

Near2Me: Design and Evaluation of a
Personalized Recommender and
Explorer for Off-the-beaten-track
Travel Destinations

María Luz Caballero

ISBN: 978-90-444-0960-4

A catalogue record is available from the Eindhoven University of Technology Library

ISBN: 978-90-444-0960-4

(Eindverslagen Stan Ackermans Instituut ; 2010/040)

GENERAL SUMMARY

This project consists of the design and evaluation of Near2me, a personalized recommender and explorer for off-the-beaten-track travel destinations, developed within PetaMedia (a European Network of Excellence devoted to the study of the synergetic combination of social network structure, media content analysis, and user-generated annotations).

This project was directed by two goals:

1. The design and evaluation of a novel application that makes use of PetaMedia-developed technologies in a context meaningful to users.
2. The exploration of a design methodology that puts the user at the center of the design process of technology-driven novel applications.

We had to define the exact use case and functionality simultaneously while building the application, because our design was technology-driven and initially shaped by the capabilities offered by PetaMedia-developed technology. To make sure that the final product responding to our designed use case would still address a realistic, existing scenario, we guided and documented the design through a framework based on our design judgments and conducted a preliminary user study to identify the motivations of individuals that engage in Independent Travel that we used as guide for our design judgments:

1. Independent travelers' search of an authentic exchange, away from "manufactured" tourism experiences.
2. Independent travelers' search of a more personal interaction with the city, an experience that arises from the unique combination of the personality of the visitor and the place.

We developed Near2me, a personalized travel recommender and explorer for off-the-beaten-track travel destinations, which uses six different PetaMedia developed technologies. Near2me takes its cues for recommendation from its users' spontaneous photography capturing behavior while traveling, obtaining a pool of crowd-sourced information of recommendable places from the geotags attached to its users' photographs. Its recommendation algorithm uses media capturing-related situational knowledge to produce recommendations based on the users' interests rather than general popularity. Functionality is added to allow the user to (1) explore and evaluate these recommendations to decide if they are worth following, (2) control recommendation direction and follow leads within recommendations by applying filters to recommended items when browsing, (3) link recommendation to external information to obtain background and practical information, and (4) save and use selected recommendations while visiting a region.

From the abstraction of the information structure in Near2me's design, we were able to derive a new model for resource-user-tags relationships that exploits the technologies developed within PetaMedia to exploit resource-user-tags relationships beyond the particular domain of Independent Travel.

Finally we conducted a user test that demonstrated that Near2me is consistent with the user motivations found in the preliminary user research: participants recognized the value of the different components designed into Near2me for the fulfillment of these motivations. Participants mentioned that they would use Near2me and were able to offer a wide variety of different scenarios based on real past travels or present travel planning situations in their lives. The concept design was validated, suggestions for improvement in future iterations were obtained, and directives for the evaluation of the next iterations are provided. The success obtained in Near2me's evaluation also suggests there is potential in the further exploration of the methodology used during Near2me's development to keep the user at the center of the technology-driven design of novel applications.

EXPLOITATION RESULTS

- We **designed Near2me**, a personalized travel recommender and explorer for off-the-beaten-track travel destinations, which uses six different PetaMedia developed technologies, and was evaluated positively by users in terms of its ability to fulfill their motivation while engaging in Independent Travel (Part 1 - Section 5).
- We **obtained suggestions for the improvement of Near2me** in its subsequent design iterations through user evaluation (Part 1 – Sub-section 7.2), notably: adding recommendation situational information beyond location (e.g. time and date, seasonality, weather), providing the option to opt out from any personal identification by other users, addressing recommendation multiple-user travels.
- We **offer detailed proceedings for the evaluation of a working prototype** of Near2me, as planned for the next design iterations (Part1 – Section 9).
- We formalized the abstract information structure and user-resource-tag relationship that lies at the core of Near2me, providing a theoretical model to facilitate PetaMedia technology portability to other problems conforming to the same characteristics, and making a **contribution to the current understanding of the user-resource-tag model** found in the literature (Part1 – Section 8).
- We conducted a study of the current landscape of travel recommenders and explorers through which we provided an extensive account of the state of the art in the field, and **distinguished Near2me as a distinctive offer** in terms of target users and capabilities for its simultaneous emphasis in: personalization based on user interests, off-mass-tourism pool of recommendable places, authentic portrayal of recommendations for evaluation (Part 1 – Section 3).
- We **identified the motivations of users** who engage in Independent Travel through a user study: search of an authentic exchange, away from "manufactured" tourism experiences, and search of a more personal interaction with the city that arises from the unique combination of the personality of the visitor and the place (Part 1 – Section 2). These **should remain the guiding principle** behind Near2me in all subsequent iterations.
- We found support for the **further study of the design methodology that keeps the user at the center of the design process of technology-driven novel applications, developed for Near2me**, which demonstrated to provide successful design guidance in the case of Near2me (Part 2)

TABLE OF CONTENTS

GENERAL INTRODUCTION	7
1. Objectives	7
<i>1.1. Design and Evaluation of a Novel Application that makes use of PetaMedia-developed Technologies in a Context Meaningful to Users</i>	7
<i>1.2. Exploration of a Design Methodology that Puts the User at the Center of the Design Process of Technology-driven Novel Applications</i>	8
2. Structure of the report	8

Part 1- Near2Me: a Personalized Recommender and Explorer for Off-the-beaten-track Travel Destinations **10**

ABSTRACT	10
1. BACKGROUND	10
2. PRELIMINARY USER RESEARCH	12
3. CURRENT LANDSCAPE OF TRAVEL RECOMMENDERS AND EXPLORERS	12
4. OBJECTIVE AND REQUIREMENTS	17
5. NEAR2ME: THE CONCEPT	18
5.1 Basic Description	18
5.2 Component Parts	19
<i>5.2.1 The Recommendation Algorithm</i>	19
<i>5.2.2 Identification and Evaluation</i>	20
<i>5.2.3 Recommendation Control and Lead Following</i>	21
<i>5.2.4 Further Exploration of the Recommendations</i>	22
<i>5.2.5 Recommendation exploitation</i>	22
5.3 Comparison to the Other Recommenders and Explorers studied in our Survey	23
5.4 Interface and Interaction Paradigm	24
5.5 Use Scenario	24
6. THE PETAMEDIA TECHNOLOGIES BEHIND THE CONCEPT	28
7. CONCEPT EVALUATION	29
7.1 User evaluation Methodology and Choice of Prototyping Technique	29
7.2 Results and Conclusions	30
8. CONCEPT ABSTRACTION: INFORMATION STRUCTURE BEHIND NEAR2ME	34
9. FUTURE WORK	35
9.1 Development of a Working Prototype	35
9.2 Recommendations for the Field Trials	35
10. REFERENCES	36

Part 2- Putting the User at the Center of the Design Process of Technology-driven Novel Applications **39**

ABSTRACT	39
1. THE PROBLEM	39

2. BACKGROUND	40
3. DESIGN METHODOLOGY	40
4. EVALUATION METHODOLOGY	42
5. CONCLUSIONS	42
6. REFERENCES	43
<hr/>	
APPENDIX 1: RECOMMENDATIONS FOR THE FIELD TRIALS	44

Near2Me: Design and Evaluation of a Personalized Recommender and Explorer for Off-the-beaten-track Travel Destinations

Author: María Luz Caballero
luzcaballero@gmail.com

User System Interaction
Technische Universiteit Eindhoven

Work done at: PetaMedia - TU Delft

Academic advisor: Prof.dr. L. Hardman (TU Eindhoven)

Company coach: Ms. M.A. Larson Ph.D. (TU Delft)

GENERAL INTRODUCTION

This report describes the work behind the design and evaluation of Near2me, a personalized recommender for off-the-beaten-track travel destinations, developed within PetaMedia. This project was directed by two goals, aligned with PetaMedia's dual vision of research (simultaneously technology- and user-driven¹):

1. The design and evaluation of a novel application that makes use of PetaMedia-developed technologies in a context meaningful to users.
2. The exploration of a design methodology that puts the user at the center of the design process of technology-driven novel applications.

These goals are explained in more detail below.

1. Objectives

1.1. Design and Evaluation of a Novel Application that makes use of PetaMedia-developed Technologies in a Context Meaningful to Users

PetaMedia is a Network of Excellence funded by the European Commission's Seventh Framework Programme. Four partners form the core of PetaMedia, each representing a national network: Delft University of Technology (NL), Ecole Polytechnique Federale de Lausanne (CH), Queen Mary University London (GB) and Technical University of Berlin (DE). PetaMedia is directed towards exploring the synergetic combination of

- user-based collaborative tagging,
- social peer-to-peer networks,
- and multimedia content analysis

(called from now on “*the triple synergy*”).

"[PetaMedia's] research questions are all driven by the [...] needs of real-life users in real-world multimedia environments."² Therefore, the PetaMedia workplan calls for "user trials that measure the success of the system in terms of its ability to satisfy users both in terms of ease of use and in terms of meeting their [...] needs."³ In order to test PetaMedia technologies in a real-world context to evaluate the experience of actual users, an application that makes use of PetaMedia developed technologies in a context meaningful to users has to be developed. This application has to help users fulfill their goals through the completion of low-level tasks afforded by the in-context deployment of PetaMedia technologies.

In relation to this goal, the scope of this project includes:

- The conceptual design of the application, and the evaluation of this concept
- Guidelines for the evaluation of a working prototype that will be developed based on the above mentioned design

¹ PetaMedia Integrative Research Plan (IRP), Deliv6.1-150306-TUD, page 8.

² Petamedia IRP, Deliv6.1-150306-TUD, page 9.

³ Petamedia IRP, Deliv6.1-150306-TUD, page 17.

1.2. Exploration of a Design Methodology that Puts the User at the Center of the Design Process of Technology-driven Novel Applications

Traditional User-centered Design practices are inadequate for PetaMedia's purposes. As a Network of Excellence, the goal of PetaMedia is to produce and lead the development of new technologies and systems in its field. This is why PetaMedia, when it comes to the development of systems, complements its user-driven vision with technology-driven vision⁴. However, the reconciliation of these two approaches is not straightforward.

In the first place, traditional User-Centered practices start with an "analysis" phase⁵ that consists of studying users to identify who will use the product, what they will use it for and in which context they will use it. The Usability Professionals Association lists as suggested activities for this phase: creation of user profiles, task analysis, documentation of user scenarios, and documentation of user performance requirements, among others. If answers to "*by whom, how and when/where* the system will be used" are expected to be found through research, this presupposes the existence of a system whose users, mode of operation and context of use closely resemble the one to be created and whose performance has somehow been found lacking and hence the new "performance requirements". In other words, User-Centered Design methodology seems to be oriented towards the *improvement* of existing systems rather than the creation of novel ones, since novel systems may not only allow users to carry out previously impossible-to-attain goals, but also dramatically change how and in which context users carry out their current goals by offering new and different possibilities to act.

In addition to the problem posed by the analysis phase, traditional evaluation of User-Centered designs focuses on "usability testing". Indeed, the Usability Professionals Association not only suggests this as a method of evaluation to ultimately check the design against the objectives but also in between design iterations. Several authors have discussed the problem posed by traditional usability evaluation when dealing with novel designs. The following is a summary paraphrasing the main points found in the literature:

- Premature usability evaluation of early designs can eliminate promising ideas by focusing on what the problems are (usability bugs) and not on the benefits (usefulness). In early designs usefulness may overshadow usability, usability should be allowed to eventually come through gradual refinements in later stages (Greenberg, 2008).
- Traditional usability evaluation of inventions does not provide meaningful information about their cultural adoption over time. By restricting usability evaluation to a set of tasks and contexts, the experimenter may be missing the tasks and contexts for which the system may prove more useful (Greenberg, 2008).
- Complex systems generally do not yield to simple controlled experimentation. By only creating what a usability test can measure, we will produce only simplistic progress that is not necessarily meaningful (Olsen, 2007).

However, there are important advantages to a design process with the user at its center. User-centered design, with its focus on **the user, his goals and his circumstances** (the rhetorical situation), drastically increases the likelihood that a product or system will be accepted-by and acceptable-to users by including these premises as requirements in the design process; without it, the risk of investing budget and time developing non-viable innovations runs high. Designing according to Human Factors and with knowledge of the interaction affordances of new platforms, although necessary, is not a sufficient condition for user acceptance and acceptability, as illustrated by what is informally called "inventions ahead of their time"—products and systems that are correctly designed from the Human Factors and technological perspective, hence their later success, but were inadequate for users and their circumstances at the time—and the large number of tools and machines that, although once useful, were later discarded because they no longer fitted any human purpose.

In conclusion, although the traditional User-Centered Design methodology does not favor the creation of innovation, its philosophy to keep the user in focus provides clear advantages with respect to user acceptance and acceptability. Therefore we suggest there is a need for the formalization of a new methodology that specifically deals with innovation while maintaining the user at the center of the design process. This project explores a new way to keep the focus on the user during the design of novel systems by maintaining the user-centered philosophy while revisiting the methodology.

2. Structure of the report

The present report is divided in two parts that correspond to the two goals that direct this project.

⁴ Petamedia IRP, Deliv6.1-150306-TUD, page 10.

⁵ http://www.usabilityprofessionals.org/usability_resources/about_usability/what_is_ucd.html

- Part 1, **Near2Me: a Personalized Recommender and Explorer for Off-the-beaten-track Travel Destinations**, describes the Near2me concept and its evaluation. The need for such a system is outlined, the concept itself is described with a mention to how it differs from existing travel destination recommenders and explorers and its alignment with user motivations. The abstracted the data structure used in Near2me in order to connect users to recommended items is described, placing it in the context of related work, and suggesting other uses the same data structure may serve. The PetaMedia-developed algorithms that compose the design are enumerated, and it is explained how they connect to produce the desired outcome. The concept's evaluation and its results are presented. And finally, recommendations for the evaluation of a working prototype in the next design iteration are provided.
- Part 2, **Putting the User at the Center of the Design Process of Technology-driven Novel Applications**, describes the design and evaluation methodologies developed and put into practice during this project in order to keep the user at the centre of technology-driven design.

Part 1- Near2Me: a Personalized Recommender and Explorer for Off-the-beaten-track Travel Destinations

ABSTRACT

This part of the report describes the design and evaluation of Near2me. Near2me is a personalized recommender of travel destinations for people who like to travel outside mainstream tourism. Near2me is targeted at users who would like to discover a destination in an ambiance of authenticity, focusing on the intersection of their own interests and the identity of the place rather than on experiences manufactured by the mass tourism industry.

Near2me takes its cues for recommendation from its users' media capturing behavior, obtaining a pool of spontaneous information uncontaminated by ulterior motivations, which reflects destinations purely in terms of their special character, offering of experiences, and aesthetic value. Its recommendation algorithm uses media capturing-related situational knowledge to produce recommendations based on the users' interests rather than general popularity. Functionality is added to allow the user to (1) explore and evaluate these recommendations to decide if they are worth following, (2) control recommendation direction and follow leads within recommendations by applying filters to recommended items when browsing, (3) link recommendation to external information to obtain background and practical information, which is not provided by Near2me, and (4) save and use selected recommendations while visiting a region. Near2me's functionality and features are all powered by PetaMedia-developed technologies.

Background research was done identifying that *search of an authentic exchange, away from 'manufactured' tourism experiences* and *search of a more personal interaction with the city, an experience that arises from the unique combination of the personality of the visitor and the place* are indeed travelers' motivations to engage in creative tourism. A study of the current landscape in travel recommenders revealed that there is no application that can assist travelers who wish to engage in personalized non-mainstream tourism.

Finally, an evaluation with users of the Near2me concept was conducted using a video prototype, resulting in the validation of the concept. Participants understood the concept and how they could benefit from it in the circumstances Near2me is built to support, they identified the functionality and linked it to realistic past or plausible future scenarios of their own experience and mentioned they would adopt Near2me. Participant's understanding and immediate familiarity with the concept was evidenced by their ability to appropriate the concept and suggest new uses and features, providing suggestions for development in future iterations.

Guidelines for the evaluation of a working prototype, being developed for the next iteration, are also provided.

1. BACKGROUND

According to a 2006 report of the European Travel Commission, titled *Tourism Trends for Europe*⁶, current tourism trends include **increasing demand for independent holidays**, with a relative falling demand for traditional package holidays, and a change in the relationship between host and guest, with **people seeking genuine experiences rather than staged ones**. Tourists try and achieve "deeper" and more "meaningful" experiences by changing their role – engaging in "volunteer tourism" or "creative tourism". UNESCO (2006) defines Creative Tourism as "travel directed toward an **engaged and authentic experience, with participative learning in the arts, heritage, or special character of a place**, and it provides a **connection with those who reside in this place and create this living culture**," as opposed to "beach tourism, in which people come to a place for relaxation and leisure;" and "cultural tourism, oriented toward museums and cultural tours."

These trends represent not only a real, but also a desirable development. UNESCO (2006) and the European Travel Commission (2006) have identified the following benefits that Creative Tourism has for society. Good for the economy: spreads the benefits of tourism distributing wealth; and works towards evening out occupancy in seasonal destinations, improving infrastructure efficiency and tourist sector employment stability. Good for culture: prevents "touristification" and provides an incentive for the preservation of local cultures, and improves travelers' culture. Good for society: connecting travelers with locals augments travelers' and hosts' exposure to different cultures and their mutual understanding, creating a more tolerant society.

