DANIEL SCHUSTER, TU Dresden, Germany ALBERTO ROSI and MARCO MAMEI, Unimore, Italy THOMAS SPRINGER, TU Dresden, Germany MARKUS ENDLER, PUC Rio, Brazil FRANCO ZAMBONELLI, Unimore, Italy

As pervasive computing meets social networks, there is a fast growing research field called pervasive social computing. Applications in this area exploit the richness of information arising out of people using sensorequipped pervasive devices in their everyday life combined with intense use of different social networking services. We call this set of information pervasive social context. We provide a taxonomy to classify pervasive social context along the dimensions space, time, people, and information source (STiPI) as well as commenting on the type and reason for creating such context. A survey of recent research shows the applicability and usefulness of the taxonomy in classifying and assessing applications and systems in the area of pervasive social computing. Finally, we present some research challenges in this area and illustrate how they affect the systems being surveyed.

Categories and Subject Descriptors: C.2.2 [Computer-Communication Networks]: Network Protocols

General Terms: Design, Algorithms

Additional Key Words and Phrases: Pervasive computing, social networks, context awareness, taxonomy, survey

#### **ACM Reference Format:**

Schuster, D., Rosi, A., Mamei, M., Springer, T., Endler, M., and Zambonelli, F. 2013. Pervasive social context: Taxonomy and survey. ACM Trans. Intell. Syst. Technol. 4, 3, Article 46 (June 2013), 22 pages. DOI: http://dx.doi.org/10.1145/2483669.2483679

# 1. INTRODUCTION

The advent of social networking applications has radically changed the way we use the Internet. Millions of users nowadays regularly turn to Web sites such as Facebook, Twitter, LinkedIn, or Orkut to keep in touch with friends, read and post comments, join common interest groups, and publish thoughts, opinions, recommendations, and up-to-date

© 2013 ACM 2157-6904/2013/06-ART46 \$15.00 DOI: http://dx.doi.org/10.1145/2483669.2483679

This work has been partially supported by the EU-FP7-FET Proactive project SAPERE – Self-aware Pervasive Service Ecosystems, under contract no. 256873.

Alexandre Skyrme, Simon Dobson, Graeme Stevenson, and Juan Ye contributed to the research at the basis of this article.

Authors' addresses: D. Schuster (corresponding author), Computer Networks Group, Faculty of Computer Science, TU Dresden, 01062 Dresden, Germany; email: daniel.schuster@tu-dresden.de; A. Rosi and M. Mamei, Department of Science and Engineering Methods, University of Modena and Reggio Emilia, Reggio Emilia, Italy; T. Springer, Computer Networks Group, Faculty of Computer Science, Tu Dresden, 01062 Dresden, Germany; M. Endler, Departmento de Informatica, PUC-Rio, Marques de São Vicente 225, 22453-900 Rio de Janeiro, Brazil; F. Zamboneli, Department of Science and Engineering Methods, University of Modena and Reggio Emilia, Reggio Emilia, Italy.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

information about themselves. As more and more people join and use such social networking Web sites, virtual communities are fostered and online social interaction rises.

Meanwhile, the ever-growing popularity of smartphones has caused a comparable phenomenon in the physical world. Most of these devices are always online while users carry them around in their daily life. Given the many sensors in these phones, a more comprehensive capturing and distribution of the physical context of people becomes possible, comprising how people move and interact, meet other people, as well as descriptions of visited places and performed activities.

The prospect of coupling the context provided from the social network sphere with sensing capabilities of current mobile devices creates a deep information space which we call *pervasive social context*. Applications and services generating and using this type of context information offer new possibilities to effectively adapt to people's physical and social situation. This leads to (hopefully) more meaningful social interactions between them, be it in the virtual or physical world.

There is already a large number of applications and services in the field of pervasive social computing producing and/or consuming context from both social networks and pervasive sensors. Each of them provides its own view about context information from the application perspective. Looking at the characteristics of this information, there are many similarities leading to clusters of applications and services. Understanding these characteristics well is thus crucial to be able to fully realize and exploit their potential. We want to classify and differentiate this diverse set of information asking the popular 5WH questions: *who, what, where, when, why,* and *how* with regard to pervasive social context.

Some of these questions are inherently application dependent. They relate to what type of context information is used in a given application, and why it is used (i.e., related to the goal of the application). We elaborate on these two questions in Section 2, giving an overview of the diversity of information that may be part of the pervasive social context.

The application-generic aspects *where*, *when*, *who*, and *how* are the basis of our STiPI (Space, Time, People, Information source) taxonomy described in Section 3 which is the main contribution of this article. We apply the taxonomy by surveying recent research in Section 4, providing a number of example systems for each of the dimensions of the taxonomy. This further clarifies the exact meaning, proper use, and usefulness of each dimension of the taxonomy. We discuss open research challenges in Section 5 and the relation of our taxonomy to other surveys in Section 6. Finally, Section 7 presents some conclusions and an outlook to future work.

# 2. TOWARDS PERVASIVE SOCIAL CONTEXT

Before we further elaborate on our notion of pervasive social context, there needs to be some clarification of the terms used in this article.

