<http://www.diva-portal.org>

This is the published version of a paper published in *Künstliche Intelligenz*.

Citation for the original published paper (version of record):

Richter, K-F. (2017)

Identifying Landmark Candidates Beyond Toy Examples: A Critical Discussion and Some Way Forward.

Künstliche Intelligenz, 31(2): 135-139

<https://doi.org/10.1007/s13218-016-0477-1>

Access to the published version may require subscription.

N.B. When citing this work, cite the original published paper.

This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>)

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:umu:diva-137162>

Identifying Landmark Candidates Beyond Toy Examples

A Critical Discussion and Some Way Forward

Kai-Florian Richter^{1,2}

Received: 14 September 2016 / Accepted: 18 November 2016 / Published online: 1 February 2017
© The Author(s) 2017. This article is published with open access at Springerlink.com

Abstract Incorporating references to landmarks in navigation systems requires having data on potential landmarks in the first place. While there have been many approaches in the scientific literature for identifying landmark candidates, these have hardly been picked up in actual, running systems. One major obstacle for this to happen may be that most—if not all—approaches presented so far are not scalable due to their underlying data requirements. In this paper, I will critically discuss existing approaches in light of their scalability. I will then suggest a way forward to more scalable solutions by combining in a smart way aspects of different approaches.

Keywords Landmark identification · Personalization · Human–computer interaction · User-generated content

1 Introduction

Landmarks are crucial in human conceptualization and understanding of an environment; they are also omnipresent in human communication about space. Getting computers to use landmarks in their communication about space as well would make for a much more natural, and much richer human–computer interaction [9]. However, despite a lot of research in this area, landmarks are hardly utilized in running (commercial) systems. One reason for this may be that

it is still very hard to reliably determine suitable landmark references uniformly across environments.

Generally, there are two major steps necessary in order to enable computational systems to incorporate landmarks into their interaction with human users: (1) the identification of what may serve as a landmark in principle—termed *landmark candidates* in the following; (2) the identification of which of these candidates is most suitable in a given situation [9]. This paper will focus on the first step.

Most generally and most usefully landmarks are defined to be everything that sticks out of the surroundings [6]. Also, it is important to note that being a landmark is a (graded) property of an object; arguably, there are no genuine ‘landmark’ objects [9]. Over the years, several methods have been suggested for identifying objects that have landmark characteristics, i.e., stick out of their surrounds and, thus, may be assigned a *landmarkness* property.

In the following, I will have a closer look at these methods and discuss their advantages and disadvantages (summarized in Table 1; see also [8]). In particular, I will discuss whether and how these approaches are suited to be used beyond the case studies used in the respective publications, which are often rather restricted in their scope—thus, the possibly somewhat snidely reference to ‘toy examples’ in this article’s title. I will also assess their potential for providing personalized landmark information, which may further increase the effectiveness of landmark references. I will then present a possible way forward by outlining how the different methods may be combined in a smart way to develop a more scalable solution, which uses aspects of personalization to ensure usefulness both for the individual user and the community at large.

✉ Kai-Florian Richter
kfr@kfrichter.org

¹ Department of Geography, University of Zurich,
Winterthurerstrasse 190, 8057 Zurich, Switzerland

² Department of Computing, Umeå University, 901 87 Umeå,
Sweden

Table 1 Advantages and disadvantages of the different approaches to determining landmark candidates

Advantage	Disadvantage
Feature-based salience	
Detailed	Data intensive
Explanatory	Parameter tuning
Extendable, adaptable	Transfer to other types
Personalization	
Category-based salience	
Little data needed	Every individual treated the same
Easy to compute	Ambiguous categories
Personalization	Depends on classification scheme
Harvesting UGC	
Potential for global coverage	Biased data
Existing methods	Suitability issues
Machine learning approaches	
Simple discrimination task	Dependency on base data
Strong personalization	Not explanatory
Feedback mechanism	
Collecting landmark candidates	
Dedicated approach	Dependency on base data
Truly uniform coverage	Dedicated user base
Personalization	

2 Identifying Landmark Candidates: A Review of Existing Methods

To the best of the author's knowledge, the earliest approach to computing the landmarkness of geographic objects was proposed by Raubal and Winter [7]. Focusing on building façades, they calculate salience of an object as a weighted sum over a range of attributes, which are classified as being either visual, structural, or semantic [11]. These attributes represent various aspects of a façade, such as its color, size, or whether any (storefront) signage is present.