⁶ <http://www.etc-corporate.org/resources/uploads/>

ETC_Tourism_Trends_for_Europe_09-2006_ENG.pdf

As a result, it makes sense, from individual and collective points of view, to develop technologies that assist the pursuit of Creative Tourism.

Since 2004, CouchSurfing, a non-profit organization that operates through the website of the same name⁷, allows travelers to meet locals wherever they travel, effectively allowing them to connect with the living culture of the place. Although the organization is mostly known for allowing its members to exchange accommodation by hosting or being hosted by other members, only 30% of its member base offers a place to stay⁸ and purely social exchanges constitute an important part of its network's activities. Its authenticity is based on the self-motivation and spontaneity of its membership.

However, no analogue alternative allows travelers to explore the special character of a place and discover meaningful experiences they may want to engage in, in a similarly authentic and independent manner. The difference between connecting travelers with local people and connecting them with the special character and meaningful experiences of a place—and the main challenge to port CouchSurfing's success to its sister trends in independent travel—lies in the fact that while local people can reach out to contact travelers (i.e. by uploading a profile to CouchSurfing) while the locations and aspects that constitute the special character and meaningful experiences of a place cannot. The question is how to find and present to travelers a list of aspects representative of the spirit of a place and potentially meaningful experiences, without falling into the “artificiality” characteristic of the “staged” experiences that independent travelers want to avoid. More precisely, the authenticity of the experience should be preserved in both **the selection of items and their method of presentation**.

The recommendation of “authentic” items is not trivial. Based on interviews with travelers that frequently engage in Creative Tourism⁹, we define as “authentic” that which may or may not be a commercial product, but (1) has not been intentionally exploited as a tourist attraction to the point that it has become massively done/visited by tourists, (2) is representative of the heritage or the special character of the place. This definition presents three challenges. In the first place, where exactly the threshold lies to determine if something is or is not a massively visited tourist attraction, or something does or does not represent the special character of a place, is a subjective matter that varies from traveler to traveler. And secondly, the definition implies that the items to be recommended are not popular with the average tourist, that is, they lie in the long tail. Recommending from the long tail is a challenge because the probability of a successful random recommendation is very low, hence it requires heavy and very accurate personalization that has to be achieved while allowing for discovery and novelty, which is a general requirement of recommenders. Finally, the effect of travelers' behavior arising from recommendations needs to be taken into account, to avoid further “touristification” that would otherwise result in a non-sustainable system. In addition, the recommendations have to be presented in a way that allows the user of the application to gain knowledge about the recommended items and evaluate them from an independent perspective, without feeling being led on in an artificial manner.

The *Triple Synergy*, the synergetic combination of the information extracted from user-generated annotations, media content analysis and social networks structure, which constitutes the area of study of PetaMedia, presents several opportunities that suggest its application to the creation of an application with the above-mentioned characteristics. To begin with, the activity of traveling; documenting the trip through media; annotating the media in order to organize it, describe its content and add context; and finally sharing this information is widespread. As a consequence, **media and its annotations hold a considerable volume and diversity (in content and point of view) of information that describes places**, which could be used to produce and evaluate recommendations outside the mainstream. Moreover, we conducted a study on 500 Flickr pictures tagged with city names—a typical indicator of travel—belonging to 16 users and depicting 3 different cities, in which we found that travelers, even those who do not engage in Creative Tourism, **document not only famous landmarks, but also anonymous off-the-beaten-track objects and locations they found interesting**, which they have presumably discovered by chance while exploring the region. These pictures were qualitatively coded, their subjects resulting in three (non-exclusive) clusters: descriptive of the nature of the place, connected to some interest or activity that the author relates to, and aesthetic value. These findings suggest that **travelers' pictures are a valuable means to obtain information about non-massively touristic locations and experiences that are representative of a place**. But most importantly, this information has been produced by a group of individuals whose aggregated behavior has imprinted it with patterns related to the portrayal of different places by the same individual and the portrayal of each particular place by different individuals, and the situational context around the creation of each piece of media. These **patterns in media capture may be able to provide a clue to relate travelers to the right new experiences and locations**.

⁷ <http://www.couchsurfing.org>

⁸ <http://www.couchsurfing.org/statistics.html>

⁹ This study is described in Section 2. **PRELIMINARY USER RESEARCH**.

2. PRELIMINARY USER RESEARCH

In 1992, Lee and Crompton published a study on the motivations of tourists signaling novelty seeking as the main drive for leisure travel. However, although some of the conclusions of this study are very general and based on human cognitive psychology, having been published before the current Creative Tourism trend, others cannot be taken for granted.

In 2003, acknowledging that little research had been done towards understanding how tourists choose and arrange various activities, Brown and Chalmers conducted an extensive ethnographic study on tourists. However, since they conducted their study around major tourist areas, the focus of their study were, again, practitioners of traditional tourism. Also, unlike Lee and Crompton (1992) they focused their study on the mechanics of tourism rather than the motivations behind their participants' behavior.

To complement the relevant part of Lee and Crompton's conclusions, we conducted our own study with participants known to engage in Creative Tourism, with the aim of finding the motivations behind this activity.

An online questionnaire was distributed to eight people who enjoy and often engage in Creative Tourism. They were asked to describe their Creative Tourism trips, to mention what they liked about them and to state what they sought when engaging in this kind of travel. No questions were asked about how they plan their travel and activities, maintaining the focus on motivation alone. The focus on motivation, rather than task analysis, was chosen with the intention to liberate the design of a future application based on emerging technology from possible artifacts of current technology limitations.

The following motivations could be identified. They are listed below, accompanied of a sample of supportive quotations.

- **Search of an authentic exchange, away from "manufactured" tourism experiences.** "I had the best experience when I trusted the people who offered me a place to stay or go to, because it is unexpected and more genuine. If you go where the book recommends, you will meet people that have met thousands of tourists like you and they are not as eager to share experiences", "I feel I have seen something that few people have, because it was more genuine and not mass-tourism like", "I even feel a bit disgusted when I'm in the middle of a tourist crowd and everyone is doing the same thing, taking the same pictures, moving too fast to grasp the essence of the place".
- **Search of a more personal interaction with the city, an experience that arises from the unique combination of the personality of the visitor and the place.** "I also enjoyed getting to know my friend's friend, and being part of their daily lives (for example, we helped out one of them with some silk-print for his graduation in graphic design, or I went to the basketball match to support an other guy for an important match" [from a graphic designer that is interested in sports], "We headed off to one of the friends' house where he has a workshop (he's a carpenter) in the ground floor (that's also where they build rockets, fix firearms, play around with electronics and so on), and where he actually lives", "The most surprising thing was discovering that there was a different way to ride a horse. They did the same thing but differently, I had never even thought there was something particular about how to do it" [about riding a horse in the Middle East vs. in Argentina], "I remember with pleasure a long cigarette break/chat with the security man of the Modern Art Museum".
- **Immersion in place's culture, through interaction with the locals, as opposed to observation.** "Seeing and (sort of) experiencing the life of a modern farmer living in a spectacular place like Hardanger, and fantasizing about what would living a life like that would be like", "What I mainly like about visiting friends is that you can better catch the flavor of some places, and you feel cozy and almost at home. In this case, I really enjoyed to be introduced to my friend's friends and family, and be treated as their acquaintance, as if I knew them already", "The main goal of this trip was to experiment how people live in these places", "I also like to go to countries where I am welcomed by a local family and simply live like them, go to the market place and cook with them for instance."

These motivations were used as a guide to direct the design presented in this work. These are the high level motivations behind the way people who engage in Creative Tourism decide, arrange and behave during their travels. Although a new application may bring about radical change in the manner people do these things, respecting and/or providing better means to fulfill these motivations will be the key to its success.

3. CURRENT LANDSCAPE OF TRAVEL RECOMMENDERS AND EXPLORERS

We performed a survey and critique of over 50 different systems that allow users to get recommendations, explore and discover places, activities or items related to travel. These applications were collected from recent literature and the web, to reflect the current variety in this domain. We clustered them according to the goals they allow the user to achieve, the type of interaction they require, and the nature and structure of the information available for recommendation and exploration as a consequence of system design. We obtained five clusters, defined as follows.

1. **Classic travel planners** are applications that offer the user personalized recommendations from a pool of known tourist destinations, travel routes, tourist attractions, and/or services (e.g. restaurants, hotels, transportation); based on the user's interests, situational context, and/or practical preferences (e.g. time, price, accessibility). The data used for

recommendations usually comes from tourist boards, travel guides, or travel agencies; and is often organized according to an ontology.

These applications are aimed at users that otherwise would have used a travel guide, browsed through Tourist Board brochures or contacted a travel agent (the sources of the data used for recommendation). The benefits these systems provide are centered around how much faster they are and effort reduction (need to evaluate less options), less likelihood to miss popular or must-see items (as they are expected to become more prominent through expert selection or collective reinforcement), and real or perceived choice safety (e.g. taste neighbors also chose this, many people rated this as good, perception of the choice as being endorsed by the recommender, alleged similarity an item the user previously liked).

In terms of interaction, users' interests are input explicitly (building a profile, or giving feedback before or after putting into practice a recommendation), or implicitly through the interaction of the user with the system (for example, when the user saves, books or prints one particular recommendation). Creating a profile demands some effort and time on the part of the user (it could be argued that this effort and time are minimal compared to those saved by a frequent user of the system, but it does introduce a barrier to adoption). And, more importantly, preference self-portrayals suffer from a large number of shortcomings (as extensively described by Pu and Chen, 2008), notably that "users are not aware of all preferences until they see them violated", which in the case of preferences as fuzzy as "interests" is hard to evaluate before the user puts the recommendation into practice and for travel can result in mistakes costly in time and money. Implicit feedback provided by the user's interaction with the system and explicit feedback prior to the implementation of a recommendation both present the same problem: the user is evaluating an artifact of the system and not the actual object of the recommendation. Finally, feedback provided after the user has implemented the recommendation can be considered more relevant but can be hard to legitimate in cases where the user has not engaged in a concrete transaction (e.g. the purchase of a service or transportation to a particular destination) and is subjected to bias as it comes from a self selected population.

Also, as a consequence of the above-described characteristics, these systems do not provide the user the opportunity to discover places they would not have found through traditional means, and are prone to several complications depending on the nature of the recommending algorithms used (Hinze and Junmanee, 2006):

- New user problem, when lack of sufficient information on a new user hinders the ability of the system to produce a good recommendation.
- Sparse feedback scores or cold start problem, when the items available for recommendation do not have enough associated user feedback for the system to recommend it to users or declare its neighborhood to other items.
- Users with particular preferences, meaning users who have few taste neighbors or users who have liked disparate items, typically do not get good recommendations.
- Overspecialization, when the recommender systematically recommends items similar to those that have been highly scored by the user or tight communities of neighbors are created. As a result, discovering diverse items or preferences that are not explicitly known by the user is not possible.
- Need for the users to understand the recommendation model in order to modify their preferences as a step to get more accurate recommendations (in the case in which recommendations come from direct elicitation of preferences).
- Difficulty to obtain feedback from final user satisfaction, which usually results in the measurement of the user's satisfaction with the recommendation rather than with its object.
- Scalability.

Examples of classic travel planners are: TSIPPS (Wu, 2008), INTRIGUE (Ardissono, 2003), SPETA (Garcia-Crespo, 2008), The Travel Channel's¹⁰, TripAdvisor's¹¹, and TripSay¹².

2. **Web-based travelogues** are the online version of the travel literature (or *travelogue*) genre. They consist on the traditional account of a trip, generally exhibiting a coherent narrative or aesthetic beyond the logging of dates and events, and in this case also enhanced by photographs, interactive maps, links to relevant material and other types of media afforded by modern technologies. Usually travelogues are put up by amateur expert travelers and depict experiences more intense, engaged and authentic than those typically found in mainstream tourism. Travelogues vary in the kind and detail of information they provide, some focus solely on the description of the places and the people encountered along the way, while other offer more practical information (e.g. camping sites, restaurants, directions to places).

¹⁰ http://www.travelchannel.com/Places_Trips

¹¹ <http://www.tripadvisor.com/Inspiration>

¹² <http://www.tripsay.com>

By browsing travelogues, one can explore travel destinations from the point of view of the author, and get hints on what to do or where to go. Travelogues are a useful means to explore places for people who are interested in getting a general idea of what a place can offer. Also, because of their content, travelogues are especially attractive for people looking for more intense, engaged and authentic experiences than those typically found in mainstream tourism.

The interaction with travelogues is usually restricted to browsing their content. Travelogues are usually scattered through the Internet, and considerable part of the interaction consists of the user searching and finding them. Search engines can help finding something specific, and some directories group some travelogues by categories and locations.

Because of their nature, travelogues offer no personalization of the content they display, there is no guarantee that the author's point of view coincides with the reader's, and, as they are the account of one particular trip by one particular traveler, they are usually limited in breadth.

Hao et al. (2010) developed a method that mines travelogue content to (1) recommend destinations based on flexible queries, (2) provide a characteristic summary for a destination, with tags and snippets, and (3) identify the informative parts of a travelogue and enrich them with related images. These would partly address the limitations mentioned above. The recommendation based on flexible queries, which the authors describe as the search of a location similar to previously visited locations that the user liked and locations relevant to an intention (e.g. hiking), could amount to some form of personalization. However, after the information in the travelogues has been mined, the nature of the recommendation method (based on comparing keywords found in travelogues looking for semantic similarity) is not different from the methods underlying the classic travel planners, hence it suffers from the same limitations. Also, it depends on direct user input for the search query, which limits the recommendations to what can be anticipated by the user. The aggregated summaries and the image enrichment both address the unique-point-of-view and breadth problems: the aggregation from various travelogues multiplies the points of view and increases breadth of experiences depicted, and the images illustrate different perspectives. How much these techniques can contribute to actually solve these problems is uncertain, since they have not been well developed yet.

In any case, the information the user can receive from travelogues will always depend on the content other user have included in their travelogues, which poses two inconveniences. In the first place, travelogues are usually highly personal and, as mentioned, tend to present an aesthetic beyond the logging of dates and events, which means that travelogue content strongly reflects on the author. As a consequence, authors may (consciously or not) censor/doctor their content to show an embellished image of themselves. Secondly, it takes hard work to document a trip in a travelogue; therefore, there is only a tiny and self-selected population of travelers that does it. As a result, the available content is limited in size and style.

Travelogues.net¹³ is an example of a travelogue, more examples can be found in Yahoo's¹⁴ and Google's¹⁵ directories.

3. **Location-based annotators** are applications that allow users to create on-location multimedia annotations that remain linked to that particular location. These annotations can later be discovered by other users (on location or usually also as a result of a location, social or keyword-based search), who can benefit from the information they contain.

Location-based annotators are similar to travelogues in the sense that they contain places that other users have visited and their (albeit brief) descriptions. Although the information each annotation contains is too limited to support the exploration of a place, the aggregation of annotations from multiple users (that results from a location-based search, for example) can help fulfill this goal. Collections of location-based annotations are a social phenomenon so, as opposed to travelogues, they lack the bias towards non-mainstream locations and activities. Because the creation of a collection of annotations for a given place is usually the sum of individual contributions and not a coordinated effort, the result is usually less detailed and sensemaking than in travelogues but shows many different points of view in one go. Therefore, location-based annotators are useful to explore places for people who want to have relatively shallow but broad descriptions of a place in a fast, effortless way. Also, because annotations usually consist on small units of information, a location+keyword(s) search can produce relevant and concise tips or reviews with minimum effort.

Regarding interaction, location-based annotators allow users to obtain information in three different ways: as on-site alerts, through a (social, location or keyword-based) search (which can be carried out on-site or remotely), or by browsing the complete collection of information for a region (e.g. a city or a neighborhood). On-site alerts imply that the user already knows and is visiting the place, in which case the information he/she receives is meant to enhance the experience

¹³ <http://www.travelogues.net>

¹⁴ <http://dir.yahoo.com/recreation/travel/travelogues>

¹⁵ <http://www.google.com/Top/Sports/Cycling/Travel/Travelogues>

or suggest activities or aspects of the place that the user may want to explore. Usually, this information can be personalized in the sense that it was produced by the user's friends or is related to a keyword to which the user subscribed. On the other hand, search relies on the explicit input of the user, which again limits the results to what can be anticipated by the user. Finally, browsing, usually aided by tag filtering, allows the user to get a more general outlook but takes more time and the lack of personalization can be frustrating.

Some features of their design induce certain particularities in location-based annotators that affect recommendation (in the form of alerts) and exploration. Most location-based annotators publically display a user's profile, detailing all the places where the user has left annotations and the annotations themselves (these annotations can be video, photos, drawings, text, links, tags, a combination of these, or simply check-ins). Also, most location-based annotators allow users to share their annotations through other social media where they reach their friends or followers directly. Finally, in all the location-based annotators reviewed during the course of this survey, annotations were explicitly intended for people other than their author. As a consequence of all this, users are conscious that they are broadcasting a message and that this message will reflect on them, both of which can lead to annotation subjectivity due to users' ulterior motivations or social desirability bias. Moreover, in the most popular location-based annotators the main incentive for contribution seems to be the existence of rewards (reputation or actual prizes) for people who annotate, review or check-in in particular places, according to a particular pattern, or a certain number of times. These incentives further distort the users' motivations for annotation, creating user-generated annotation spam. And finally, collections of annotations related to a place are the sum of individual contributions that take place in the blind of others' and not a concerted effort, and annotations are usually short and simple (because they are created on the move). As a consequence of the combination stated above, annotations that do not arise spontaneously and are meant to be an extension of a check-in (with the purpose to let an audience know or record in the user's profile that the user has been there) tend to state the obvious or the most prominent and do not show great variety. This behavior is similar to that of tags, which follow a power law, only that users tend to add less annotations than tags making the effect more pronounced; and tags consist of more complicated and diverse forms, which makes grouping under concepts, to be able to show a number of them that displays variety, more difficult.