### 2.1. Definition of Terms

We use the term Pervasive Social Computing in the sense introduced by Mokhtar and Capra [2009] and further defined by Zhou et al. [2010]. This comprises systems and applications combining functionality from the areas of pervasive computing and social computing, moving the focus from a user's physical environment (pervasive computing) to his/her social environment and from online interactions (social computing) to interactions with colocated people.

On the other hand, we are talking about different variants of context. The most popular definition of *context* in computing is given by Dey [2001], binding context to entities (persons, places, objects) relevant to the user and the application.

In social computing, the most important such entities are other people. Stressing on this issue, *social context* in computing is often used as a term commonly referring to the people, groups, and organizations with which an individual is interacting. There are some variations of this term: Adams et al. [2008] broaden it by adding important locations and activities while Groh et al. [2010] use social context only in the narrower sense of small time intervals and space regions. The latter definition is close to what we call situated social context in an earlier work [Endler et al. 2011].

In pervasive computing, we could equally phrase the term *pervasive context* (or individual context as it is called by Groh et al. [2010]), comprising the information characterizing the entities around a pervasive device. This information can be gathered using sensors for temperature, light, sound, and the like. As soon as these sensors can measure the location and proximity of other devices and are able to identify the person carrying the device, we can also sense information about people. Thus the distinction between social context and pervasive context disappears, leading to the term pervasive social context.

*Definition* 1. Pervasive social context of an individual is the set of information that arises out of direct or indirect interaction with people carrying sensor-equipped pervasive devices connected to the same social network service.

This comprises the explicit links, profile information, and activities of people within the social graph, the joint sensor information of the pervasive devices, as well as implicit information that can be inferred by combining the two.

In this definition we use the term *social network service* in the same way as Quan [2011] and other authors, that is, naming any service delivering social networking functionality. It is furthermore important to mention that Dey's definition of context is broad enough to still contain pervasive social context as a subset.

Discussing about gathering this set of context information, in Rosi et al. [2011] we used the term social sensing to refer to the process of getting information from social network services. In this vision a social network becomes similar to a sensor providing information about people and the environment. Other authors [Olguin and Pentland 2008; Eagle et al. 2009] use the term social sensing as a process where the sensors present in mobile devices are exploited to infer data about activities of people, Miluzzo et al. [2008] use the term people-centric sensing for the same meaning.

As we rather need distinctive terms to define how information is gathered (see Section 3.4), we use the terms *social mining* for mining social network services to gather information about other persons, and *pervasive sensing* for gathering information from sensors on pervasive devices.

Explicit examples of social mining include information extraction from Twitter posts, Facebook status updates, location check-ins on Foursquare, or pictures posted on Flickr. Examples for pervasive sensing are getting GPS, temperature, light, or sound measurements as well as proximity information of other devices measured via Bluetooth.

#### 2.2. What and Why: Type and Reason of Pervasive Social Context

Starting with our 5WH questions, the what and why of pervasive social context need to be elaborated. *What* refers to the type of context which is inherently application-specific. We do not try to build a classification for this aspect as this would only be a snapshot of current applications not covering future developments.

Table I mentions recent examples for the type of context, namely user activity, location, proximity, real-world events, famous places, preferences, etc. This list is by no means complete but just offers a better understanding about pervasive social context in practice.

Example context	System		
current user activity or motion pattern	[Grob et al. 2009]		
patterns of recently visited places	[Gupta et al. 2009]		
detection and classification of events of the real world	[Zhao et al. 2006]		
unusual crowding in physical locations	[Fujisaka et al. 2010]		
people that are both socially and physically proximate	[Mokhtar et al. 2009][Gupta et al. 2009][Mehmood et al. 2009]		
mining famous city landmarks from blogs for personalized tourist suggestions	[Ji et al. 2009]		
common preferences of a group of co-located users	[Gartrell 2008]		
social knowledge spaces (i.e. people, content, information, calendar, authorization rules)	[Zhdanova et al. 2008]		
interconnected graph of socially relevant information from different sources	[Toninelli et al. 2010]		
real-time updated contact information	[Mehmood et al. 2009]		
sharing of live videos of events	[Kokku et al. 2008]		
crowd-sourced sensing and sharing of the status of the physical environment	[Demirbas et al. 2010]		
awareness of closest, most meaningful contacts	[Lugano 2008] [Ankolekar et al. 2009]		

Table I. Examples of Pervasive Social Context

The type of context will be better illustrated by the many systems surveyed in Section 4. Within these systems, the most popular type of context is proximity of friends, friends of friends, or yet unknown individuals. Proximity information can further be used to infer more complex information such as common activities, interests, or visited places, frequent interaction, or social link patterns. On the other hand, there is lots of diverse context information that can be detected by crowdsourcing like weather, noise, unusual crowding of places, or traffic information.