Raubal and Winter developed quantitative measures for each attribute. The attributes are explicitly represented, so it would be easy to create an explanatory model for why something is (not) considered a landmark. A weighted sum also makes it easy to extend and adapt the model, or to provide personalized settings. But these advantages at the same time point to the model's weaknesses. It requires a lot of detailed (geographic) attribute data to populate the different measures with values. And since each attribute is weighted against the others, it also requires a lot of parameter tuning, which may make it difficult to transfer the model to other kinds of objects than building façades, or other contexts more generally [10].

The need for detailed data can be overcome by using categories rather than individuals [3]. This then only requires data on an object's type and geographic location in order

to determine an object's suitability as a landmark. It still requires parametrization though, since the different categories need to be ranked according to their general, average landmarkness. However, this seems less problematic as in the Raubal and Winter approach, and may, for example, be done via expert interviews [3]. On the other hand, the heuristic approach of treating the same every individual of a given category may get landmarkness very wrong for some of these individuals. Also it may not always be unambiguously possible to assign a single category to every geographic object, and there is a clear dependency on the chosen categorization scheme, i.e., changing the scheme will also change landmarkness of objects—even though the objects themselves did not change.

Given the increasing availability and spread of user-generated content (UGC)—a lot of which has geographic components—it seems promising trying to exploit such data for determining landmark candidates. Several approaches exist using documents (e.g., [12]) or (annotated) photographs (e.g., [2, 13]) to extract landmarks. Using such data leads to a potentially global coverage. It also becomes possible to make use of established methods from data mining or geographic information retrieval. However, since the data was not specifically designed for covering landmark candidates, it will likely contain biases towards specific regions, specific types of geographic objects, or specific attributes.

Since user-generated content in the end may fall short in replacing dedicated data sets of landmark candidates, but these data sets do not really exist, another potential pathway is to let users create such data sets. One option is to learn landmark candidates from user behavior [4], for example, by having users identify geographic objects that they deem suitable landmarks in a training phase. Later, the system may then pick similar objects, which would also need to work in previously unencountered environments. What 'similar' means needs to be defined, but may, for example, use feature vectors based on the objects' attributes [4].

Such a learning system, which in its implementation can rely on a relatively simple discrimination task, clearly leads to strong personalization. However, the identification and selection of objects still relies on some underlying base data (e.g., some topographic data set), and the system will not be able to provide explanations as of why some specific object is a landmark candidate beyond some similarity value to a previously learned object.

Another option for creating a dedicated data set for landmark candidates is employing principles and methods of user-generated content, i.e., to have users collect data on potential landmarks, which then can be used with any of the existing methods for determining the most suitable landmark in a given situation. Such collection most likely needs to happen in-situ, i.e., in the field, and may ask the users directly to contribute landmarks [14], or may have

a game-like character [1]. Clearly, if successful, such an approach will lead to data that is specifically tailored to determining landmark candidates. It also has the potential for a truly uniform coverage; users may collect landmarks in city centers as much as in residential neighborhoods or rural areas. The resulting landmark candidates may also be personalized, most simply by preferring those contributed by a specific user, for example, those collected by oneself.

However, determining which geographic object some user is meaning to add to the data set either still requires some comprehensive geographic base data or some elaborate interaction steps in adding geographic attribute data on the fly. And as any project relying on user contributions, collecting landmark candidates this way requires a dedicated user base in order to reach sufficient coverage.

Table 1 summarizes advantages and disadvantages of the presented kinds of approaches to identifying landmark candidates.

3 Towards Scalable Solutions

As the discussion in the previous section has addressed several times (see also Table 1), a major challenge to identifying landmark candidates reliably and to sufficient numbers across environments is the lack of data that consistently provides detailed enough information on geographic objects and their attributes. Accordingly, it seems rather optimistic to base the identification of landmark candidates on such data if this is to be done on a large scale spanning whole cities, countries, or even globally. Raubal and Winter's approach has been very important conceptually for driving research, but it is not scalable.

As we have also seen, generally a lightweight approach to identifying landmark candidates seems more promising, as for example the one chosen by Duckham et al. [3]. Relying only on type and location information has very few computational demands. It also reduces demands posed on the underlying data. But as discussed in the previous section, such an approach has some disadvantages as well, namely potential ambiguity in categorization and the fact that not all individuals of a category will be equally suitable as landmarks. Therefore, such an approach would ideally be augmented with some mechanisms to flexibly adapting both category and suitability ratings. Overall, a smart combination of principles implemented in existing approaches might present a solution here, further discussed in the following.