Examples of location-based annotators are: A Location is Worth a Thousand Experiences (El Ali, 2010), HyCon (Hansen, 2004), Foursquare¹⁶, and Gowalla¹⁷.

4. **Recommenders based on location tracking** are services that track users locations automatically (e.g. using Cell-ID, Cellular Positioning, or GPS) and use the data from their community of users to derive patterns that allow them to recommend new items. The recommenders based on location tracking encountered during this survey also allow users to make text/tag annotations on-location, although their creators mention that their test participants made very seldom use of this feature (Zheng, 2010).

Recommenders based on location tracking use user-generated data (annotations) and known geographical coordinates-place pairs to derive knowledge about location-activity, activity-location, location-location, and user-user correlations. These correlations allow these recommenders to provide users with the following kinds of recommendations: recommending an activity given a location-query and recommending a location given an activity-query, recommending a location given the user's location history and recommending taste neighbors. Also, users' individual location histories can be used to personalize location-activity and activity-location recommendations predicting the user's interest on one location/activity through its correlation to the user's previous locations or how popular the location/activity is with his/her neighbors.

In terms of interaction, users can access recommendations by different means: search queries, spontaneous suggestions by the system based on the user's current location or activity, or browsing of a personalized pre-selection of items. What is more interesting in this case, is the way users input information about the environment and about themselves: almost completely implicitly. Because location is tracked automatically, and because it is in turn used to provide recommendations to the user, users have less both opportunity and incentive to let ulterior motivations and social desirability bias take over their input. Also, input is effortless: users just need to carry a location-aware device, which does not need to be obtrusive (i.e. can be wearable or embedded in a mobile phone). Because of the nature of the input method described above, motivations other than the true motivations people have to visit a place or perform an activity are not introduced artificially by the system. As a consequence, the data in the system can be expected to be more genuine than the data in explicit-input location-based annotators. Also, because input is effortless, a larger and more diverse

¹⁶ <http://foursquare.com>

¹⁷ <http://gowalla.com>

population can be expected to contribute to the system (compared to travelogue authors, or even location-based annotators, for example).

Other design characteristics of recommenders based on location tracking also influence the kind of recommendations it can provide. In the first place, this kind of systems bases its recommendation algorithms in data that is not human-friendly: latitude and longitude coordinates. Each time a location recommendation is produced, it has the form of a pair of geographical coordinates. In the surveyed applications, work is still ongoing and for now it is not clear how and what information is presented to users to be able to evaluate the suggestions the system offers; and, in any case, this remains a not straightforward challenge and a central point to the success of the system. Secondly, the number of people who provide manual annotations for the locations they visit is negligible. As a consequence, it is very unlikely that the system will perform adequately for locations in the long tail, for which external geographical coordinates-place information pairs cannot be found in external sources (e.g. in gazetteers, Wikipedia). Finally, although the surveyed systems offer a way to address this issue, it remains unclear how effectively it is possible to separate the locations where people stay because they are enjoying themselves from the locations where people stay because they have to or they have no choice, without more feedback than the stay's duration.

All the surveyed examples of recommenders based on location tracking were developed by Microsoft Research Asia (Zheng, 2009a and 2009b; and Zheng, 2010) as a part of the Microsoft GeoLife project¹⁸.

5. **Photography-based explorers** are applications that allow users to browse through images of locations in order to get a feeling of what these locations are like. Photo-based explorers arrange images according to the location where they were taken, using information from geotags. The photographs are supplied by users and, usually, also carry user-generated annotations.

Photo-based explorers are useful for people who want to explore places, without regard for practical considerations such as price, travel time, opening times and sometimes even point of interest names. Photo-based explorers help their users find out if a place appeals to them and what is there to see and to do by visual inspection and using their own judgment. In this sense, photo-based explorers are not very different from what the user would experience by taking a real "sample" tour of the location. Tags in images enhance this experience, allowing the user to filter by topic or providing extra information on what the user is seeing. Thanks to tags, and initiating the exploration based on a topic rather than a geographical location, photo-based explorers also allow users to discover where they can find or do something that interests them, later being able to evaluate the options by looking at pictures of that instance corresponding to different places.

Photo-based explorers display the geolocated images in a format that allows the user to browse through them and carry out simple searches and filtering, i.e. a map or a hierarchical tree of locations, typically accompanied by a search box and/or a tag collection, both of which can be used to refine according to tags and location (usually search only consists of tag-based indexing and some ranking algorithm). Sometimes only a selection of images is shown at each level of detail according to different criteria: for example, diversification or collective image ranking.

In photo-based explorers, there is no personalization beyond what can result from semi-permanent tag-based filtering. Also, because photo-based explorers are image-based and depend on user-generated tags to add content and descriptions, browsing can result in a dead end that occurs when the user wants to know more about a location depicted in an image but does not have any textual relevant information (notably, the place name) to proceed with the exploration by other means. On the other hand, in all the reviewed photo-based explorers, the photographs had not been intentionally supplied by the users with the goal of contributing to the application: all the photographs came from online image repositories where users upload their photographs in order to store or share them, giving permission for public access. As a consequence, there is no bias introduced by the system and photographs can be assumed to be taken with the original motivation of documenting the trip, reflecting the locations and situations that aroused the traveler's interest, for remembrance and/or social conversation through sharing. In addition, travel photography is universal, and online storing and sharing is a popular activity, resulting in the largest and, in terms of points of view, the most complete and diverse source of information about places (compared to the other systems found during the survey). Also, images have a very high descriptive power in relationship with the effort the user has to invest in creating them (compared to location tracking, which represents zero effort but is much less descriptive, and location-based annotations, which would require much more effort to match the descriptive power of an image). And finally, although a traveler chooses which aspects to portray through photography and there is some editing effort when not all images are uploaded for public consumption, photographs have an aura of authenticity, and they are perceived by viewers to be less subjective than textual descriptions

¹⁸ <http://research.microsoft.com/en-us/projects/geolife>

in the sense that they may not be reflecting the whole truth but they are undoubtedly portraying something that indeed exists or occurred.

Examples of photography-based explorers are: World Explorer (Ahern, 2007), World Wide Media Exchange (Toyama, 2003), Flickr Places¹⁹ and Panoramio²⁰.

6. Recommenders based on trends sourced from photography are systems that examine travelers' collective behaviors related to photo capturing, and the images and metadata in photos, and analyze and use this information in order to find popular places of interest (i.e. the places that people photograph the most or the places that experience certain photography patterns). These systems are capable of using the wisdom of the crowd to discover and identify new and interesting landmarks and the images that best portray them, as long as collective behavior shows they are popular with the majority of people.

In a sense, these recommenders are very similar to photography-based explorers (i.e. they also offer no personalization, are based on image browsing and obtain their information from photographs and their geotags), it is the subsequent analysis they make of the images that sets them apart. By identifying the landmarks they can be either restricted or focused solely on items that are more likely to interest the user (since they have been selected because of having this effect in large numbers of other people). And by identifying the best or more informative images from each landmark, they can also increase the probability that the landmark is shown in a manner that is most appealing or descriptive to the user.

These recommenders are targeted at users similar to those of photography-based explorers (who want to explore places, without regard for practical considerations such as price, travel time, opening times and sometimes even point of interest names), but who want to have a more limited, typical and iconic view of a place, composed by iconic landmarks and their canonical views, to the detriment of variety, little known aspects, and more original points of view.

The surveyed systems did not emphasize the way of interaction, which due to content-type similarity can be assumed to afford any functionality that has been developed or suggested for photography-based explorers.

In the same vein, data acquisition results from similar procedures as in photography-based explorers, although the authenticity, spontaneity and multiplicity of points of view are washed out by the subsequent analysis and sanitization of what is presented to the user. The descriptive power of images over words, however, is still an advantage of this kind of systems being photography-based.

Examples of recommenders based on trends sourced from photography are: Tour the World (Zheng, 2009) and Mapping the World's Photos (Crandall, 2009).

4. OBJECTIVE AND REQUIREMENTS

The goal is to create an application capable of assisting travelers who engage in Creative Tourism in exploring and discovering experiences, objects and locations that authentically represent the spirit of a place and match their personal interests.

This application should:

1. Present the user with objects, places and experiences that “authentically” represent the spirit of places, which means objects, places and experiences that, being representative of the culture and heritage of the place, are not the object of mass tourism popularity.
2. Adjust the threshold it uses to determine where to draw the line between “touristy” and “authentic”, according to each user's criteria.
3. Personalize the selection of items according to each user's interests to avoid recommending items that, although authentic and representative of the place, belong to niches that do not resonate with the user.
4. Supply enough “authentic” descriptive information pertaining to each recommendation to allow the user to evaluate it and decide if he/she wants to pursue it. This is particularly important given that recommended items are outside the mainstream.

A quick look at the current landscapes of recommenders shows that there is yet no recommender capable of satisfying the above-mentioned requirements.

¹⁹ <http://www.flickr.com/places>

²⁰ <http://www.panoramio.com>

Also, on a different note, as this project was developed within PetaMedia²¹, it is subject to practical restrictions that originate from the needs to (1) adapt its workload to the available resources, and (2) test with users the component technologies, which are part of PetaMedia's ongoing research. The application was designed under the following constraints:

- The application should be restricted to the use of PetaMedia-developed technologies, without the possibility to develop any components from new technologies or technologies existing outside PetaMedia.
- The design should also serve as user-testing platform for the Petamedia technologies: the different technologies that compose it should be decoupled in different features so the feedback concerning each of them can be singled out during the PetaMedia Field Trials (trials that will be performed in a subsequent iteration, with the use of a working prototype).
- The design should result in a prototype that is/includes a web application that can be developed by a fulltime programmer in 4 months.

5. NEAR2ME: THE CONCEPT

5.1 Basic Description

Near2me is a personalized recommender of travel destinations for people who like to travel outside mainstream tourism. Near2me is targeted at users who would like to discover a destination in an ambiance of authenticity, focusing on the intersection of their own interests and the identity of the place rather than on experiences manufactured by the mass tourism industry. Near2me is focused on allowing the user to get the feeling of what the spirit of each place is, what is there to see and what is there to do, rather than providing detailed background information on the culture and heritage or practical information to assist the planning of a trip: the goal of Near2me is to give the user a glimpse of what he/she would experience as a traveler by being at the place.

Near2me looks at the geotags present in the metadata of the photographs users have taken while traveling. By looking for co-occurrences of geographical-coordinates within individuals, Near2me detects places that are related based solely on the statistical correlation between a user liking one of these places (i.e. documenting it with photographs) and liking the other one as well. In this manner, Near2me can recommend to users that are known to have visited and liked one particular location (because they took photographs at this location and submitted them to the system) another location that they may like based on the statistical correlation described above. This is the basis of Near2me's personalization. The photographs of a place, as well as their user-generated annotations, are shown to users in order to allow them to explore and evaluate these recommendations. User-generated annotations that describe the images (i.e. tags added by the authors of the photographs) and the authors of the photographs themselves act as pivots for the user to direct the recommendations in one particular direction: choosing a topic or aspect from the tags in order to filter recommendations towards these subjects, or choosing another traveler in order to see the current recommendation through his/her perspective. Finally, as the recommendations are predominantly expressed through imagery, an image-based search, which retrieves textual information and external links to related topics, is offered for the user to be able to connect what he/she is seeing in Near2me to external sources of information on the Internet and satisfy the demand for background knowledge and practical information that is out of the scope of Near2me.

Concerning the interaction during the recommendation exploration, Near2me focuses on browsing, as opposed to search, browsing being better suited for discovery and the type of ill-structured quest on which Near2me aims to assist its users.

However, the most consequential characteristics of Near2me's interaction paradigm correspond to the manner in which its users contribute, at once, the information that will be used as a source for recommendations and information about themselves that will be used to target these recommendations. The information from which the recommendations are sourced comes from the geotags in the photographs that users have collectively taken. As a consequence, this information is all implicitly generated, drawing its authenticity from the spontaneity of photograph-capturing on the part of travelers: the information from where recommendations are sourced presents no ulterior motivations apart from the ones behind travel photography. Moreover, given that the same information contributed by one user to the pool from where recommendations are sourced is also used to target recommendations to that user, users are discouraged to tamper with the information they submit to the system. Also, information sourcing is effortless, since photography, tagging and sharing are a spontaneous behavior for many travelers.

As it will be further explained in the next section (sub-subsection 6.2.1. *The Recommender Algorithm*), the way the recommender algorithm within Near2me is designed allows it to produce recommendations not only sourced from authentic information and based on personal history, but also with a focus on personalization rather than popularity, favoring recommendations outside the mainstream with an emphasis on the user's personal interests.

²¹ <http://www.petamedia.eu>

5.2 Component Parts

5.2.1 The Recommendation Algorithm

Near2me exploits users media capturing as a signal of interest aroused by an object, place or experience. When people travel, they take pictures at locations where they find something descriptive of the nature of the place, something connected to some interest or activity that they relate to, something that has aesthetic value, or a combination of these. Travelers, even those who do not engage in Creative Tourism, document not only famous landmarks, but also anonymous off-the-beaten-track objects and locations they found interesting, which they have presumably discovered by chance while exploring the region²². As an increasing number of cameras and camera-phones come equipped with location-aware technology, these locations become available in the form of geotags embedded in the image metadata. Many of the travelers, later upload their photos to an image sharing website and annotate them in order to organize them, describe them and add context (Ames, 2007). As a result, a collection of locations *each traveler* found interesting, **including off-the-beaten-track locations** that do not correspond to any known landmarks, together with images that portray these locations and annotations that describe and explain what he/she found interesting, becomes accessible. As opposed to photography-based explorers and recommenders based on trends sourced from photography, which aggregate the images from many travelers in order to present a general overview of a place, Near2me keeps the information connected to the traveler that created it, in order to enable personalization of content based on individual preferences.

Near2me uses the information in the images' geotags to find "places" by geographically clustering these tags, and finding the center and delimiting the expanse of the region enclosing a single item of interest. With knowledge of these "places" Near2me looks for co-occurrence of place pairs in each traveler's histories. When the conditional probability that travelers who visited *place A* also visited *place B* exceeds the general probability of visiting *place B*, Near2me detects a match between these two places. It is this information about matches between locations that is later used for recommendations (i.e. a traveler who photographed *place A*, who is assumed to have enjoyed this place, which was his motivation for documenting it, will be recommended *place B*).

As places are only considered related when the conditional probability for a pair exceeds the general probability for an individual item, **the relationship between places is privileged over popularity**. Also, by design, because the relationships between places are obtained from co-occurrences in individual travelers' behaviors; rather than through ontological relationships or "similarities" based on coincidence of annotations, classifications or any kind of descriptive keywords; these recommendations exhibit some particular characteristics. In the first place, they can be thematically more diverse and more serendipitous than recommendations arisen from keyword-based similarities or ontological relationships. The first, because they can arise from relationships whose level of abstraction would typically not result in successful recommendations and as a consequence would be discarded; the latter, being simply the result of the fact that they can arise from relationships that are not foreseen. This means that recommendations based on co-occurrence computed through conditional probability are best suited for travelers who are interested in more diverse and serendipitous recommendations, and items that are not necessarily *like* the ones they have visited before. On the other hand, recommendations based on co-occurrence cannot provide on-demand *variety of a kind* (e.g. given a golf course, other golf courses; or, more loosely, given the birthplace of a writer, the houses where a personality has lived), which traditional ontology based and keyword similarity-based methods can do. In recommendations, variety of a kind can be of interest in two ways: items similar to the ones that the user has liked before; or, once a recommendation is made, to see more items in the same vein. Several of the recommenders we have surveyed specifically address the first question already; Near2me is targeted at users who are interested in more serendipitous recommendations, producing recommendations for locations that are similar to the ones that the user has visited previously only when this is present in users behavioral patterns (which has been shown to occur by Clements et al., 2010). Concerning the second use case, when the user gets a recommendation and is interested in exploring more items of the like, Near2me addresses this by using the tags related to the original recommendations (i.e. the tags that annotate the photographs whose geotags constitute the place that is being recommended) to allow the user to guide recommendations towards more objects that have been annotated with those same tags. This mechanism will be explained in detail in *Section 5.1.3*.