Like the type of context, the reason (why) for producing and consuming context information differs from application to application. Often there is the goal to predict future behavior by past behavior like forecasting box-office revenues for movies from Twitter posts [Asur and Huberman 2010]. Another goal is to improve emergency response in disaster situations like earthquakes [Mendoza et al. 2010] or to improve communication at big events like scientific conferences [Van den Broeck et al. 2010]. In the area of traffic detection, Traffic AUS (http://itunes.apple.com/au/app/aus-traffic/) and Waze (http://world.waze.com/) propose social networks for car drivers, in which the data produced by drivers about the traffic situation can be exploited by other drivers with the goal to enable and improve real-time navigation.

Again, this only shows the diversity of reasons thus making it hard to assess these viewpoints within a taxonomy. Other aspects are more application generic and should be elaborated in the following.

#### 3. WHERE, WHEN, WHO, AND HOW: THE STIPI TAXONOMY

As already introduced, we try to discuss pervasive social context on the basis of the who, what, where, when, why, and how questions. In Section 2, we discussed the role of what and why in the context of specific application domains. In this section we will use the remaining questions to create a taxonomy to classify pervasive social context from a general (application-independent) perspective.

The goal of the taxonomy is twofold. On the one hand, we try to clearly explain what we mean by social interactions enriched with pervasive social context. On the other



Fig. 1. Space dimension.

hand, it will be the basis for organizing our survey of existing systems and approaches that will be presented in the next section.

We elaborated the following four-dimensional *definition space* for social interactions, where each specific definition is characterized by the following.

- (1) *Space* (S). This is the spatial extent in which pervasive social context is produced and accessed. This answers the *where* question at the basis of our taxonomy.
- (2) *Time (Ti).* This is the temporal extent in which pervasive social context is produced and accessed. This answers the *when* question at the basis of our taxonomy.
- (3) *People (P).* This is the target population that is described by the pervasive social context. This answers the *who* question at the basis of our taxonomy.
- (4) Information (I). This is the information source from where data about pervasive social context comes from, and where such data is put together. This answers the how question at the basis of our taxonomy.

We call it the "STiPI" taxonomy. The choice of these four key dimensions allows to capture many facets and aspects of social interactions and allows to orderly classify a large number of existing proposals trying to support this kind of interactive systems.

# 3.1. Space Dimension (S)

The S dimension determines where—physically—social links and interactions are consuming or producing context information. In particular, for each level of geographical distance among the peers, the spatial dimension defines the kind of relations they could establish (see Figure 1).

S1: Small scope. This mode of pervasive social context comprehends only colocated/nearby people,that is, people having the potential to interact in direct, faceto-face mode within close proximity. Examples are people participating in the same event like a conference or experiencing a common situation such as a car accident at some street. The idea of small scope applies also for people being at the same place at different points in time interacting via that place (e.g., by leaving messages/traces to be retrieved later on).

*S2: Medium scope.* From a quantitative perspective this mode distinguishes from S1 for the major extents of people interacting, referring to all citizens of a city, region, or country. From a qualitative perspective it encompasses people sharing the same social and geo-political aspects of life, for example, all inhabitants of a village struck by an earthquake, or all people working in the same building during a fire alarm.

D. Schuster et al.



Fig. 2. Time dimension.

*S3*: *Anywhere*. In this mode, the current place/location of the peers is of no importance to establish interaction, but people can interact with everyone.

# 3.2. Time Dimension (Ti)

The Ti dimension characterizes interactions between peers and the context information they produce from the perspective of when, in the temporal scale, they happen. More precisely, this dimension determines the temporal aspect of the social interactions, in terms of peer discoverability and maintenance of the social links or interactions, which are directly related to the granularity of the user activity, role, and group membership (see Figure 2).

T1: Short-term activity. This mode includes all people that meet (or have the potential to meet) during a short period of time (e.g., while waiting at the bus stop, by bumping into each other on the street, or being online at the same time in the same social application).

*T2: Mid-term activity.* This includes individuals who may interact for some time because of their mid-term goals, and typically differentiates from T1 for a priori planning of the activity (e.g., a scheduled work meeting).

T3: Long-term activity. This mode includes people that have a direct or indirect relationship due to their long-term activity, role, or participation in a common project or common interest. Examples include families, fellows and friends, club members, employees, etc.

# 3.3. People Dimension (P)

The P dimension expresses who are the players involved in situations, consuming and/or producing context information. In particular, from the "granularity", or extent, of peers involved, growing levels of social interactions arise (see Figure 3).

*P1: Individuals.* The social context of a person consists of individual persons (i.e., friends, friends of friends, or unknown individuals) he/she wants to interact with.

*P2: Groups.* The social context of a person consists of a group or multiple groups of people he/she wants to interact with. A group can be a cluster of friends from the social network, a group gathering at some place, or the like.

*P3: Anonymous community.* The social context of a person consists of an anonymous community with the same social application.



Fig. 3. People dimension.



Fig. 4. Information dimension.

# 3.4. Information Dimension (I)

The I dimension determines how—at information level—social interactions build and compose the pervasive social context. In particular, information could either come from the physical world (i.e., pervasive devices) or from the virtual world (i.e., social networking platforms), or from their composition and integration (see Figure 4).