In the proposed new approach, uniformly assigning the same landmarkness value to all objects of a specific category will form the base assessment of landmark suitability. Any application using this landmark data may then include feedback mechanisms that would allow users to mark the

usefulness of a given object up or down, and also to disagree with its categorization. These proposed changes can initially be kept to the user who made them, i.e., personalize their landmarkness settings. Aggregated over multiple users, these proposed changes may also change general settings of both suitability ratings and categorization.

In some more detail, while using types provides an easy, lightweight approach, uniformity of landmarkness in a given category will not hold in the real world. For example, some places of worship will be more salient than others; compare St. Peter's Cathedral with a small 'place of worship' room hidden away at an airport. These differences may be captured by enabling users of a system to provide such feedback. If users are presented with a landmark they deem unsuitable, for instance, if they cannot even detect it, there may be a simple mechanism to mark them as not useful in the system. In the same manner, they may also mark referenced objects as particularly useful landmarks (e.g., by using simple '+' and '-' or 'thumbs-up' and 'thumbs-down' buttons). This would then change the landmarkness value for the individual object, initially only for the individual user. If a specific user repeatedly marks down (or up) objects of the same type, say street furniture or retail outlets, a system may also infer user preferences from this behavior and, thus, adapt globally landmark selection for this user accordingly.

Initially, this will lead to a type-based, but more personalized landmark selection for individual users. However, as so often with such approaches, user behavior may also be aggregated to perform general adaptations. For example, if repeated rejections of an object occur across multiple users, this may be taken as indication that the specific object is generally not suited as a landmark. Following the same reasoning, such behavior may also lead to adapting landmarkness values for a whole category of objects. In case repeatedly multiple objects of the same category get marked down by multiple users, this may indicate that the initial judgement of the category's suitability as landmark candidate needs to be re-evaluated.

There are some caveats with the proposed approach. Clearly, also a type-based approach to identifying landmark candidates depends on an underlying data set. While this set has less demands on object attributes, it would still need to provide a reasonably uniform coverage of objects of various categories with their geographic location. It is doubtful whether such a data set currently exists even for a single city. For example, experiments presented in [14] and an analysis of the Swiss OSM data [5] have shown that even in a highly developed country, such as Switzerland, geographic data may not be uniformly fit for use in such an approach. But similar to some existing social network platforms, users may be encouraged to submit additional landmark candidates themselves, either integrated into a

navigation application, or probably more usefully as a standalone application. Such user-generated content comes with the usual issues, such as potential errors or even malicious user behavior. But again, firstly this contributed data may be used to improve navigation experience for the contributing user. As such, it may not be necessary to make the data available to other users immediately, but some moderation mechanisms could be incorporated. One such option may be to set up a game-like application, where newly contributed landmark candidates would first need to be ‘found’ by other users, before they will be used globally; similar to the approach in [1].

When using a type-based approach assigning landmarkness values also strongly depends on the underlying categorization scheme—the ‘object ontology’ if you will. And since the type names will most likely also be used when referring to landmark objects in user interaction (e.g., ‘turn left at the church’, ‘move towards the museum’) this scheme also has a strong influence on user interaction. It is highly likely that not all users will agree with how an object is referred to all the time, i.e., they may have a different conceptualization of what kind of object it is than what the system assumes. Again, it would be possible to implement some feedback mechanism that allows users to change an object’s categorization (or just its label). This would first and foremost result in personalization, i.e., and adaptation for an individual user. But as with usefulness, these changes may be feedback into the overall system and with multiple users providing the same, or very similar, feedback, categorization may change globally.

Clearly, implementing such a new approach to identifying landmark candidates requires thorough evaluation and testing. This should be done on at least three levels: targeted studies that test the usability and usefulness of the new approach’s individual elements; a medium-term study that tests how and where individual users employ the feedback mechanisms or add new landmark information; and finally a medium- to long-term study with multiple users observing the effects and interplay of the different feedback mechanisms on global landmarkness settings. The first level of evaluation is mainly meant to ensure that the implemented procedures and interaction mechanisms actually work. It may follow ‘standard’ procedures of user and usability testing and should also be preceded or accompanied by software testing and some geo-spatial analysis of the underlying data—the base landmarkness assessments and their distribution. The second level will evaluate how the different implemented components interplay in the longer run, for example, whether some of them counter each other and how (much) personalization will occur. It will also allow for assessing user acceptance of the different mechanisms and their willingness to continuously use the system. Finally, the third level of evaluation will provide similar

insights to the second level, but in addition will shed some light on desired and undesired effects of user and software components interplay when multiple users with potentially conflicting interests are involved. It will also show whether user contributions will be reasonably uniformly distributed or whether there are similar biases to data distribution as we observe in many UGC data sets. The latter case would then require some counter-measures, for example, by setting up incentives to explore potential landmarks in less covered areas in some game-like settings¹—which would need to be evaluated again of course.