Personalization is a vital aspect of recommendation of items when popularity is to be left aside. Lee and Crompton (1992), in their study *Measuring Novelty Seeking in Tourism*, paraphrase Berlyne (1950, 1960, 1966) stating that "extremely novel stimuli may discourage exploration" and "we are indifferent from things that are either too remote from our experience or too familiar." Items in the long tail, in this case consisting on off-the-beaten-track or non-mainstream tourism locations, occupy such a place because they are indeed "*remote*" for the majority of the population. They appeal only to a small niche of the total population appearing bizarre or incomprehensible to rest, which remains indifferent to them. It is because of this that,

²² Based on a study we conducted on 500 Flickr pictures tagged with city names—a typical indicator of travel—belonging to 16 users and depicting 3 different cities.

when recommending items from the long tail, some degree of personalization is essential to make sure that the individual to whom the recommendations are made is not completely foreign to the recommended items. On the other hand, these niche items tend to be very specifically particular—probably this being the cause of their appeal to a small subset of the total population—and, as a consequence the sequential recommendation of several items of the same kind within them would result in the over-familiarity that according to Berlyne also results in indifference. The kind of algorithm at the heart of Near2me, one that allows personalization, but does not result in overspecialization by focusing on ontological relationships or “similarities” based on coincidence of annotations, classifications or any kind of descriptive keywords is ideal for recommendation of non-mainstream items.

5.2.2 Identification and Evaluation

Recommendation evaluation is a central part of a recommender. Once recommendations are produced, they are assessed by the user to determine if they are relevant and worth following. Especially in the case of a recommender that favors diversity and serendipity, which is bound to produce more risky recommendations, it is important to provide the user with the tools to go through a set of recommendations and effectively determine the ones that are a good match. It is also important that the process to do so is seemingly effortless and agreeable, in order to (1) provide an incentive for the user to thoroughly engage in evaluation, to increase satisfaction with the experience resulting from the recommendations, and (2) make the recommendation process an enjoyable activity in itself, to increase satisfaction with the experience of using the recommender.

Near2me produces its recommendations as geographical coordinates, which are not suitable for recommendation evaluation on the part of users. As a consequence, in order to present the recommendations to the user, some information that allows the user to *identify* the place and *evaluate* the relevance of its recommendation has to be added to the interface.

To provide a suitable means for the user to identify the place, Near2me automatically assigns a name to the place. This is done in one of two ways: by soft-matching the recommended place’s geographic coordinates to those of some known landmark (whose geographic coordinates can be found in a public repository, e.g. Wikipedia), with the assistance of keywords found in tags, or, if the previous method does not provide a match, retrieving the name of a region that contains the area of interest from a gazetteer. In this way for example, a group of graffiti close to the Père Lachaise Cemetery, in Paris, will become “*Père Lachaise area*”, in the first case, or “*Paris 20ème arrondissement*”, in the second case. Better-defined landmarks, like the *canal Saint-Martin*, also in Paris, whose geographical coordinates can be found in sources like Wikipedia, will carry their proper names. In all cases, Near2me will also provide the exact point in a map where the exact geographical coordinates of the recommended place lie. These two identification methods, a name and a point in a map, have complementary functions. A name, allows the user to individualize or instantiate the place in his/her mind to better remember it, to engage in conversation about the place with others, and to look for more information about the place on the Internet or elsewhere. The exact location on a map, allows the user to place this location in context within the general region and with respect to other locations of interest, to incorporate the place in an itinerary, and to navigate to the destination.

However, a name is not a proper means for recommendation evaluation, especially in the case where the object of the recommendation is not a defined landmark and, as a result, the name is vague and non descriptive. This case can be expected in recommendations produced by Near2me, given the presence of non-touristic, or off-the-beaten-track, items in the recommendation source pool. Another means has to be found in order to describe to users the recommended place. To portray each place, Near2me uses photographs taken there by other users. The authorship held by the user’s peers guarantees the authenticity of the view that is being presented, in the sense that what the user is seen is what has been experienced in reality by other travelers, like him, who have visited the location and it has not been edited or doctored by third parties like tourist promotion organisms. Also, as the images being shown have multiple authors, the pool of photographs available to describe each place can be expected to be diverse and represent the place from different points of view (e.g. different activities that can be done in the place; different things to be seen; emphasis on different dimensions like aesthetics, shopping, local people, etc.) and in different circumstances (e.g. season, time of day, weather).

Near2me makes sure that the images that are shown to the user to describe a place, and allow him/her to evaluate it, retain the spirit behind the recommendation (i.e. they are *relevant*) and that they are, within this spirit, *diverse* enough for the user to form a comprehensive idea of what is there to see and do and the place. For example, if a travelers who visited the Hunterian museum of the Royal College of Surgeons of England, in London, which houses a collection of anatomical specimens, also visit the Oude Kerk, in Delft, where they take photographs of the tomb of Antonie van Leeuwenhoek, the father of Microbiology, and they tag their pictures with “*Leeuwenhoek, microscope, microbes, scientist*”; it is important that, when the Oude Kerk is recommended to a user who has visited the Hunterian Museum, the reason behind this recommendation (that it is the resting place of Antonie van Leeuwenhoek, the father of Microbiology,) should remain evident. Also, if travelers who have visited the same location as a user have tagged Covent Garden, in London, with “*market, shops, New Wave,*

restaurants, architecture, Punk, people, Royal Opera House, street artists, pub, shoes,” the images that are shown to the user should reflect that variety. In order to achieve relevance and diversity, Near2me uses the images’ textual tags. As mentioned before, each recommended place consists of a set of geographical coordinates marking its center, and the information about the expanse of the place. Within the limits of the place lie the geotags from the photographs that triggered the recommendation, the photographs taken by travelers who also traveled to some place the user visited. These photographs have textual tags, which were originally added by their authors with the intention to organize them, describe them and add context to them. The keywords in these textual tags explain what and why the author of the photograph wanted to document in this particular place. These are the keywords used by Near2me, when retrieving images geotagged within the place’s geographical expanse, to ensure that the images displayed are relevant and at the same time diverse enough for the user to adequately grasp the spirit of the place and what is there for him/her to see.

5.2.3 Recommendation Control and Lead Following

When users receive a group of recommendations and proceed to identify the recommended items and to evaluate them, it is also desirable that they are able to follow some leads that aroused their interest and see ***more items of a kind***. The possibility to do this, places users in control of the system, as opposed to mere observers, and allows them to engage in active exploratory behavior. The ability to actively pursue one or another form of recommendation allows users to communicate to the system information about their preferences at the moment (i.e. for that particular trip, or that particular day), complementing the system’s knowledge of the users’ more permanent interests (which have been derived from their image capturing history).

It is our view that the input of immediate preferences should occur (iteratively, if necessary) as a refinement of the recommendations brought about by permanent preferences, after an initial set of recommendations has been produced and shown to the user. The reasons for this are multiple. In the first place, refining after having had an overview of what is there in a region to see and do allows users to take hints from the environment about what is typical or special about the region rather than focus solely on their interests, which should widen the landscape of choices in comparison to what the user would typically search for if unassisted. Secondly, it eliminates the mental exertion, the uncertainty and the frustration that could occur if the user has to choose from a restricted number of options while being blind about what these options are, this problem being akin to choosing a meal at a restaurant without any knowledge of what is in the menu: it is hard to “come up” with an imaginary menu from which to choose, once a choice is made it is uncertain if it would be a possible option, and finally even if it is there is always the chance that one is missing something one did not think about but would have preferred if it had been a visible option. Also, if the input of preferences takes the form of a dialog, rather than an up-front choice dump, effort is made in smaller amounts and immediately followed by a visible response from the system, which makes the effort more seamless and worthwhile. Incremental amounts, with immediate responses, also render errors more unlikely and easier to recover from, as it is evident which action caused the undesired response. Finally, a continuous step-by-step dialogue between the user and the system shapes the interaction as a continually user-driven exploration, through which the user has the feeling to be the *discoverer* of place, rather than the passive recipient of information, reinforcing the authenticity, the individuality and the unique character of the user’s connection with the place.

Near2me affords this dialogue of discovery in two ways: through the use of tags that allow the user to filter the results, and by offering the user the possibility to follow expert users’ perspectives of a place, a topic, or a combination of both.

As mentioned before, each recommended place consists of a set of geographical coordinates marking its center and the information about the expanse of the place. Within the limits of the place lie the geotags from the photographs that triggered the recommendation, the photographs taken by travelers who also traveled to some place the user visited. These photographs have textual tags, which were originally added with the intention to organize them, describe them and add context to them. These textual tags contain the information of what is in these photographs and, by extension, what is there to see in these places. If a user is recommended a collection of places, a group of tags describing these places comes with it. Near2me analyzes these tags to retain only the ones that are relevant (e.g. removing highly personalized tags and camera equipment related tags), groups these tags into thematic clusters assigning one representative tag for each cluster, and then displays these category-representative tags alongside the recommendations. In this way, the user can have a textual overview of what is being recommended, without the information overload that would result from displaying all the tags. But most importantly, these tags also serve as a navigational tool. By selecting one or more of them simultaneously, the user can direct the recommendations he/she receives by forcing them to include the selected tags, hence seeing ***more items belonging to one particular category***. For example, if a user is browsing recommendations within Glasgow and one of the tags is “*park,*” by clicking on it, the user will see all the parks in Glasgow and all the tags grouped under “*park,*” which describe these new recommendations: “*Glasgow Green, Nelson memorial, winter gardens, fence, Kelvingrove, skate park, botanical gardens, Strathclyde, river bank, bench, people.*” Any number of tags can be cumulatively added as filters, or removed, at any time, modifying the recommendations issued by the system.

It may be also possible that the user is not only interested in seeing more items in the same category, but *more items seen from one particular point of view*. Rather than explicitly offering taste neighbors, Near2me addresses this issue by presenting the user with the identities of the travelers who have taken the photographs that compose the recommendations the user is seeing at any given time. If the user finds that any of the recommendations he/she is presented with shows a particular perspective that calls his/her attention, by selecting the author the user makes the author a filter such that the same recommendation is only portrayed from the point of view of—by the photographs taken by—this author, (eliminating other possibly existing constraints allows the user to see more general recommendations based on increasingly broader items documented by the author of interest.) To guide the user as to the credibility and competence of an author's take on a particular place or theme, Near2me provides an indication of which photographers, or place chroniclers, are experts at portraying the different aspects a user may be interested in exploring. In the same way that people who have coincided in one or more places they have visited are not necessarily absolute taste neighbors (i.e. do not necessarily agree in most of their appreciations of other places), people who have been deemed to be good by the user community at documenting or portraying one aspect or one place, are not necessarily good at documenting any aspects of any place. Near2me can analyze the tags attached to each photograph to determine what aspects, themes or place it portrays. It can also analyze comments left by the photo's viewers at the photo-sharing website where it is published, to know how much this photograph is appreciated by those who encounter it. By collecting both these kinds of information for each photograph for each particular author, Near2me can determine which authors are experts at portraying which subjects and places. When a user is browsing recommendations for a place or subject, and seeing photographs for these places, the authors of these photographs appear displayed alongside the images. Authors who are expert chroniclers and photographers of the subjects or places that the user is browsing appear highlighted, to let the user know that these authors' perspectives may be interesting to explore.

5.2.4 Further Exploration of the Recommendations

Unlike other recommenders analyzed in our survey, Near2me is more of a tool for discovery, which allows the user to obtain a predominantly visual overview of a place and to find out what is there to see and do at a region, than a tool to *learn* about the place (in the sense of getting deep background information about the place, its culture and heritage) or to obtain practical information necessary to plan an itinerary (e.g. transportation to get somewhere, when is it better to visit a destination). To allow users to obtain these kinds of information about the objects, locations and experiences they can discover through Near2me, a link has to be established between visual and textual information—the latter being the format in which most other information is found, searches can be conducted on the web, questions can be addressed within the users' social networks or to their travel partners, etc. Sometimes, place names can be of help, identifying the tags corresponding to each place and to each photograph is also useful. But it is still possible that the user is interested in something he/she is seeing in one of the images, with which he/she is unfamiliar with and no sufficient information is given through the available textual cues (name and tags).

In order to address this, and allow users to more easily obtain further information on items that they can only identify visually, Near2me includes an image-based search feature. By selecting an object in an image—drawing a contour around it—the user can launch a search that will retrieve more information about the object in the form of: more tags that can be found within the system annotating other images that contain the exact same object (found through image content analysis) or external links to pages where more textual information can be found around images that contain the exact same object (again, found through image content analysis). These new tags, which may have not been displayed before for a variety of reasons (e.g. photograph's geotag not matching the location the user is exploring, tag belonging to a photograph taken by a user who did not participate in the recommendation process), and the information found in the external links provide a connection between the visual images that prompted the user's interesting and the textual information the user needs to deepen his/her knowledge of the item or proceed to putting in practice the recommendation.

Image based search, however, as implemented by Near2me according to an algorithm developed within PetaMedia, has its limitations. As it depends on the visual identifiability of the object on which the search is being conducted, Near2me's image-based search will only work satisfactorily for iconic entities that have a distinctive shape (e.g. typical buildings, musical instruments, traditional craft styles, brand logos) and not for generic items. In principle, this should not represent a major impediment for its expected use within Near2me, since users are expected to require further information on items that they cannot recognize, which are expected to be things very particular from the region they are exploring and not generic everyday things. This assumption should be tested in subsequent iterations, through the evaluation of spontaneous behavior of users of a working prototype.

5.2.5 Recommendation exploitation

Apart from allowing the user to obtain, even if externally, more background and practical information about the recommendations, a recommender needs to offer the user a suitable way to store, carry and put to use the information the

recommender itself provides about the recommendations the user wants to put into practice when the user needs to (i.e. while the user is visiting the city and when he/she is planning to see one or several of the recommended locations).

This aspect of Near2me is not developed to its full potential in this first iteration of the concept. Ideally, a scenario in which the user comes back to use the information gained from the recommender should be supported by the ability to merge in the information that the recommender does not provide (should the user be interested in doing so), in this case background and practical information, and should offer more connectivity with other systems that the user could also be interested in connecting to his/her activities (social networks, location-based services, microblogging platforms, etc.) However, as a first iteration, we have decided to concentrate in aspects pertaining Near2me only, and the information it provides. Connectivity with other applications and compatibility with other systems, nonetheless, are acknowledged as a fundamental part of any web-based system's success.

With the focus on allowing the user to store, carry and make use of Near2me recommendations on locations, we added the possibility to save a recommendation (consisting on an item or category of items) to an *itinerary* (which is simply the collection of a user's saved locations for one defined region that is being explored. Also, we added to the concept a mobile application that can be accessed from any mobile device with internet connectivity that allows the user, not only to access the items saved to an itinerary, but also provides some extra (basic) functionality that makes it simpler for the user to access and act on this information: (1) the possibility to set alerts that will notify the user when he/she is in the proximity of one of his saved locations, (2) an option to navigate from the user's current location to the saved location. All the recommendation-related functionality present in the web application (described in the previous sub-sub-sections) is also present in the mobile application, also allowing the user to explore and obtain recommendations for a place on-site.

5.3 Comparison to the Other Recommenders and Explorers studied in our Survey

Near2me differs itself from other travel recommenders and explorers in multiple dimensions. Some have been mentioned throughout its description in the previous sections, below is a summary of the main aspects that distinguish Near2me and its consequence for users.

- In the first place, Near2me differs from some recommenders and explorers in the object of its recommendations. Near2me aims at **helping the user discover what is there to see and do in a place only**, offering a glimpse of what he/she would experience as a traveler by being at the place, without offering any deep background on the place's culture or heritage (like travelogues do) or information to help the user with the practical considerations of planning a trip (like classic recommenders do). Near2me only offers a means to help the user look for this information externally (through its image-based search).
- Unlike other methods (i.e. travelogues, location-based annotators, or classic travel planners) locations do not depend on someone making an **effort to introduce information to the system**, a few passer-bys who thought that the location made an interesting memento is all that it needed: people who would have taken these photographs anyway.
- The effortless submission of a place to the system, and the fact that the submitting user need no more information in order to enter a location for potential recommendation than having passed by it and liked it, increases the likelihood that **more (and more obscure, off-the-beaten-track) locations will be in the pool of recommendable places**.
- Another result of the implicit sourcing is that photograph geotags are also **untainted by the ulterior motivations** that may be present in deliberate entries and descriptions (in the case of travelogues or location-based annotators, for example). This makes the information from which recommendations are created and which is shown to the user in order to evaluate the recommendations **more authentic**.
- On the other hand, unlike in recommenders based on location tracking, which also gather information implicitly, in Near2me an implicit entry—a photograph—is charged with enough deliberation on the part of the user to be **clearly considered up for recommendation** (whereas this is not the case for points where the user has stayed a given amount of time).
- Both photography-based examples in our survey, photography-based explorers and recommenders based on trends sourced from photography, present the benefits derived from implicit sourcing coupled to unambiguous signaling of locations as potentially recommendable, however they offer absolutely no **personalization**, whereas Near2me does.
- Moreover, recommenders based on trends sourced from photography produce recommendations deliberately based on popularity only, being useless for travelers who seek off-the-beaten track destinations. On the other hand, Near2me is able to produce non-mainstream or **off-the-beaten-track recommendations** because it focuses on the conditional probability that one individual user would visit a pair of places.
- Near2me's focus on co-occurrences of pairs of geographical-coordinates within individuals produces recommendations based solely on the statistical correlation between a user liking one of these locations and also liking the other. As a result of this, Near2me's recommendations **do not suffer from overspecialization** as its counterparts based on description, category or user-generated annotations coincidence.