*I1: Information from pervasive sensors.* This source provides pervasive social information by exploiting pervasive devices. These devices can monitor the activities of the user (and thus indirectly their "social" profile) and her interactions (and thus indirectly her "social" network). More in detail, this source encompasses off-the-shelf sensors possibly embedded in portable and mobile devices like modern smartphones, tablets, and notebooks. Such kind of pervasive devices—provided with sensing capabilities like GPS, Bluetooth, gyroscope, microphone, camera, etc.—observe people interacting with the physical world, recording coordinates for the places they visit, the MAC address of people inside Bluetooth or Wi-Fi range, or the pictures of people with whom they interact.

12: Information from social networks. This source provides pervasive social information by exploiting social networking platforms. Such kind of data can contain information about what the user is doing (e.g., Twitter), where he/she is (Foursquares), and also other contextual information [Sakaki et al. 2010]. More in detail, with the term social mining, as already stated in Section 2, we consider mining any source of information that can be identified in social networking services that expresses individual

		, ,	
Dimensions	1	2	3
(S)pace	Small	Medium	Anywhere
	scope	scope	
(Ti)me	Short-term	Mid-term	Long-term
	activity	activity	activity
(P)eople	Individuals	Groups	Anonymous
			Community
(I)nformation	From pervasive	From social	From sources
Source	sensors	networks	integration

Table II. The STiPI	Taxonomy	Summarized
---------------------	----------	------------

interests, preferences, and activities in the social sphere. Examples of this source could be exemplified in Twitter posts, Facebook status updates, pictures posted on Flickr, etc.

13: Information from integration. The prospect of coupling social networks with pervasive sensors makes it possible to provide pervasive applications with higher degrees of context awareness, further shortening the gap between digital and physical worlds, and thus providing users with more personalized services. Such integration typically follows three main directions [Rosi et al. 2011]. First, one can exploit social networks to extract social information and have such information feeding pervasive services and applications, thus bringing social networks to the same network level of pervasive sensors. Second, one can exploit social network services as a way to collect and organize the data coming from pervasive sensors. In this way, the social network infrastructure is assimilated to the role of middleware for pervasive services and applications, that is, as a socio-pervasive medium to distribute and fetch pervasive content and information. Third, one could think of exploiting existing social network infrastructures as the ground upon which to build an overlay to perform the integration. Extracting information from both pervasive-services and social networks and bridging them in application-specific social network overlays accounts for the specific needs of applications and services.

# 3.5. Summary

The STiPI taxonomy, summarized in Table II, helps us to more precisely define *perva*sive social context. This term can be used for any system that makes use of both social and physical information (I1, I2, or I3) to define a pervasive social context enabling interaction with individuals, groups, or anonymous communities (P1, P2, or P3), in small, medium, or large scope regarding location of participants (S1, S2, or S3) within short-term, mid-term, or long-term activities (T1, T2, or T3).

# 4. APPLYING THE TAXONOMY

To more clearly show the usefulness of the STiPI taxonomy, we present a short survey of pervasive social computing applications ordered according to the four introduced dimensions. This further describes how these applications make use and/or produce pervasive social context.

# 4.1. Space Dimension

In this dimension, most research can be categorized into those that support only smallscope (S1) proximity detection, the ones that focus exclusively on medium scope (S2), and some works that support both small- and medium-scope colocation detection. There are also some approaches dealing with (S3), but given their global scale, they tend to be more Web applications rather than pervasive computing systems.

In the first category (S1), some of the most prominent systems/approaches are: SAMOA, MobiClique, VENETA, and WhozThat. SAMOA [Bottazzi et al. 2007] is a middleware system enabling inference of previously unknown social patterns by some semantic analysis that tries to match activities/attributes of users and places profiles. The goal is to determine potential peers for social interaction based on these profiles and physical small-scope colocation. MobiClique [Pietiläinen et al. 2009] employs a decentralized proximity detection, where geographical and social context is associated by means of a social networking service based on a store-carry-forward content dissemination. A similar decentralized approach is also adopted by VENETA [von Arb et al. 2008]. It is a mobile social network platform able to explore the social neighborhood of a user by detecting common friends of friends which are in the user's current physical proximity. This is done by Bluetooth-based proximity detection and by comparing their phone contact entries. Similarly, WhozThat [Beach et al. 2008] also uses phone proximity to exchange social network (e.g., Facebook) IDs, which are then used to fetch personal profiles from the social network.

Among the systems belonging to S2, one should mention PeopleNet [Motani et al. 2005], which is a P2P architecture for information search in a set of user-specific geographic locations, named bazaars. Because a bazaar can span a quite large geographical region, PeopleNet's notion of coloction is clearly of medium scope (S2). Another interesting work in this category is MobiSoc [Gupta et al. 2009], which is able to infer previously unknown social patterns by analyzing people profiles and their mobility traces. Colocation of users is verified on the server side, rather than employing direct proximity detection. Since the goal is to discover new social links based on the set of the user's commonly visited places, rather than particular places for a "here and now people discovery", the spatial dimension is S2. A similar spatial scope is also considered in the work on mobile social ecosystems [Toninelli et al. 2010].