To conclude, using a type-based approach ensures that there is a reasonable base level of useful landmark candidates, which can be determined quickly and with low effort. Providing a range of feedback and interaction mechanisms then allows for fine-tuning such a system to accommodate individual differences, but also mis-classifications that are bound to occur in such a heuristic approach. Clearly, we cannot expect users to evaluate all landmark references all the time, but providing feedback will have immediate benefits, particularly for those references that did not work well for a user. Thus, given an engaging, unobtrusive user interface a smart combination of a simple, but well-balanced base selection of landmark candidates with elaborate inference mechanisms based on user feedback may prove to be the scalable solution missing so far.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Bell M, Reeves S, Brown B, Sherwood S, MacMillan D, Ferguson J, Chalmers M (2009) Eyespy: supporting navigation through play. In: Proceedings of the 27th international conference on human factors in computing systems. ACM, New York, pp 123–132
2. Crandall DJ, Backstrom L, Huttenlocher D, Kleinberg J (2009) Mapping the world’s photos. In: Proceedings of the 18th international conference on World Wide Web. ACM, New York, pp 761–770
3. Duckham M, Winter S, Robinson M (2010) Including landmarks in routing instructions. *J Locat Based Serv* 4(1):28–52
4. Götze J, Boye J (2015) Learning landmark salience models from users’ route instructions. In: 12th symposium on location based services, Augsburg, Germany

¹ Maybe similar to how Pokemon Go (<http://www.pokemongo.com>) works, even though its success is most likely very hard to recreate even with professional game developers involved.

5. Hauser L (2014) OpenStreetMap in der Schweiz – Thematische, räumliche und zeitliche Analyse von “Points of Interest”. MSc thesis, Department of Geography, University of Zurich
6. Presson CC, Montello DR (1988) Points of reference in spatial cognition: stalking the elusive landmark. *Br J Dev Psychol* 6:378–381
7. Raubal M, Winter S (2002) Enriching wayfinding instructions with local landmarks. In: Egenhofer M, Mark D (eds) *Geographic information science*, vol 2478. *Lecture notes in computer science*. Springer, Berlin, pp 243–259
8. Richter KF (2013) Prospects and challenges of landmarks in navigation services. In: Raubal M, Mark DM, Frank AU (eds) *Cognitive and linguistic aspects of geographic space—new perspectives on geographic information research. Lecture notes in geoinformation and cartography*. Springer, Berlin, pp 83–97
9. Richter KF, Winter S (2014) *Landmarks—GIScience for intelligent services*. Springer International Publishing, Berlin
10. Sadeghian P, Kantardzic M (2008) The new generation of automatic landmark detection systems: challenges and guidelines. *Spat Cognit Comput* 8(3):252–287
11. Sorrows ME, Hirtle SC (1999) The nature of landmarks for real and electronic spaces. In: Freksa C, Mark DM (eds) *Spatial information theory*, vol 1661. *Lecture notes in computer science*. Springer, Berlin, pp 37–50
12. Tezuka T, Tanaka K (2005) Landmark extraction: a web mining approach. In: Cohn AG, Mark DM (eds) *Spatial information theory*, vol 3693. *Lecture notes in computer science*. Springer, Berlin, pp 379–396
13. Wither J, Au CE, Rischpater R, Grzeszczuk R (2013) Moving beyond the map: automated landmark based pedestrian guidance using street level panoramas. In: *Proceedings of the 15th international conference on human–computer interaction with mobile devices and services*. ACM, New York, pp 203–212
14. Wolfensberger M, Richter KF (2015) A mobile application for a user-generated collection of landmarks. In: Gensel J, Tomko M (eds) *Web and wireless geographical information systems. Lecture notes in computer science*, vol 9080. Springer International Publishing, Berlin, pp 3–19