- Finally, unlike in systems where only short snippets (location-based annotators) or one-author descriptions (travelogues and classic recommenders) are available to allow the user to form an opinion about a place or images are chosen deliberately, including editing and doctoring, with the goal of illustrating the location (classic recommenders and travelogues) description of places through images taken by a large group of peers presents the following advantages: (1) **images are highly descriptive**, more than words for the same input effort, (2) images are generally more **faithful to reality** than textual descriptions (since even if they may show only part of reality they cannot show that which was not there), and (3) they express **multiple points of view**. These advantages contribute to successful recommendation evaluation in Near2me.
- On a different note, compared to other systems that offer recommendations personalized based on user behavior (like some classic travel planners or recommenders based on location tracking), Near2me **does not suffer from the new user problem**, since the assumption can be made that every beginner user counts with a collection of photographs taken during his/her travels. Although it is true that in many cases these photographs may still not be geotagged (automatic geotagging being relatively new and not universal), there is a trend for all mobile devices (cameras and camera-phones) becoming location aware in the near future, most photo-sharing websites offer the possibility to manually geotag images, and algorithms such as the one described by Serdyukov et al. (2009) have proved that it is possible to geolocate photographs based on image metadata.
- And, on the other hand, as opposed to recommenders that rely on explicit up-front user input in order to personalize their recommendations (which of course also do not suffer from the new user problem) Near2me is **not limited by the need for users to understand the recommendation model** in order to modify their preferences as a step to get more accurate recommendations.
- And, finally, although Near2me's recommendation algorithm does not include any learning from user feedback, as opposed to what happens in systems that only have access to the user's interaction with the system (e.g. classic travel planners) **it would be possible to include feedback from users interaction with the object of the recommendation itself** (the location) assessing if the user did follow the recommendation or not and how much they found it interesting through the study of subsequent user photographs during a trip to the vicinity of a recommended item. (In addition to the always-available feedback of a user saving a location for future reference.)

5.4 Interface and Interaction Paradigm

The Near2me interface (**Figure 1**) was only developed to a conceptual level, with emphasis on content rather than form, defining only those components that are essential and define the flow of information during user interaction.

The Near2me interface is based on five panels that show information consistent with the same group of recommendations at any given time. (1) Photographs of the places, (2) place locations on a map, (3) a list of place names, (4) descriptive tags, (5) travelers who have photographed the places. The group of recommendations being displayed at a given time is defined by constraints shown through *tag crumbs* (**Figure 1**). These constraints are added by the user, through filtering while browsing the recommendations. The browsing takes place through the panels. Each panel is, simultaneously, a display and an input element (i.e. each image that depicts a place, will lead the user to more recommendations in that place when clicked, each tag that describes a place will lead the user to recommendations filtered according to this tag when clicked). Once the user chooses the destination that he/she wants to explore, the system is designed to work through browsing, as opposed to search, browsing being better suited for exploratory search. Exploratory search involves ill-structured problems and more open-ended goals, with persistent, opportunistic, iterative and multi-faceted processes aimed more at learning than answering a specific query (Marchionini, 2006), the kind of tasks with which Near2me aims to assist its users.

In order to allow the exploitation of recommendations when the user is at the destination, Near2me includes also a mobile application, which affords the same exploratory capabilities of the web interface described above, and the following location-based functionalities:

- The possibility to set alerts that will inform the user when he/she is closed to a saved location.
- A navigation service to take the user from his/her current position to a saved location.

No mobile interface is specified given that mobile interfaces are highly platform dependent.

5.5 Use Scenario

This section relates a typical use scenario. For a more complete and illustrated demonstration of a Near2me use-case, please refer to the video that accompanies this report.

John is our typical user. John likes to travel, but when he travels he seeks authentic experiences that allow him to connect to the spirit of the place. He doesn't like manufactured, massive tourism activities and locations. However, the living culture of a city is as complex and multidimensional as the cultural expressions of people who inhabit it and it can provide authentic

experiences in many domains, for example: food, extreme sports, architecture, technology, fashion, environmentalism, etc. For this reason, the random choice of activities outside the mainstream can result in bizarre choices that are too far from John's interests for him to engage. John best enjoys situations that take place in the intersection of his own identity and that

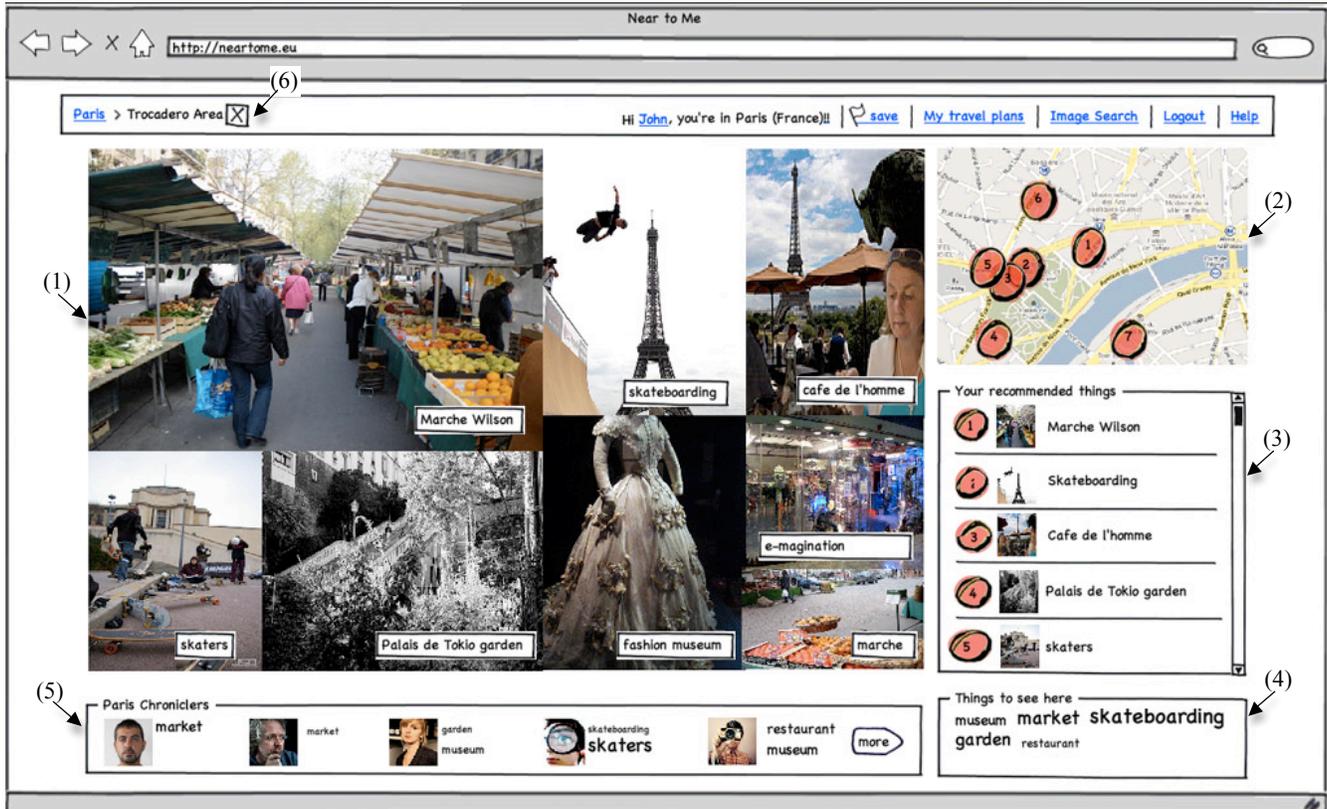


Figure 1. The five panels in Near2me's interface: (1) Photographs of the places, (2) place locations on a map, (3) a list of place names, (4) descriptive tags, (5) travelers who have photographed the places. And (6) the tag crumbs showing the constraints for recommendations.

of the place.

When John visits places he takes pictures, not only of landmarks he intends to visit but also of interesting and curious things he finds along the way to them. The subjects of these pictures appealed to him because of: their descriptive value of the spirit of the place, their relation to some activity or interest that he likes, their aesthetic value (or, usually, a combination of these). He uploads these pictures to a photo-sharing website where he tags them to organize them, describe them and add context. Each of the pictures he takes also has a geo-tag that marks where the picture was taken.

John wants to visit Paris. He turns on his laptop, logs into Near2me and chooses Paris as his destination. Near2me knows John's interests based on the locations where he took pictures during his previous travels. For example, we can see from pictures of John at the Borough Market, in London, and other markets that John has visited, that John enjoys going to local markets to discover local food and crafts. Other travelers also share this interest. Travelers who enjoy visiting Markets are likely to do so in many of the cities that they visit; because of this, geo-tags marking the places where markets are are likely to co-occur within travelers' picture sets. In this case, travelers that visited the Borough Market in London, are much more likely than other people to have also visited the Marché Wilson in Paris, to Near2me this means that these two places are related. We can assume that if John enjoyed the Borough Market, he will enjoy visiting the Marché Wilson during his trip to Paris and recommend it to him, since this market is not a mainstream tourism destination it is unlikely that John would have found it by traditional means.

Let's see how this is done. The Near2me interface shows John five panels (**Figure 2**).



Figure 2. The five panels in Near2me’s interface.

(1) In the first panel, there are pictures of the places Near2me recommends for him. (2) In the next one, a map showing the places’ locations. (3) The third panel shows a list containing the names of these places, which are produced automatically by Near2me based on the pictures’ geotags. (4) Next, a tag cloud describes the things that he can find in these places. The tags are clustered into topics, with only one representative tag showing for each topic, to ensure that John can have a broad view without clutter. (5) And last, a list of travelers who share John’s interests and have photographed these places.

John sees a picture he finds interesting. When he hovers the cursor over it, all the information related to this place highlights in each panel to provide context for what he is seeing (**Figure 3**).



Figure 3. When John hovers over a picture, the information related to the picture highlights in each panel to provide context for what he is seeing.

Clicking on the picture, John dives into this location, the panels adapt their content and John is now exploring the Trocadéro area. The tag cloud shows what Near2me recommends for John in Trocadéro. John clicks on “market”, he is now in Paris > Trocadéro area> market. The tags now show: “Marché Wilson” among others, this is what he’s seeing. The pictures depict the place. Near2me makes sure that the pictures John is seeing are diverse enough for him to have a good idea of what can be found or done in this place. John likes the place and saves this to his *travel plans* and naming it Marche Wilson (**Figure 4**).



Figure 4. John saves a place to his travel plans.

John wonders if there are any other interesting markets in Paris. He removes the Trocadéro Area label and now he can see all the markets in Paris. Some travelers appear starred (**Figure 5**). This means that they are expert market chroniclers and photographers. By clicking on a user's avatar, John can experience markets in Paris from her perspective.



starred travelers

Figure 5. Green stars appear next to some travelers' avatars, which means they are expert market chroniclers and photographers.

He can see several markets he's interested in. He adds the group to his *travel plans* and continues exploring.

In one picture he sees something unknown to him, the information from the tags is not helpful to know what it is. He launches the imaged-based search. He picks the brush, he draws a contour around the unknown object, and he clicks on the search button (**Figure 6**). Image-based search will look for related tags that can help John figure out what the object is.



Figure 6. He launches the imaged-based search. He picks the brush, he draws a contour around the unknown object, and he clicks on the search button.

He continues to explore Paris and save stuff to his travel plans.

Now John is in Paris. He logs into Near2me from his mobile phone. From his phone, he can not only continue exploring places and things as usual, but also access his saved travel plans and use them as his personalized guide while in Paris. He can set alerts, for Near2me to notify him when he's close to his saved locations or use a navigation function for Near2me to guide him there. He looks for the Marché Wilson in his travel plans and selects *take me there* (**Figure 7**).



Figure 7. John logs into Near2me from his mobile phone. He can use a navigation function for Near2me to guide him to his saved locations.

The navigation service opens to give him instructions.

On his way to the Marché Wilson, John finds an interesting building and wants to know what it is. He takes a picture of it with his phone and uses Near2me's image-based search to learn more about it. Near2me shows the building is the Assemblée Nationale and more information is provided.

Finally, the market is around the corner. Another opportunity for discovery, whose photographs—on location and along the way—will fuel more recommendations and exploration.

6. THE PETAMEDIA TECHNOLOGIES BEHIND THE CONCEPT

Six PetaMedia-developed technologies provide the technical basis for Near2me.

- The recommender algorithm was extracted from Clements et al. (2010) article *Personalised Travel Recommendation based on Location Co-occurrence*.

- The place names for the recommendations are found by matching photograph geotags and textual tags to a database of landmarks compiled from public sources, defaulting to gazetteers if there is no match²³.
- The images depicting the recommendations, provided for recommendation evaluation, are made subject-relevant and diverse by retrieving them using a combination of their geotags, textual tags and image content²⁴.
- The expert chroniclers and photographers (starred travelers) are found through an algorithm for the determination of subject-related authority based on comments made by peers on a user's uploaded content. Work based on four publications by Wartena (2010), Neubauer and Obermayer (2009), Raaijmakers and Kraaij (2010) and Ramzan et al. (2010).
- The tag clustering is carried out by a set of algorithms based on the unpublished and published the work of Wartena et al. (2008a, 2008b, 2009).
- The image-based search is provided by a PetaMedia-developed object duplicate detection algorithm²⁵.

7. CONCEPT EVALUATION

7.1 User evaluation Methodology and Choice of Prototyping Technique

There are two dimensions that is desirable to evaluate in an application with users: concept and performance. The concept refers to that which is intended, a **concept** is validated when it is proved that it responds to user requirements (there is a group of target users for whom it supports the concretion of a desired goal in the required circumstances) and it is technically feasible (FAA, 2008). Concept validation aims at *answering are we building the right system?* As such, it is not focused on **performance** issues like *are we building it in the right way?*, which are relevant only in the implementation phase. In the case of Near2me, which has only yet reached its concept development stage, evaluation will be focused on concept validation.

The development of Near2me was highly technology-driven, with its component technologies as starting point. As such, it is considered to be feasible, with the evaluation focus resting solely on adequacy to user requirements.

More details on the rationale behind the evaluation methodology for Near2me are provided in **Part 2** of this report. Here, only a brief account of the research questions, the evaluation methodology and the choice of prototype are given.

The goal of this evaluation is

1. To assess if potential users understand the concept behind Near2me.
2. To determine if users recognize within Near2me the motivations identified in the **Preliminary User Research** in **Section 2**:
 - Search of an authentic exchange, away from "manufactured" tourism experiences.
 - Search of a more personal interaction with the city, an experience that arises from the unique combination of the personality of the visitor and the place.
3. To identify what are the aspects of Near2me that users find more salient (positive and negative).
4. To identify if there are any reasons or circumstances that would prevent users from using Near2me, and if users picture themselves using Near2me before/during their travels.
5. To evaluate if potential users are able to appropriate the concept, thinking of custom modifications or additions and suggesting novel uses for it.
6. To research how the concept behind Near2me compares, in the mind of its target users, to other applications of its kind (i.e. determine if users see distinctive value in the proposal of Near2me).
7. To obtain suggestions for improvement in subsequent iterations.

In order to focus on the conceptual aspects only, without any shortcomings in performance affecting the evaluation, we have opted for a video prototype, justified by its mixed fidelity (McCurdy, 2006) along the following dimensions:

- Low visual refinement, not needed for an evaluation of the concept.
- High breadth of functionality, to show users comprehensively what Near2me can do and give them an idea of its usefulness. Every major functionality, each of them derived from one PetaMedia technology, is featured in the video.
- High depth of functionality, to show users the manner of execution of tasks afforded by each functionality and give them an idea of its range of capability, acceptability, ease of use and enjoyability. One task related to each functionality is depicted to completion in the video.

²³ Unpublished work by Yue Shi (y.shi@tudelft.nl) and Pavel Serdyukov (p.serdyukov@tudelft.nl).

²⁴ Unpublished work by Stevan Rudinac (s.rudinac@tudelft.nl).

²⁵ Unpublished work by Peter Vajda (peter.vajda@epfl.ch) and Ivan Ivanov (ivan.ivanov@epfl.ch).

- Low richness of interactivity, not needed for a concept evaluation and will be the object of evaluation in subsequent iterations when a working prototype is available.
- Medium richness of data model, to show users the nature of the relationships that prompt recommendations in Near2me and the variety of information with which they may be provided, without falling into the explanation of exceedingly complicated mechanisms that may confuse them.

Twelve participants, aged between 22 and 35, gender and occupation randomized, were shown a video prototype that describes Near2me through a use case (for the prototype, please see the video accompanying this report). Once shown the video, the participants took part in a semi-structured interview, guided by the research questions described above. The information provided verbatim by the participants was coded according to the research question to which it relates. And finally the quotations were collectively subjected to a qualitative analysis to address each research question.

7.2 Results and Conclusions

Immediately after watching the video prototype, participants were asked to describe Near2me and what it does in their own words, focusing on *who* would use it, *what* they would use it for and *in which circumstances* they would use it. All participants demonstrated their understanding of the concept in terms of the rhetorical situation it addresses (its users, users' goals, and users' circumstances). There was little mention of Near2me's mechanisms of operation (e.g. *how* recommendations are produced, how results are presented) with emphasis on *what* Near2me allows users to do. The following quotes exemplify how participants describe Near2me:

- “It is a website that recommends places for people who don't want to do touristic stuff and what to see things that are more authentic. You go there to plan your trip and then you can use [your plans] when you're traveling.”
- “It shows you recommendations for when you travel based on what you like instead of what is popular, so if you are not into seeing famous things that are really touristic you know what to do in a city.”
- “It is for people who want to see a place not from a touristic perspective but also not seeing whatever, so it takes your interests and shows you things you may like. Then you explore this and can choose what you want to do.”