Among the works on pervasive social context a few have a rather flexible definition of colocation, and thus span the S1 and S2 modes. One of them is Dodgeball [Ziv and Mulloth 2006], that was probably the first application combining LBS (Location-Based Service) and social networks. Dodgeball's goal was to share the user's location (as a symbolic place name) in the social network and to send text messages to friends and friends of friends within a "distance of up to ten blocks". Therefore, Dodgeball could be used both for small-scope and medium-scope interactions, depending on the size of the coverage set by the user. The other work with a spatially flexible approach is CenceMe [Miluzzo et al. 2008]. CenceMe is now a commercial social network available on the iPhone that is able to collect, classify, and infer the user's present status and activity from the mobile device's sensors and export this information, in realtime, into social networks. It also has a social context classifier running at a central back-end server that computes a user's neighborhood condition, that is, the CenceMe buddies in a user's surrounding area. Since it is aimed both at detecting new social ties based on similar visited places (like in Gupta et al. [2009]) and at identifying user proximity, it supports modes S1 and S2.

As already introduced, S3 systems are more in the realm of Web applications than pervasive computing. For example, Web collaborative filtering and recommendation tools [Linden et al. 2003] are S3 systems in that they take advantage of people in indirect interaction on a global scale. The information produced by these systems can be regarded as pervasive social context in that it describes the preferences of users in different scenarios.

#### 4.2. Time Dimension

With regard to the temporal dimension, many of the research systems we reviewed fall in category T1, that is, supporting short-term activities. A typical example can be

seen in PeopleTones [Li et al. 2008], where users of the service get an alert whenever a friend or buddy is in close proximity. The goal is simply to be informed about such a "nice to know" situation and to be able to contact a person directly, enabling some kind of spontaneous interaction not possible before (e.g., to have a cup of coffee together right now).

Most of the systems with proximity detection can be classified in S1 and T1, as proximity is most often used to involve any here and now activity which is made possible by the proximity situation. So, many of the systems already introduced for the spatial dimension like MobiClique [Pietiläinen et al. 2009], SAMOA [Bottazzi et al. 2007], and VENETA [von Arb et al. 2008] fall in category T1.

FLORA [Kokku et al. 2008] is an example of T1 systems beyond proximity, where users collect information for real-time collaboration. One possible application scenario is the collection of traffic updates to alert users entering a region with a traffic jam. Further scenarios include measuring people density at public places, tracking lost people, or public transport support, for example, letting the bus driver know that a potential passenger is waiting close to a particular bus stop.

As the distinction between categories T1 and T2 is not too sharp, we only found some systems supporting mid-term activities as their main focus (T2). Matching of requests and offers within a region, like in PeopleNet [Motani et al. 2005], is a typical example of T2 functionality as this is a mid-term activity, with a typical time frame of some days or weeks. One user may want to sell his used car, while another user is looking for a car with somehow matching preferences. Other examples for T2 are the framework for mobile social ecosystems [Toninelli et al. 2010] and MOSS [Zhdanova et al. 2008].

Another important area of research is represented by life-logger- and diary-based systems [Bell and Gemmell 2009; Ferrari et al. 2011]. They are tools for recording, in a browsable and machine-processable format, the everyday activities and events of people, communities, objects, and places. Since they can potentially span the whole lifetime of an individual they belong to T3 category.

# 4.3. People Dimension

If we look at the people dimension, most of the systems we surveyed aim at individuals (P1), that is, the social context consists of concrete persons as the target of interaction. Within P1, the largest number of systems helps the user to contact people from his list of friends within a social network or phonebook. Examples are again proximity-based services, like PeopleTones [Li et al. 2008] and SAMOA [Bottazzi et al. 2007], or systems tracking friends' position and their social status, like Google Latitude [Google 2009] or CenceMe [Miluzzo et al. 2008]. The latter type of systems will only be used within a smaller circle of close friends as a lot of information is shared to create a high level of social awareness.

Friendlee [Ankolekar et al. 2009] also targets at close friends only, as it creates a list of contacts the user often communicates with by analyzing his phone behavior. It thus provides a possibility to filter out the "real" friends of a mobile user's phonebook.

If we extend the scope to friends and friends of friends (still within the P1 category), Dodgeball [Ziv and Mulloth 2006], which was previously described in this article, represents an early system enabling interaction based on proximity where location has to be shared manually. Friends and friends-of-friends get an update whenever they are close to each other. This is inferred centrally using symbolic locations (like names of restaurants) only. VENETA [von Arb et al. 2008], also previously described in this article, is an example for serverless friend-of-friend detection based on Bluetooth technology.

The last subgroup within P1 is the encounter of unknown individuals like in akaaki [aka-aki networks GmbH 2010]. Users of such application are interested in finding new acquaintances. If two users of the system are in close proximity, they exchange profile information via Bluetooth and are thus able to check if the other person might be interesting and to contact him or her directly.

