Schmidt, A. (2011). Digifieds: Insights into Deploying Digital Public Notice Areas in the Wild. In *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia - MUM '11* (pp. 165–174). New York, New York, USA: ACM Press. doi:10.1145/2107596.2107618
  13. José, R., Pinto, H., Silva, B., & Melro, A. (2013). Pins and Posters: Paradigms for Content Publication on Situated Displays. *IEEE Computer Graphics and Applications*, 33(2), 64–72.
  14. Xavier, A. I., & José, R. (2018). I am in here : Implicit assumptions about proximate selection of nearby places. In *WorldCist'18 - 6th World Conference on Information Systems and Technologies*. Naples, Italy: Springer.
  15. Robusto, C. C. (1957). The cosine-haversine formula. *The American Mathematical Monthly*, 64(1), 38–40.