However when participants were later asked about the most salient aspects of Near2me and what they like or dislike about the concept, they made reference to the recommendation, evaluation, recommendation control, lead-following and recommendation exploitation methods, which suggests that *the way in which the rhetorical situation is addressed* is also important for them. Participants recognized within Near2me the possibility to fulfill the motivations identified in the preliminary user research, and they explicitly appreciated the components deliberately included in its design in order to provide this possibility. Below we present some quotations obtained during our user evaluation that describe participant's perception of the different aspects of Near2me.

About Near2me's recommendation method, participants recognized how it takes into account the traveler's *personal interests*:

- “It's more personal than other ways, and also more accurate because it looks at the things that you like. It is hard for me to trust recommendations from other people [taste neighbors] because I don't know if their opinion will be my opinion.”
- “I like that] things are not recommended because they are the same, your example [in the video prototype] is not maybe the best because it's market-market, but I don't like about recommenders is, for example in Imdb²⁶ that you look for a movie about Elizabeth the queen, and they recommend only movies about Elizabeth and other kings and queens, and it's obvious they are getting very similar things, but I want variety. I can see that if these are places that the same people go to, you'll get variety because it's not about being the similar, it's about being both liked by the same people.”

About finding recommendable items through crowd-sourcing, participants mentioned how this method *broadens the pool of places available for recommendation*:

- “I'm sure you can find a lot more places that in any guide or website, it's like Wikipedia: as soon as things appear, they will be there immediately because there are a lot of people doing it.”
- “I'd expect to find graffiti in Near2me [...] because of other people my age or that are graphic designers, this is something I'd expect to find in Near2me and not in a guide.”

About finding items worth recommending through photo capturing, participants mentioned how this method makes these recommendations be perceived as *authentic* and *valuable*:

- “The experiences are provided by real people, things explored have value because some people have thought they were valuable and so captured them.”

²⁶ The Internet Movie Database (<http://imdb.com>)

- “I like that it is based on pictures, I have many friends that have their Foursquare²⁷ connected to Facebook and it’s like people are not sincere at all, they just check-in in places to look cool, pictures are more real.”
- “The fact that there is a stream of pictures tells a lot about not just the place or the person [who took the pictures] but about the context. It’s like a conversation of people about the place. Views are expressed and exchanged. These people might capture something that you have not thought of [referring to a different view of the same place].”
- “I like the fact that it connects images with interests, that because you find it nice you took a picture.”

About providing user-generated photographs for recommendation evaluation, participants showed their appreciation and remarked how this made images *authentic* and *diverse*:

- “I like the pictures, that you can see everything and you can trust the criterion, because you ‘see’ it.”
- “[I like that] the things it shows are real, they are not pictures from brochures that you have to take from a helicopter and be a professional photographer and, of course, if you go to the place you will never see. I remember I was so disappointed when I went to Chartres and there is no way to take a picture of the labyrinth or see it from above but that appears in all guides!”
- “In my case I like going to aquariums. I think it’s not so out of the mainstream. I can find aquariums in a guide, but a guide will not tell me that some are tiny and for children full of rooms with interactive activities and no animals, and others are awesome and have giant tanks, what I like is this is something I can see from other people’s pictures.”
- “The thing I like best is that finally you can see real photographs of a place you want to go to. There are two things, that the images are real and not trying to advertise the place and that [...] the images are diverse so I can have a real and good idea of everything that is there and all that can be done.”

Participants showed their appreciation of Near2me’s *implicit learning of user’s interests* (through their own photograph capturing behavior) and how this facilitates recommendation when users have *non-specific queries*:

- “I like that you don’t have to make a search, because if I don’t know the place, how will I know what to search for?”
- “A good thing is that you don’t have to specify things that you want to see, when I see the travel websites that ask you to say what you want is like the chicken and the egg, if I know what I want so I can specify it, why would I go to the website?”

Participants mentioned that they value and enjoy that Near2me allows the user to carry out his/her *own discovery* of the place by browsing:

- “I like that you can browse yourself, that is dynamic, that you browse and you discover yourself and go as far as you want to see everything.”
- “It is important to me that apart from getting recommendations, I browse and see and explore.”
- “I like that you can spend time exploring, because this is a part of traveling I enjoy.”

Participants appreciated that users are *in control of the subjects* of the recommendations:

- “I also like that you can filter by tags, because I think there is one thing that is who I am and what I usually enjoy but also I want to take into account what I want to do right now. Choosing from tags allows me to say ‘OK you showed me all things I like but today I feel like markets’.”
- “I like the way you can jump from one subject to another, for example see the Trocadéro , then Trocadéro markets and then say ‘I want to see all the markets’ and be able to remove Trocadéro without going all the way back.”
- “I like that you can jump between filters, select and deselect regions.”

And they also valued that users are *in control of the recommendations’ point of view*, by seeing places and things from one particular (*expert*) traveler’s perspective:

- “I trust more something based on my own interests than other people. It is important that in Near2me I see which people are experts, but I’d only trust this through time if I see they’re right.”
- “One of the most interesting things is exploring the pictures of specific people by topic or by area, you can have their view for each place and thing.”
- “I would like it to recommend me things I will like but, as I see them from other people [‘s perspectives], I can see the same things in a different way and also learn something.”

Participants mentioned the importance of being able to *connect image-based recommendations with a textual information-based Internet*:

²⁷ A location-based annotator (<http://foursquare.com>)

- “The image-based search is also important, because if you see something you like but know nothing about and you want to go deeper, how would you do it? I like architecture and this happens to me all the time with buildings.”
- “I also like image-based search, because it solves the problem of when you’re there you see something you like and you think ‘this must be something’ but how can you know? You know nothing to start your search.”

Participants appreciated that Near2me offers means to help users to *exploit these recommendations*:

- “I like that you can make your travel plans and then use them.”
- “I like that you can save items so you don’t get lost in everything that you have browsed and explored.”
- “I am not a person that plans so much before going to a place, if it’s not a big trip I’m not going to search, so it’s nice to have a mobile Near2me that I can use when I’m there and will tell me things I’m going to like and are close by.”

Participants also mentioned being attracted to the *“social” aspect of Near2me*. This social aspect had not been deliberately intended as an incentive for the use of Near2me by design, but was mentioned several times throughout the interviews:

- “[I like that] it looks like some organic mapping of the world’s wonders for people by people.”
- “I like that although people are not present it makes me feel part of a community of travelers, because I know other real people took the pictures, these are their experiences and they are sharing with me, telling me something about the place like hobo marks [...] it feels free and adventurous.”
- “I would use it. I like the usefulness but I also like the philosophy: it’s like a Wikipedia of places [...] something I want to participate.”
- “This would motivate me to take pictures and feel like an explorer when I am visiting places, discovering off-the-beaten-path spots for Humanity.”

During the interviews, all participants expressed that they would use Near2me. All of them *could think of real past travels or current travel planning scenarios in which they would like to use Near2me and describe how they would use it*. The following quotations offer an impression of the breadth of participants’ contributions:

- “For me it would be interesting for what I’m planning now: I’m going to Florence, and I’ll stay there three weeks and I’ve already been to Florence [many] times. There are still lots of small things in the city non-mainstream that I’ve never been [to], [...] and now I’m kind of lost to search, a system like this would be good. When I go the first time to a city, I go for the mainstream stuff [...] but later on if I go several times... One thing I specially like is bridges”
- “If it is a small city [...] and I have three days, then in two days I see the main things and the third day I use Near2me to look for more personalized things to see.”
- “Last week we went to Paris and I found this walking tour by an artist by chance that went to non-mainstream places like where Oscar Wilde died and that walking tour made all the difference about visiting Saint Germain, and I see the link between that [and Near2me] and this is why I get excited [about it]. Saint Germain is just a neighborhood and it has cute stores but if you know all this about it [it’s better]. I can imagine to [build] my own tour of this kind.”
- “What was in the video, when John ended up in a place that he didn’t like, that happened to me and is something I didn’t like. I was in London and traveled for one hour to get to one place and in the end it wasn’t nice and it wasn’t what I expected, I read [about] it on the Internet and people said it was nice. Near2me would have helped me by showing me real pictures of it before I go.”
- “I went to New York and we had tight time, and I wanted to go and see Harlem, and I knew that there were interesting graffiti by Banksy (a graffiti artist) there we walked around and I didn’t find them. With Near2me I would have found exactly where they are.”
- “I’m not a non-mainstream person, but the way I understand the system is that if you like mainstream places it will recommend all the mainstream places because it looks at where you have been and taken pictures, so I would use it.”
- “I’m not very organized in the things I do, I don’t check for the weather and the optimal route, or prices for tickets, I try to see a general direction of where I to go and then I go. I’d mostly use it to get an overview of the places I would go to, and as a reflective tool to see how other people have experienced the places.”
- “Sometimes I am at a place and I don’t know what to do next, having [Near2me] at the tip of my fingers of course it would be useful.”
- “I would use it. Maybe I would use it in my city too, it looks stupid to do the tourist stuff in your city but this is different.”
- “When I travel I am very interested in shops [...] typical local very old grocery stores or pharmacies, these hand-made hat shops, I like to take pictures of show windows. It would be interesting with Near2me to find all these and set an alert to see when I’m passing close to one.”

Two *situations were identified in which Near2me would not be used*:

- “[I wouldn’t use it] when it’s longer holidays [and it’s worth spending time organizing.] I like the experience of having the guide and reading and reading, but when I’m not taking this time, I’d use Near2me.”
- “I wouldn’t use it if I travel with friends because the personal profile [for recommendations] wouldn’t make sense.”

However, participants defined many *challenges that Near2me would have to overcome in practice and conditions that it would have to fulfill to be acceptable*: concerning photography privacy and sharing habits, and timeliness of recommendations related to opening times and seasonality of activities. They *suggested several relevant features to overcome these problems*:

- “What I don’t like is the people who appear starred, the experts, because I wouldn’t like to be shown together with my pictures and strange people know they are from me, I am ok if it is anonymous, I would like to be able to opt out to be identified.”
- “I would use it, but can I use it without uploading my own pictures? I don’t usually upload all my own pictures, just a few. Maybe I can be shown pictures of others and say I like them? Maybe pictures I have liked in Facebook?”
- “I can see a problem because things can be recommended because other people took pictures and then they are not there anymore. For example, a market could only be on Tuesdays and Sundays. It would be interesting to take other data from the photos, not just the geotag, because photos have date and time and see that if pictures of markets appear only on Tuesdays and Sundays there, you only get this recommendation for these days. This could work with other data also, for example, time of day because museums may be close, and some things, like sledding are only good in winter.”

Participants were also enthusiastic and quickly familiar enough with the concept to *appropriate it and suggest new uses for it*, related to their own realities, when asked to do so:

- “I would also like to have the same but real time, the same information with photographs, and places and things, but to see things popping out as I go, like some sort of cross between Near2me and Twitter.”
- “I’d like to have a system like this for when you’re new in a city, that also has activities and shops.”
- “I would also like to have a system like this for “places” on the Internet, like where you have been but inside the internet and then it recommends me websites.”
- “Maybe this could exist with weather, time of day for activities, instead of with geotags for places.”

When asked to compare Near2me to other methods they know or use, participants mentioned that they did not use other recommenders and *explored places for travel through manual processes that they often described resembling Near2me*:

- “Right now it’s like you have guessed how I do it, I always start looking if I want to go to a place in Flickr, I go there and see the pictures other people took, that’s how I decide if a place is even worth it. Then I look for information that is not related to tourism at all, for example Wikipedia about the history and the culture or watch movies and read some books. Then, unless it may be dangerous (like when I went to Iran for example), I just go and try to find the practical information there asking local people.”
- “I look for pictures in Google and see if there is something nice I try to see what it is in the picture, I also use Panoramio²⁸ to see what is nice close to the are that I am. Then I look for information about experiences in blogs and travel websites (about practical things).”
- “I don’t use recommenders now, so I don’t know how it’s different. I usually ask people I trust, people I have known for a long time that their judgment will be like mine. Sometimes no one like this has been to the place and I use guides and look in Google. But it is very hard to get an idea of what the place will be like, you mostly get information about hotels and how to get there and what to do.”
- “The idea [of Near2me] I like it a lot, it doesn’t look touristy, this is the biggest difference with other [travel recommendation] websites. I don’t like to be a ‘tourist’ I like to think I’m a ‘traveler’.”

Overall, Near2me is consistent with the user motivations found in the preliminary user research (**Section 2** of this report), which participants were able to spontaneously recognize in the design during the user evaluations. Participants identified and acknowledged the value of the different components designed into Near2me. Participants mentioned being drawn to Near2me by its “social” aspect, which was not deliberately included into the system by design, signaling an interesting motivation for user adoption that should, from now, be taken into consideration and maintain in future iterations. Participants mentioned that they would use Near2me and were able to offer a wide variety of different scenarios based on real past travels or present travel planning situations in their lives. Participants identified challenges faced by Near2me for user adoption but were able to offer realistic suggestions for the solution of these acceptability problems. Participants were also able to appropriate the concept and suggest new and personalized uses for it and variations for its development, which is usually an indication of concept potential. Of particular interest is the issue of adding recommendation situational information beyond location (time and date, seasonality, etc.) that is consistent with benefits also explored and found in literature (Baltrunas 2008), is feasible within the context of PetaMedia and would add value including *events* to the current selection of *objects and places* that can be recommended.

²⁸ A photography-based explorer (<http://panoramio.com>)

All in all, Near2me’s consistency with its target users’ motivations and the discovery of a new social motivation, participant’s appreciation of the different aspects that compose the concept, and their ability to imagine it as a part of their lives and appropriation suggest there is value in the concept and it is worth developing it further. Attention should be paid to scenarios where participants would not use Near2me, in order to evaluate possible adaptations of the concept or definitely discard them, and to the challenges identified by the participants in order to address them in future iterations.

8. CONCEPT ABSTRACTION: INFORMATION STRUCTURE BEHIND NEAR2ME

The present section describes the abstract information structure and relationship between entities that lies at the core of Near2me, in order to provide a theoretical model that facilitates technology portability to other problems conforming to the same characteristics.

In 2006, Marlow et al. presented a formalized model of tagging systems that has been widely accepted, numbering more than a hundred citations, and that, to the date, has not been substantially improved upon. In Marlow et al.’s model, “users assign tags to a specific resource, tags are represented as typed edges connecting users and resources. Resources may also be connected to each other (e.g., as links between web pages) and users may be associated by a social network, or sets of affiliations (e.g., users that work for the same company)”. Later, the authors mention that “the aggregation and semantic aspects of tags have been discussed and debated at length”, suggesting that tags can also, in the same manner as resources and users, be related to each other within their same class by “folksonomies and crafted ontologies”. **Figure 8**, borrowed from Marlow et al.’s paper, illustrates this model.

The authors use this model in pursuit of their premise that “a unified user-tag-resource approach might be useful for many key web technologies, including: search and information retrieval; information organization, discovery and communication; spam filtering; reducing effects of link spam, and improving on trust metrics; identifying trends and emerging topics globally and within communities; and locating experts and opinion leaders in specific domains.”

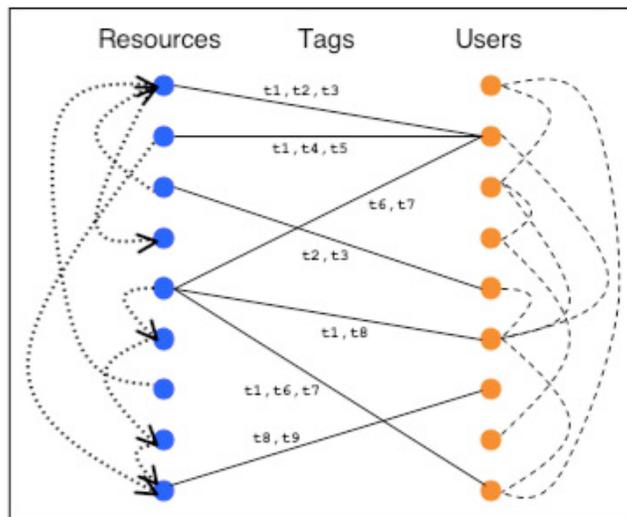


Figure 8. Marlow et al.’s model of tagging systems
Figure taken from Marlow et al. (2006).

The unified user-tag-resource approach is also at the core of Near2me. However, the model on which Near2me is based has been slightly modified to accommodate changes that since 2006 have modified our understanding of tagging systems:

1. The increasing influence of **implicit tagging**, as technology makes methods to obtain this kind of tags more readily available (e.g. location aware devices, mood assessing systems, ubiquity of weather information once location is a given, real time systems that put a new emphasis on time and date). There is literature signaling the particular role of implicit tagging in terms of providing situational information for information retrieval and in particular recommender systems (Baltrunas, 2008).
2. A 2008 study by Kinsella et al shows that “anchortext significantly overlaps with tags,” suggesting that the information derived from the links between resources that Marlow et al. see in hyperlinks between web pages and the information derived from the association on these resources with textual tags may not be different in nature, resulting in the same redundant information.