There is also a number of systems where the interaction based on pervasive social context is not targeting at individuals but at groups of people (P2). Cluestr [Grob et al. 2009] focuses on the initial group formation process and proposes a list of contacts from the same group if the user selects a contact from his phonebook. To achieve this, the tool performs analysis on the social graph extracted from a social network like Facebook. The interaction will then take place between the user and the group he selected; that's why we chose to place Cluestr in the P2 category. Another typical P2 example is FLORA [Kokku et al. 2008], where groups of persons are built location dependent to enable real-time collaboration between the users. We already mentioned some of the envisioned scenarios earlier in this section. An additional interesting P2 system is Socialaware [Gartrell 2008]. It supports groups gathered at some place, for example, to create a playlist matching the preferences of the audience near a jukebox.

P3 systems, given their global scale, tend to be more oriented to Web applications. They consist in approaches to extract information from the global population of a social network. For example, Bonchi et al. [2011] survey different approaches trying to extract strategic information from the whole population of social networks.

# 4.4. Information Dimension

The I1 category is based on proposals that consider modern computing devices as enabling technologies for humans to interact with the surrounding world (e.g., to localize and get information about entities nearby). In most cases, works in literature for this category fall on S1, T1, and P1 since from the geographical position of each single user they provide real-time localization-based services built over a GPS position, a recognized MAC address, or a read RFID tag.

Literature related to this area is extremely wide, as the majority of works in context awareness for pervasive computing falls in this category. Among the many examples, we should mention NAVITIME [Arikawa et al. 2007], a navigation service that from user GPS positions incorporates various modes of transportation and guides users to their destinations via several types of transportation. Several works use RFID tags for the purpose of tracking systems [Kim et al. 2008; Min et al. 2007] and for providing context information. Similarly, Bluetooth and wireless technologies enable mobile users to enjoy proximity services [Tang 2009; Rashid et al. 2005] in a wide range of scenarios. Moreover, further mainstream applications for the I1 category consist in environmental monitoring through Wireless Sensor Networks (WSNs), from indoor to outdoor, from urban setting to rural and hostile ones [Werner-Allen et al. 2006; Lee et al. 2008].

Systems from the I2 category consider social networks as powerful tools to detect and predict collective patterns of behavior [Asur and Huberman 2010], possibly associated with events occurring in a circumscribed region from the real world [Mendoza et al. 2010; Van den Broeck et al. 2010] for a short- and mid-term time span: that puts such works in S2, T1 (or T2), and P2.

In Crandall et al. [2009], the authors present techniques to automatically identify the location of points of high interest all over the world, by analyzing the spatial distribution of millions of geo-tagged pictures posted on Flickr. Results accord with common-sense opinions and travel guide suggestions.

Similarly, in Zhao et al. [2006] authors propose a system for detecting and framing events from the real world by exploiting the tags supplied by users in Flickr photos. The temporal and spatial distributions of tag usage are analyzed and tags related to aperiodic events and those of periodic events are distinguished. Tags are finally clustered and, for each cluster, a representing picture and tag is extracted. Ji et al. [2009] report a work on mining famous city landmarks from blogs for personalized tourist suggestions. Their main contribution is a graph modeling framework to discover city landmarks by mining blog photo correlations with community supervision. From our side, we have developed an unattended system [Mamei et al. 2010] able to extract and take advantage of up-to-date and spontaneous information embedded with pictures. With experiments on the Flickr database, we have shown that this system, by learning from past tourist user experience, is able to make effective recommendations to people visiting tourist places for the first time.

In Sakaki et al. [2010], it is demonstrated that Twitter, thanks to its real-time nature, can effectively act as a seismometer for the detection of earthquakes, simply by observing user tweets. More generally, the ability to identify global trends and events via Twitter is the core of numerous applications, such as *Tweettronics* (http://www.tweettronics.com/), oriented to identify market trends and brand awareness for marketing purposes.

In Fujisaka et al. [2010] authors propose methods for the detection of unusual crowding in physical locations from existing blog and Twitter communities. Here authors, by the analysis of common patterns of occurrence in each region over a specified time period, achieve the extraction of useful and interesting movement patterns, reflecting the occurrence of critical events in a geographic region.

Facebook is often cited for studies on network evolution and peer (as nodes of a graph) behavior. Of the many examples in literature (as the already mentioned VENETA [von Arb et al. 2008] and Cluestr [Grob et al. 2009]) of services extracting social features from a network, we further mention Hui Yi and Hung Yuan [2010], which studies and analyzes the patterns of friend-making, and the work of Viswanath et al. [2009], which studies the dynamic properties of the friendship network. Both of these works have the potential for improving our understanding of the dynamics of real-world social networks (and therefore have potentially high commercial and social impact).

Concerning the I3 category, as already explained, there is a general tendency to integrate data from social networks with information from pervasive sensing to provide a more comprehensive context description.

In Olguin and Pentland [2008] authors present the design, implementation, and deployment of a wearable social platform that can measure and analyze personal and social behaviors in a variety of settings and applications. Individual and social patterns of behavior are identified measuring face-to-face interaction, conversational dynamics, physical proximity to other people, and physical activity levels.

The approach of integrating real-world data from face-to-face proximity, with identities in online social networks, has also been followed by Van den Broeck et al. [2010] who have developed an application for people attending scientific conferences. Personal profiles of the participants are automatically generated using several Web 2.0 systems and semantic data sources, and integrated in real time with face-to-face proximity relations detected via RFID badges.

Authors from Lovett et al. [2010] present two heuristic methods for data fusion that combine the user's personal calendar with social network posts, in order to produce a real-time multisensor interpretation of real-world events. This study shows that the calendar can be significantly improved as a sensor and indexer of real-world events through data fusion.

Another set of proposals in the I3 area considers social network infrastructures as a sort of socio-pervasive middleware in which to merge and consolidate data from different sources, specifically pervasive sensors, and from which to exploit the functionalities for data and event management. As from before, works in this category typically involve large region of analysis (S2), consider mid-term time span as weeks or months (T2), and aggregate people on groups (P2).

S-Sensors [Baqer and Kamal 2009] provides a framework to globally share locally measured sensory readings. Authors propose to employ microblogging to publish and share sensory data and resources, where short messages depicting the status of the environment are used to convey sensory data of the physical world.

Automated murmurs [Freyne et al. 2009] presents a mobile platform that leverages the popularity of mobile and social computing to produce a location-sensitive messaging system which delivers user-generated content to the public on the basis of both physical location and social relations.

Authors from Patterson et al. [2009] present a prototype system that automatically infers the users' place, activity, and availability from sensors on her handheld devices or laptop computers. Data is then reported to buddies through embedding information in commercial instant-messaging profile status.

In Demirbas et al. [2010] authors have designed and implemented a crowdsourced data mining and collaboration service over Twitter, for two application scenarios: a crowdsourced weather radar, and a participatory noise-mapping application. The whole system is based on the intuition of exploiting Twitter as a publish-subscribe system for the storing and the diffusion of information and events about pervasive sensors and user-provided data.

Finally, there are proposals related to interconnecting and sharing data sensed from personal devices with the rest of the world. Accordingly, overlays over existing social networks, as well as brand new application-specific networks, are realized to interface with such local networks and, to support specific application requirements, implement or extend existing functionality.

SenseFace [Rahman et al. 2010] is a software overlay suitable for capturing the sensory data produced from user personal devices, processing and storing the sensory data in his/her personal gateway (which is a mobile device), and sending the data to a remote Internet gateway. Finally, the sensory data is disseminated to a list of his/her social networks.

In Anwar et al. [2005] authors propose an overlay constructed on top of the Orkut social network. Their aim is to demonstrate an alternative model to query the social network, where each node chooses its peers to query using metrics that can account for data coming from pervasive sensors. In authors' opinion such an approach not only improves the overall search time but also gives a sizable improvement in lookups, average round-trip delay, and scalability.

#### 4.5. Discussion

In summary, Figure 5 illustrates how the presented systems fit in the STiPI taxonomy. As already discussed, it can be noticed that most of the presented works focus on the first two values of the taxonomy. Systems in the (S3, T3, P3) part of the taxonomy, dealing with global aspects, lose the notion of "pervasive social" to some degree.

It is also rather natural to see a general correlation among values in the four dimensions, so a lot of systems are classified in (S1, T1, P1) and (S2, T2, P2). Again this is rather simple to explain: the more the spatial scope, the more people and time can be considered. The (I) dimension, representing the information source on which systems are built, tends to be more orthogonal to the other dimensions. For example, profile information coming from social networks can be used both in confined environments (S1, T1, P1) for the sake of better contextualizing information and services (e.g., Van den Broeck et al. [2010]) or in large-scale scenarios to identify areas of interest (e.g., Mamei et al. [2010]).

In conclusion, the STiPI taxonomy can effectively classify a wealth of works in the area. The chosen dimensions allow also to identify the general application scenario (e.g., large or small scale) that is the target for a given system.

Dimensions	1	2	3
(S)pace	WhozThat [Beach et al. 2008] MobiClique [Pietilainen et al. 2009] SAMOA [Bottazzi et al. 2007] VENETA [von Arb et al. 2008]	MobiSoc [Gupta et al. 2009] Mobile Social Ecosystems [Toninelli et al. 2010] Dodgeball [Ziv and Mulloth 2006] CenceMe [Miluzzo et al. 2010]	Web Collaborative Filtering [Linden et al. 2003]
(Ti)me	PeopleTones [Li et al. 2008]	PeopleNet [Motani et al. 2005] MOSS [Zhdanova et al. 2008]	life-logger [Bell and Gemmell 2009; Ferrari et al. 2011].
(P)eople	Google Latitude [Google 2009] Friendlee [Ankolekar et al. 2009]	Cluestr [Grob et al. 2009] Flora [Kokku et al. 2008] SenseFace [Rahman et al. 2010] Socialware [Gartrell 2008]	Web Data Mining [Bonchi et al. 2011]
(I)nformation Source	NAVITIME [Arikawa et al. 2007] RFID [Kim et al. 2008, Min et al. 2007] Bluetooth & Wireless [Tang 2009, Rashid et al. 2005] WSN [Werner-Allen et. Al 2006, Lee et al. 2008]	[Asur and Huberman 2010] [Mendoza et al. 2010] Flickr, [Crandall et al. 2009] Flickr [Zhao et al. 2006, Mamei et al. 2010] Blogs [Ji et al. 2009] Twitter, [Sakaki et al. 2010] Twitter [Fujisaka et al. 2010] [Viswanath et al. 2009] [Yi et al. 2010]	S-Sensors [Baqer and Kamal 2009] Automated Murmurs [Freyne et al. 2009] [Patterson et al. 2009] Twitter [Demirbas et al. 2010] SenseFace [Rahman et al. 2010] Orkut [Anwar et al. 2005] [Van den Broeck et al. 2010] [Olguin and Pentland 2008] Calendar [Lovett et al. 2010]

Fig. 5. Applying the STiPI taxonomy. Surveyed systems are places in the four STiPI dimensions.