Near2me maintains the **user-user relationships** in terms of the social network that connects them, in this case an *utilitarian* social network determined by the visit of one or more places by the same individual and in which the relationship has the sole purpose from one individual to benefit from the information possessed by the other (as opposed to purely social or common interest-based social networks where users form more permanent bonds, for example).

In terms of **resource-resource relationships**, Near2me replaces the hyperlinking by relationships found through *media content analysis*, as in the case of image-based search (where an object’s description, further information and links to related resources are found through the link established between an image that contained this object but was deprived of these resources, and another image of the same object—found through object duplicate detection—from which these resources could be harvested).

User-resources relationships remain also based on tags, but a distinction is made between explicit and implicit tagging. The role of explicit tagging continues to be that of describing a resource and the user who applied the tag (in this case, describing a location, and indirectly the interests of the user who visited this location). However we maintain that the information garnered in implicit tags is of a different nature: it contains, instead of declarative knowledge, *situational knowledge* that defines *in which circumstances* it is relevant to take into account—or even consider the existence—of a relationship between an user and a resource (implicit tags are present in Near2me in the manner of geotags, in this case it is only relevant to recommend a location to a user when this location is in the region the user is interested in, the same could be said for other sources of implicit tags—that were brought up by participants in the user study: weather, time and date, season, and even mood, company or any others one can think of).

Figure 9 shows the tag-user-resource underlying Near2me.

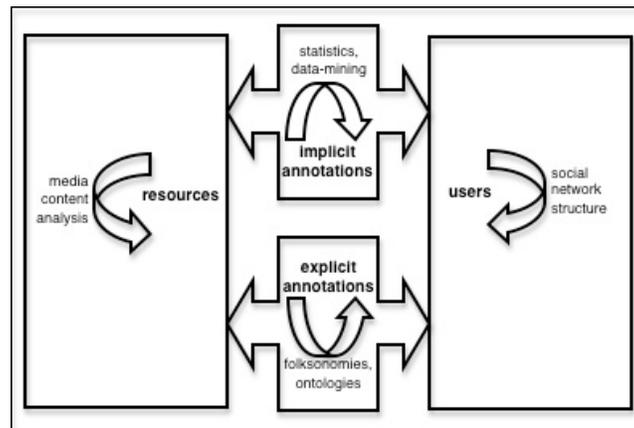


Figure 9. Data structure and user-resource-tag relationships behind Near2me

9. FUTURE WORK

The next steps for Near2me consist on the development of a working prototype to carry out the *PetaMedia Field Trials* (i.e. tests with users that will allow us to evaluate aspects specifically related to the **interaction** and user satisfaction with the **performance of the implemented PetaMedia technologies**). In this work we include general recommendations for the execution of the Field Trials.

9.1 Development of a Working Prototype

The development of a working prototype is already being carried out at the time this report is being written. The working prototype follows the guidelines described in this work in section 5. **THE CONCEPT**, and incorporates the technologies detailed in section 6. **THE PETAMEDIA TECHNOLOGIES BEHIND THE CONCEPT**.

9.2 Recommendations for the Field Trials

A preliminary plan for the execution of the PetaMedia Field Trials is presented in **APPENDIX 1: RECOMMENDATIONS FOR THE FIELD TRIALS**, at the end of this document. The PetaMedia field trials consist of a user test carried out on a working prototype to evaluate aspects of Near2me specifically related to the **interaction** and user satisfaction with the **performance of the implemented PetaMedia technologies**. User evaluations, like the PetaMedia Field Trials, are needed to

evaluate the performance of recommenders if the effectiveness of the recommendations, as well as their usefulness and quality form part of the assessment (Celma, 2008).

10. REFERENCES

- Ahern, S., Naaman, M., Nair, R., and Yang, J. H. 2007. World explorer: visualizing aggregate data from unstructured text in geo-referenced collections. In *Proceedings of the 7th ACM/IEEE-CS Joint Conference on Digital Libraries* (Vancouver, BC, Canada, June 18 - 23, 2007). JCDL '07. ACM, New York, NY, 1-10.
- Ames, M. and Naaman, M. 2007. Why we tag: motivations for annotation in mobile and online media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA, April 28 - May 03, 2007). CHI '07. ACM, New York, NY, 971-980.
- Ardissono, L., Goy, A., Petrone, G., Segnan, M., Torasso, P. 2003. INTRIGUE: Personalized Recommendation of Tourist Attractions for Desktop and Handset Devices. In *Applied Artificial Intelligence* 17(8-9). Special Issue on Artificial Intelligence for Cultural Heritage and Digital Libraries, 2003, 687-714.
- Baltrunas, L. 2008. Exploiting contextual information in recommender systems. In *Proceedings of the 2008 ACM Conference on Recommender Systems* (Lausanne, Switzerland, October 23 - 25, 2008). RecSys '08. ACM, New York, NY, 295-298.
- Berlyne, D. E. 1950. Novelty and Curiosity as Determinants of Exploratory Behavior. In *British Journal of Psychology* 41. 1950, 68-80. [Cited in Lee, T., Crompton, J. Measuring Novelty Seeking in Tourism. In *Annals of Tourism Research* 19, 4. 1992, 732-751.
- Berlyne, D. E. 1960. Conflict, Arousal and Curiosity. New York, McGraw-Hill, 1960. [Cited in Lee, T., Crompton, J. Measuring Novelty Seeking in Tourism. In *Annals of Tourism Research* 19, 4. 1992, 732-751.
- Berlyne, D. E. 1966. Curiosity and Exploration. In *Science* 153. 1966, 25-33. [Cited in Lee, T., Crompton, J. Measuring Novelty Seeking in Tourism. In *Annals of Tourism Research* 19, 4. 1992, 732-751.
- Brown, B. and Chalmers, M. 2003. Tourism and mobile technology. In *Proceedings of the Eighth Conference on European Conference on Computer Supported Cooperative Work* (Helsinki, Finland, September 14 - 18, 2003). K. Kuutti, E. H. Karsten, G. Fitzpatrick, P. Dourish, and K. Schmidt, Eds. ECSCW. Kluwer Academic Publishers, Norwell, MA, 335-354.
- Celma, Ò. and Herrera, P. 2008. A new approach to evaluating novel recommendations. In *Proceedings of the 2008 ACM Conference on Recommender Systems* (Lausanne, Switzerland, October 23 - 25, 2008). RecSys '08. ACM, New York, NY, 179-186.
- Clements, M., Serdyukov, P., de Vries, A., Reinders, M. 2010. Personalized Travel Recommendation based on Location Co-occurrence. In *ACM Transactions on Information Systems* 2010.
- Crandall, D. J., Backstrom, L., Huttenlocher, D., and Kleinberg, J. 2009. Mapping the world's photos. In *Proceedings of the 18th international Conference on World Wide Web* (Madrid, Spain, April 20 - 24, 2009). WWW '09. ACM, New York, NY, 761-770.
- El Ali, A., Nack, F. & Hardman, L. (accepted). A location is worth a thousand experiences: Design implications for location-based multimedia messaging. To be published in *Proceedings of the 12th international conference on Multimodal Interfaces*, 2010, Beijing, China.
- FAA 2008. Federal Aviation Administration, Operational Concept Validation Strategy Document, 2008.
- Garcia-Crespo, A., Chamizo, J., Rivera, I., Mencke, M., Colomo-Palacios, R., Gomez-Berbis, J. M. 2009. SPETA: Social pervasive e-Tourism advisor. In *Telematics and Informatics*. 26, 3. August 2009, 306-315.
- Hansen, F. A., Bouvin, N. O., Christensen, B. G., Grønbaek, K., Pedersen, T. B., and Gagach, J. 2004. Integrating the web and the world: contextual trails on the move. In *Proceedings of the Fifteenth ACM Conference on Hypertext and Hypermedia* (Santa Cruz, CA, USA, August 09 - 13, 2004). HYPERTEXT '04. ACM, New York, NY, 98-107.
- Hao, Q., Cai, R., Wang, C., Xiao, R., Yang, J., Pang, Y., and Zhang, L. 2010. Equip tourists with knowledge mined from travelogues. In *Proceedings of the 19th international Conference on World Wide Web* (Raleigh, North Carolina, USA, April 26 - 30, 2010). WWW '10. ACM, New York, NY, 401-410.
- Hinze, A. and Junmanee, S. 2006. Advanced Recommendation Models for Mobile Tourist Information. *On The Move To Meaningful Internet Systems 2006: COOPIS, DOA, GADA, AND ODBASE. Lecture Notes in Computer Science*, 2006, Volume 4275/2006, 643-660.

- Kinsella, S., Budura, A., Skobeltsyn, G., Michel, S., Breslin, J. G., and Aberer, K. 2008. From Web 1.0 to Web 2.0 and back -: how did your grandma use to tag?. In *Proceeding of the 10th ACM Workshop on Web information and Data Management* (Napa Valley, California, USA, October 30 - 30, 2008). WIDM '08. ACM, New York, NY, 79-86.
- Lee, T., Crompton, J. Measuring Novelty Seeking in Tourism. In *Annals of Tourism Research* 19, 4. 1992, 732-751.
- Marchionini, G. 2006. Exploratory search: from finding to understanding. *Commun. ACM* 49, 4 (Apr. 2006), 41-46.
- Marlow, C., Naaman, M., Boyd, D., and Davis, M. 2006. HT06, tagging paper, taxonomy, Flickr, academic article, to read. In *Proceedings of the Seventeenth Conference on Hypertext and Hypermedia* (Odense, Denmark, August 22 - 25, 2006). HYPERTEXT '06. ACM, New York, NY, 31-40.
- McCurdy, M., Connors, C., Pyrzak, G., Kanefsky, B., and Vera, A. 2006. Breaking the fidelity barrier: an examination of our current characterization of prototypes and an example of a mixed-fidelity success. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Montréal, Québec, Canada, April 22 - 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI '06. ACM, New York, NY, 1233-1242.
- Neubauer, N. and Obermayer, K. 2009. Towards Community Detection in K-partite K-uniform Hypergraphs. *Workshop on Analyzing Networks and Learning with Graphs, NIPS 2009*.
- Pu, P. and Chen, L. 2008. User-Involved Preference Elicitation for Product Search and Recommender Systems. *AIMag.* 28, 4 (Winter 2008), 93-103.
- Raaijmakers, S. and Kraaij, W. 2010. Classifier Calibration for Multi-Domain Sentiment Classification. In *Proceedings of ICWSM'10* (Washington DC, May 2010).
- Ramzan, N., Larson, M. Dufaux, K. and Cluever, K. 2010. The Participation Payoff: Challenges and Opportunities for Multimedia Access in Networked Communities. *MIR 2010*.
- Serdyukov, P., Murdock, V., and van Zwol, R. 2009. Placing flickr photos on a map. In *Proceedings of the 32nd international ACM SIGIR Conference on Research and Development in information Retrieval* (Boston, MA, USA, July 19 - 23, 2009). SIGIR '09. ACM, New York, NY, 484-491.
- Toyama, K., Logan, R., and Roseway, A. 2003. Geographic location tags on digital images. In *Proceedings of the Eleventh ACM international Conference on Multimedia* (Berkeley, CA, USA, November 02 - 08, 2003). MULTIMEDIA '03. ACM, New York, NY, 156-166.
- UNESCO. *Discussion Report of the Planning Meeting for 2008 International Conference on Creative Tourism*, Santa Fe, New Mexico, U.S.A. October 25-27, 2006.
- Wartena, C. and Brussee, R. 2008a. Instance-based Mapping between Thesauri and Folksonomies. In *Proceedings of the 7th ISWC 2008*, Karlsruhe, Germany, October 26-30, 2008, 356-370
- Wartena, C. and Brussee, R. 2008b. Topic Detection by Clustering Keywords. In *DEXA Workshops* 54-58. IEEE Computer Society, 2008.
- Wartena, C., Brussee, R. and Wibbels, M. 2009. Using tag Co-occurrence for Recommendation. In *ISDA*, 273-278, 2009.
- Wartena, C. 2010. Automatic classification of social tags: a case study. In *Proceedings of ECDL 2010*.
- Wu, K. C. and Chiu, D. K. 2008. Toward Tourist Service Integration and Personalization with Semantic Web Services: A Case Study in Hong Kong. In *Proceedings of the 2008 IEEE international Conference on E-Business Engineering* (October 22 - 24, 2008). ICEBE. IEEE Computer Society, Washington, DC, 687-692.
- Zheng, Y., Zhao, M., Song, Y., Adam, H., Buddemeier, U., Bissacco, A., Brucher, F., Chua, T., Neven, H., and Yagnik, J. 2009. Tour the world: a technical demonstration of a web-scale landmark recognition engine. In *Proceedings of the Seventeen ACM international Conference on Multimedia* (Beijing, China, October 19 - 24, 2009). MM '09. ACM, New York, NY, 961-962.
- Zheng, Y., Zhang, L., Xie, X., and Ma, W. 2009a. Mining interesting locations and travel sequences from GPS trajectories. In *Proceedings of the 18th international Conference on World Wide Web* (Madrid, Spain, April 20 - 24, 2009). WWW '09. ACM, New York, NY, 791-800.
- Zheng, Y., Zhang, L., Xie, X., and Ma, W. 2009b. Mining correlation between locations using human location history. In *Proceedings of the 17th ACM SIGSPATIAL international Conference on Advances in Geographic information Systems* (Seattle, Washington, November 04 - 06, 2009). GIS '09. ACM, New York, NY, 472-475.

Zheng, V. W., Zheng, Y., Xie, X., and Yang, Q. 2010. Collaborative location and activity recommendations with GPS history data. In *Proceedings of the 19th international Conference on World Wide Web* (Raleigh, North Carolina, USA, April 26 - 30, 2010). WWW '10. ACM, New York, NY, 1029-1038.

Part 2- Putting the User at the Center of the Design Process of Technology-driven Novel Applications

ABSTRACT

This part of the report describes the challenge faced by the designers of Near2me in order to combine technology-driven design with a process that would keep the user at its center, and to evaluate a design in its conceptual stage with a focus on usefulness rather than usability. Traditional User-centered Design practices were not useful for our purpose because they base their initial study and final evaluation in the comparison with other systems that allow the user to perform the same tasks, which are unavailable for novel applications. To lead and document the design we used a framework structured around the design judgments proposed by Nelson and Stolterman, designing the user rhetorical situation simultaneously with the application. To make sure that as the design progressed it still applied to real users, their goals and their circumstances, we conducted a preliminary user study to identify user motivations and used the latter to guide our decisions. An evaluation designed to assess usefulness as the recognition of the application to fulfill user motivations was carried out successfully, indicating that the methodology described here presents potential for further study with the objective of defining guidelines to the design of technology-driven novel applications with the user at the center of the design process.

1. THE PROBLEM

The conception of a novel technology-driven product or application through a design process with the user at its center is not a trivial and straightforward matter.

In the first place, traditional User-Centered practices start with an “analysis” phase²⁹ that consists of studying users to identify who will use the product, what they will use it for and in which context they will use it. The Usability Professionals Association lists as suggested activities for this phase: creation of user profiles, task analysis, documentation of user scenarios, and documentation of user performance requirements, among others. If answers to “*by whom, how and when/where* the system will be used” are expected to be found through research, this presupposes the existence of a system whose users, mode of operation and context of use closely resemble the one to be created and whose performance has somehow been found lacking and hence the new “performance requirements”. In other words, User-centered Design methodology seems to be oriented towards the *improvement* of existing systems rather than the creation of novel ones, since novel systems may not only allow users to carry out previously impossible-to-attain goals, but also dramatically change how and in which context users carry out their current goals by offering new and different possibilities to act.

In addition to the problem posed by the analysis phase, traditional evaluation of User-Centered designs focuses on “usability testing”. Indeed, the Usability Professionals Association not only suggests this as a method of evaluation to ultimately check the design against the objectives but also in between design iterations. Several authors have discussed the problem posed by traditional usability evaluation when dealing with novel designs. The following is a summary paraphrasing the main points found in the literature:

- Premature usability evaluation of early designs can eliminate promising ideas by focusing on what the problems are (usability bugs) and not on the benefits (usefulness). In early designs usefulness may overshadow usability, usability should be allowed to eventually come through gradual refinements in later stages (Greenberg, 2008)
- Traditional usability evaluation of inventions does not provide meaningful information about their cultural adoption over time. By restricting usability evaluation to a set of tasks and contexts, the experimenter may be missing the tasks and contexts for which the system may prove more useful (Greenberg, 2008)
- Complex systems generally do not yield to simple controlled experimentation. By only creating what a usability test can measure, we will produce only simplistic progress that is not necessarily meaningful (Olsen, 2007).

However, there are important advantages to a design process with the user at its center. User-centered design, with its focus on **the user, his goals and his circumstances** (the rhetorical situation), drastically increases the likelihood that a product or system will be accepted-by and acceptable-to users by including these premises as requirements in the design process; without it, the risk of investing budget and time developing non-viable innovations runs high. Designing according to Human Factors and with knowledge of the interaction affordances of new platforms, although necessary, is not a sufficient condition for user acceptance and acceptability, as illustrated by what is informally called “inventions ahead of their time”—products and systems

²⁹ http://www.usabilityprofessionals.org/usability_resources/about_usability/what_is_ucd.html

that are correctly designed from the Human Factors and technological perspective, hence their later success, but were inadequate for users and their circumstances at the time—and the large number of tools and machines that, although once useful, were later discarded because they no longer fitted any human purpose.

In conclusion, although the traditional User-Centered Design methodology does not favor the creation of innovation, its philosophy to keep the user in focus provides clear advantages with respect to user acceptance and acceptability. Therefore we suggest there is a need for the formalization of a new methodology that specifically deals with innovation while maintaining the user at the center of the design process. This project explores a new way to keep the focus on the user during the design of novel systems by maintaining the user-centered philosophy while revisiting the methodology.

2. BACKGROUND

As outlined in the previous section, several authors have addressed the apparent contradiction between scientific-like usability evaluation, suggested by *User-centered Design* practice, and the more creative and life-altering activity of *Design* (Greenberg, 2008, and Olsen, 2007). Most have done so from a constructively critical perspective, with the aim of initiating a dialogue within the HCI community; others have gone one step further, suggesting guidelines and methodologies to follow to tackle the problem in practice. The methodology presented in this report builds on the works of the latter.

Wolf et al. (2006) have used Nelson and Stolterman's (2003) description of the **design judgments that occur during the design process** as a “means of documenting the creation of artifacts” using “**the creative design process as a framework with knowledge of the user as input.**” Their process consists of documenting the successive iterations of their design in the form of the artifacts that they created to materialize the design at each stage (i.e. the prototypes), which were either used for reflection by designers or for evaluation by users. They place these prototypes in a timeline in which they indicate which prototype at which stage lead to each judgment they made (of the types described by Nelson and Stolterman). In their own assessment, by doing this, they were able to achieve a “rigorous creative design process in the context of a multidisciplinary team with multiple users, technological and organizational constraints.” The ability to make judgments that went beyond refinements coupled to the rigor in keeping track of these judgments and their origins guaranteed that “the design was not merely the product of local optimizations and thrashing, and that it converged and cohered to the myriad of requirements to which it had to be responsive.”

On a different note, Fallman (2003), who defines design as “devising a course of action aimed at changing existing situations into preferred ones” rather than a mere improvement, states how, in a problem that requires *design* of a solution as opposed to the plain *engineering* of a solution, user-related requirements have to be “designed” during the whole system's design process in a manner he describes as “the process of unfolding the problem setting/problem solving pair” in which “it makes sense to see the designer as being involved in a conversation—a dialogue—rather than in a structured and linear process [...] regardless of whether or not iterations are allowed.”

Our methodology unites and builds upon these two approaches.

3. DESIGN METHODOLOGY

This section presents the design methodology used during the development of the off-the-beaten-track travel destinations explorer and recommender presented in **Part 1** of this report.

The design of Near2me was prompted by the availability of new technologies within the PetaMedia Europeam Network of Excellence. A survey of the technologies being researched inside the Network set the technical background and potential of the design. An ethnographic study of the current use of existing technologies in the areas of expertise of PetaMedia was conducted to survey state of the art functionality offered in PetaMedia's field and its applications. Simultaneously, we conducted a general socio-cultural study of the latest general trends within society to see in which areas the functionality potentially made available by PetaMedia technologies could be of use. This study aimed, not at identifying behavioral trends, but motivational ones: what makes people behave in the way they do. During this process, very general trends came up as candidates:

- People are interested in first-hand information coming from peers.
- People look for authenticity in products, viewed as a meaningful connection to the source of the product.
- People feel entitled to demand personalized experiences/products (long tail growth)

Brainstorm seeded with information from PetaMedia technologies, the functionality that these technologies can provide, how these functionalities are being used right now—what is most valued by user and what is missing, and the recognition of these trends pointed to an off-the-beaten-track recommender and explorer of travel destinations fueled by crowd-sourced implicitly obtained information.

To reflect upon and document the design process we drew inspiration from Wolf et al. (2006), also taking Nelson and Stolterman's (2003) design judgments as a base to use the creative design process as framework. The following is a brief description of the ten design judgments presented by Nelson and Stolterman and used as framework during our design process:

1. Judgment of service: choosing the audience, determining for whom the design is aimed and how they will benefit from it.
2. Deliberated offhand judgment: judgments based on procedural knowledge that can be deliberately brought to the surface in the effort to apply it to a new—different—situation.
3. Appreciative judgment: what is considered to be background and what is to be paid attention as foreground.
4. Appearance judgment: determination of the style, nature and character of the system (not necessarily involving aesthetics).
5. Quality judgment: judgments made in the quest for excellence in the creation of things of beauty, sublimity and practicality.
6. Instrumental judgment: choice and mediation of means (which techniques, instruments and technologies to use).
7. Navigational judgment: situational choices made to allow the object of the design or the design process itself to survive or take advantage of circumstances.
8. Framing judgment: defining the space of design outcomes determining what is to be included and what is out of scope.
9. Compositional judgment: how to join the elements to create a whole that serves the design intention.
10. Core judgment: composite of meanings and values formed during the designer's experience of living that can be made known when the designer is pushed with "why" questions.

However, there are three points in which our approach differs from the one described by Vetting Wolf et al.

1. Nelson and Stolterman detail ten types of design judgments. However—it is unclear why—Wolf et al. only include six of these judgments in their framework. We include all ten judgments in our version of the framework.
2. In terms of documentation, Wolf et al. place, in their timeline of judgments, the different design artifacts that they developed during the design process from which each design judgment comes from. Some of the actual judgments are explained textually in the paper. We, on the other hand, opted on a slightly different representation: we place, in our timeline of judgments, the judgments themselves. This timeline is available as a digital appendix of this report.
3. And finally, Nelson and Stolterman describe the service judgment (the judgment of whom to serve, or selection of a target audience of users) as a *conscious* judgment. Again it is unclear why—Wolf et al. reinterpret this judgment as an *unconscious* judgment. **We think it is of major importance that this kind of judgments is not only conscious but also at the forefront of the design process, so service judgments can guide the designer in the other judgments or remind him/her of the necessity to assess how the other judgments will affect the composition of the target audience**, the advantages of keeping the user at the center being already explained in Section 1 of this report.

Of these three differences, we regard the last one as the most critical. In addition to Wolf et al. considering this judgment unconscious, the only documentation they make of a design judgment is at the very beginning in which they set their user requirements which remain unchanged through the whole design process. In the case of Near2me, the design being technology-driven and with the goal to create new possibilities for users rather than incremental progress in the users' current activities, instead, in line with Fallman (2003), this type of judgment was maintained dynamic and at the center of the design process: designing the "problem" (by defining Near2me's users, the goals they would be able to fulfill with Near2me, and the circumstances in which Near2me would be of use) simultaneously with the solution.

This methodology posed, however, a new challenge: how to "design" the problem, in terms of target users, their goals and their circumstances, while keeping it realistic (i.e. matching an existing problem). Near2me's objectives allowed for an open concept design with the only constraint that it makes use of PetaMedia-developed technologies in a context *meaningful* to users, allowing the application to modify people's tasks and the way in which they are accomplished, as long as the application remains *useful*. We defined "useful" as allowing the user to fulfill a goal in accordance with his/her *motivations*. We emphasized motivations, as the driving force behind Human behavior and as reasons tightly connected to Human identity, considering that even un-voiced goals, or goals that were not apparent to users, were worth aiding by the design of an application if they represented alternatives to users' current goals to fulfill the same motivations.

To be sure that the design did not go astray in terms of motivations, we conducted a preliminary user research³⁰ aimed at identifying the motivations of travelers who engage in Creative Tourism. These motivations were used consistently throughout the development of Near2me to guide each design judgment, to decide on the acceptance or discarding of each

³⁰ Part 1. Section 1 of this report.

aspect and component of the system. Finally, the evaluation methodology was focused on assessing if users could identify in the final design the opportunity to fulfill the motivations obtained in the preliminary user study.

4. EVALUATION METHODOLOGY

The purpose of the concept evaluation phase is to determine if the goal to **design of a novel application that makes use of PetaMedia-developed technologies in a context meaningful to users** has been met. In particular—assuming that the question of novelty has been proved through the design’s comparison to other systems of its kind³¹ and inclusion of PetaMedia-developed technologies has been also been made clear³² – it is the “*in a context meaningful to users*” part that needs experimental evaluation. In other words, the purpose of the evaluation is to determine how **useful** users find the system to be: if users can identify in the final design the opportunity to fulfill the motivations obtained in the preliminary user study

As mentioned in **Section 1. THE PROBLEM**, the traditional usability evaluation practices are not appropriate for the evaluation of a novel, technology-driven design in its conceptual stage as Near2me. Greenberg (2008) suggests various other evaluation methods for when traditional usability evaluation is not appropriate: “there are many other appropriate ways to validate one’s work. Examples include a design rationale, a vision of what could be, expected scenarios of use, reflections, case studies, participatory critique, and so on.”

In the case of Near2me, the design framework provided by Nelson and Stolterman’s (2003) design judgments provides a design rationale for the concept. Each design judgment can be followed through the timeline presented as a digital appendix on this report, resulting in a structured process open to the inspection and critique of the multidisciplinary stakeholders in the project.

We successfully developed an expected scenario, which the participants of a user evaluation regarded credible and conducive to the understanding of the usefulness of the concept. Moreover, all participants to the study were able to formulate their own scenarios for their use of Near2me, based on real past travels and future travel plans.

And finally, we successfully carried out a participatory critique, in the form of user semi-structured interviews aimed at:

1. Identifying the aspects of Near2me that users find more salient (positive and negative).
2. Find if users would use Near2me and if there are any reasons or circumstances that would prevent users from using Near2me.
3. Evaluating if potential users are able to appropriate the concept, thinking of custom modifications or additions and suggesting novel uses for it.
4. To research how the concept behind Near2me compares, in the mind of its target users, to other applications of its kind (i.e. determine if users see distinctive value in the proposal of Near2me).

The results of the latter can be found in **Part 1. Sub-section 7.2. Results and Conclusions** of this report.

5. CONCLUSIONS

The methodology for design of technology-driven novel applications presented in **Part 2** of this report has successfully led to the design of Near2me, this success being defined as ***the creation of a new application***

- ***that serves a purpose otherwise unattainable by users,***
- ***whose design arose from the intention to take advantage of the availability of new technology,***
- ***for which the opportunity to fulfill the user motivations identified in the preliminary user study has been recognized by participants in a user evaluation,***
- ***whose feasibility and design judgments have been apparent to and corroborated by the different stakeholders in the project.***

This success suggests potential in the further exploration of this methodology, which only thorough use can validate as valuable guidelines for the professional practice of the technology-driven design of novel applications with the user at the center of the design process.

³¹ Part 2. Sub-section 5.3 of this report.

³² Part 1. Section 6 of this report.

6. REFERENCES

- Greenberg, S. and Buxton, B. 2008. Usability evaluation considered harmful (some of the time). In *Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy, April 05 - 10, 2008). CHI '08. ACM, New York, NY, 111-120.
- Fallman, D. 2003. Design-oriented human-computer interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, Florida, USA, April 05 - 10, 2003). CHI '03. ACM, New York, NY, 225-232.
- Nelson, H. G. and Stolterman, E. 2003. *The Design Way: Intentional Change in an Unpredictable World: Foundations and Fundamentals of Design Competence*. Educational Technology Publications, Englewood Cliffs, New Jersey, 2003.
- Olsen, D. R. 2007. Evaluating user interface systems research. In *Proceedings of the 20th Annual ACM Symposium on User interface Software and Technology* (Newport, Rhode Island, USA, October 07 - 10, 2007). UIST '07. ACM, New York, NY, 251-258.
- Wolf, T. V., Rode, J. A., Sussman, J., and Kellogg, W. A. 2006. Dispelling "design" as the black art of CHI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Montréal, Québec, Canada, April 22 - 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI '06. ACM, New York, NY, 521-530.

APPENDIX 1: RECOMMENDATIONS FOR THE FIELD TRIALS

1. Tentative Research Questions

1. How do users rate the usefulness and experience of the triple-synergy-fueled traveler recommender compared to existing methods? (Is the triple synergy capable of creating value in this domain?)
2. Do users find place recommendations produced from picture geo-tag patterns appealing? (Is the analysis of the user's previous picture taking locations a suitable method to predict the user's preference for new spots?)
3. Do users find the place names accurate and relevant? (Can the analysis of the relationship between geotags and tags produce place names that users find semantically accurate?)
4. Do the pictures succeed at portraying the place from the users' perspectives of interest (are they focused but sufficiently diverse)? (Can textual tags complement geotag-driven image search to produce diverse results for which subject focus is maintained?)
5. Do users find the top-down navigation of tags intuitive, useful and satisfactory? (Can a tag hierarchy able to support top-down browsing be derived from tag analysis?)
6. Do users find the suggested method for searching using the image as the query intuitive and easy to use? (Are the object duplicate detection input requirements acceptable for users?)
7. Do users find the information obtained through the use of object duplicate detection related data relevant? (Can the data derived from the context around other images belonging to the same object be used to construct a unit of descriptive information for an image?)
8. Do users find recommended peers more appealing than control peers? (Does the subject-related recommendation of people through a network of subject-authoritative peers produce more appealing results than the addition of tag relevance and general popularity?)

2. User Evaluation Methodology

The user test will consist of a task-directed walkthrough of the interface, in which the user's actions and intentions are observed directly and by the think-aloud method, respectively, followed by semi-structured interviews.

The participants will be asked to perform one or several tasks, yet to be determined, that allow him/her to walk through the concept's task flow and interface in a complete manner. During this stage, a researcher will observe the participant and register his/her interaction with the system and most significant comments. The participant will also be video recorded.

After the walkthrough, information will be obtained from the participants through semi-structured interviews. A question framework based on the research questions will guide the interviews. This framework will be adapted for each participant according to his/her vocabulary and to allow the researcher to explore and confirm the participant's feedback. The researcher will also make use of the notes taken while observing the participant during the walkthrough in order to tailor and direct the interview. These interviews will also be video recorded.

These video sessions will be analyzed qualitatively. The pair of video sessions for each participant will be coded together, for qualitative analysis of the participants' behaviors, comments during the interviews, and also any within-participant patterns or correspondences found between particularities in the interaction with the system and user perception. The coding will include multiple researchers, working either simultaneously coding observations that arise from consensus or asynchronously through the merging of coincidental observations.

The test will include between 8 and 15 participants, which lies within the standard sample size for qualitative analysis of a concept in its early stages. The estimated analysis time is eight hours for each hour of recorded video. Given the rapid increase in workload with the addition of each additional participant, the final number of participants, provided it lies within the above mentioned range, can be adjusted taking into consideration time constraints during the evaluation phase.

There will be two iterations of this evaluation methodology. After the first one, the prototype and evaluation criteria will be refined, with the product of these modifications being fed into the following evaluation round.

Finally, given that the coding requires two researchers working simultaneously, and the duration of each of the sequential stages in the evaluation process (from finding participants belonging to the target user group, to analyzing the results), we estimate that two researchers should be able to carry the whole evaluation process in six weeks time.

3. Selection of the Participants

The participants will be selected according to two requirements:

1. Membership to the target user group: people who like to travel outside mainstream tourism, visiting places that match their personal interests.

2. Familiarity with the use of web applications (beyond very basic ones, like email or search engines).

As the majority of independent travelers who opt for non-mainstream destinations belong to the 25-40 age group, and this group also counts with the most people familiar with non-trivial web applications, it is expected that so will the participants of the test. Other factors (i.e. gender, occupation) will be randomized.

4. Evaluation Criteria

A successful evaluation will be one that results in the gathering of relevant statements aimed in the direction of answering the research questions. Being qualitative in nature, the evaluation is not expected to answer the research questions in a complete, conclusive way (as would be the case of a quantitative evaluation that aims to prove a metric is beyond a threshold or not). A positive result is defined as the acquisition of information relevant to answering the research questions that aids the researchers in making a net advancement towards finding answers through any of the following mechanisms:

1. Validation of the research question as pertaining to an issue relevant to users. The users, although their behaviors and opinions may not agree, repeatedly raise issues within the realm of the question.
2. Acquisition of information that leads to the discarding of a research question, the reformulation of the current questions, or the creation of new research questions that may complement the current ones.
3. Acquisition of information that leads to the narrowing down of the research question into a more specific formulation, making it easier and more direct to answer it in a subsequent iteration.
4. Indication of a trend towards one possible answer for the question, although with non conclusive results, which would contribute to the design of a task better suited to answer the question in following iterations.
5. Indication that there may be no unanimous conclusive answer for the research question, and that several behaviors and opinions on the part of users can be used to cluster them into groups to which the question can be addressed individually.
6. The obtainment of a qualitative answer to the research question.

And finally, another valuable result that would define the test as successful would be the opening of new lines of research, in the form of new research questions. These research questions could be of the type more testing on the system could answer, but also broader and define new areas in which PetaMedia may want to venture.

5. Expected Research Outcome of the Field Trials

The expected research outcomes of the Field Trials can be summarized in the following categories:

1. Insights leading to the improvement of Near2me as system for aiding the recommendation and exploration of off-the-beaten-track locations in a personalized way.
2. Insights leading to the improvement of Near2me as a testing platform for PetaMedia developed Triple Synergy-related technologies.
3. Better understanding of the mechanisms through which the triple synergy can be used for the recommendation and exploration of items depicted in media.
4. Knowledge of how users react to and appropriate the technologies incorporated into Near2me, which could suggest new applications for these technologies.
5. Discovery of new research directions in the field of the Triple Synergy worth pursuing by PetaMedia in the future.