# 5. RESEARCH CHALLENGES

Many of the approaches towards the definition of a pervasive social context have been worked out in research projects and have been tested only with simple prototype systems (e.g., Ziv and Mulloth [2006], aka-aki networks GmbH [2010], and Miluzzo et al. [2008]). In spite of the several appealing functions and characteristics embodied in works from literature, it is clear that creating pervasive social context, taking advantage of a seamless integration between sources of information, still has several open issues and challenges [Mehmood et al. 2009].

As a matter of fact, the use of information from the social networks, on the one hand gives access to a wide panorama of information from the social sphere. On the other hand, it challenges researchers to deal with information sources that: (*i*) produce data often in free text with no structure nor codified semantics; (*ii*) don't provide a guarantee on the delivery of specific information about specific facts and at specific times by social networks; (*iii*) are completely out of the control loop of system managers and application developers.

In this section, we discuss the research challenges behind the realization of a pervasive social context. While some of these challenges are already present in pervasive computing and context-awareness literature, they tend to be exacerbated with the integration of social networking platforms.

#### 5.1. Sensor Uncertainty

pervasive social context relies on information obtained from both a mobile device's sensors and the social sphere. However, due to cost effectiveness and limited resources, many of the sensors found in today's mobile devices produce data that has low resolution, limited precision, and lacks accuracy (i.e., non-up-to-dateness) [Neisse et al. 2008]. Social data further exacerbates the problems, since it is by nature incomplete, inaccurate, context specific, and dynamic. A typical example is user profile data, which may be partial (e.g., user does not want to disclose some information), may be vague (e.g., user mentions only general interest in sports), and may be frequently changed according to the situation at hand (e.g., user changes his settings to achieve a specific goal). These problems are further amplified when single-source data (e.g., only data from an accelerometer) is aggregated to derive higher-level information (e.g., analyzing accelerometer data to infer the user's mobility pattern [Miluzzo et al. 2008]).

A solution comes from sources integration, where physical sensors and social sensors complete each other to resolve uncertainty. Consider, as an example, a typical localization problem, where a GPS fails to localize places that are indoor or don't have a clear sky view (e.g., the "urban canyon" problem). Here precision could be raised mining on the social sphere, where a tweet from Twitter, a post from Facebook, etc., could place somebody in a precise context (e.g., a Facebook post saying: "Me and Marco watching a football match at Irish Pub").

Despite the recent efforts to reduce the uncertainty of sensed information through information integration, much work is still required in the development of reliable and resource-efficient inference methods which can produce trustworthy pervasive sensor information.

This kind of challenge mainly applies to systems relying on pervasive sensor information (I1) and systems trying to compensate the lack of sensors with data coming from social networks (I1–I3). Similarly, it mainly applies to systems in the S1, T1, and P1 area of the STiPI taxonomy, in that in these systems the focus is exactly to make sense of the available information and compensate sensor uncertainty.

#### 5.2. Sensor Fusion

More sophisticated data mining approaches are needed in order to allow data from different sources to be fused together, or rather, combined and cross-checked, to allow trustworthy information inference. Thus, it seems that there is a strong demand for new and more powerful inference techniques and data mining methods [Memon et al. 2010], to better identify patterns and correlate data sensed from socially related users, in order to deduce, for example, a user's immediate information needs.

From our analysis, the majority of works in literature recognize the potential benefits of cross-analyzing pervasive sensor data with information from social networks. However, most of the existing inference, correlation, or matchmaking approaches so far are quite limited and concentrate at a few purposes: discovery/suggestion of new potential social ties, recommendation of places to visit, detection of social events/crowding, inference of user (group) motion patterns, and sharing data about the physical environment.

The key problem of all the preceding research is that they are often conducted as stand-alone data mining exercises and seldom exploited in a synergetic way. Instead, the real challenge is to integrate several techniques and have them cooperate in their inference activities. The point is that the knowledge that can be inferred from a multitude of different sensors can be more expressive and high level (other than overall more correct) than that inferred by independently working sensors. For instance, if a 46:16

body-worn accelerometer detects a "cyclin" activity and the geocoder detects the user is at home, the activity can be more properly corrected as "using a fitness cycle".

In any case, before social networks can widely contribute to the mining process, they have to reach a critical mass of data across many (e.g., spatial) characteristics. For instance, in our experience on Flickr [Mamei et al. 2010], only a restricted number of cities in the world already have enough information to make our tourist recommendation tool applicable.

This kind of challenge mainly applies to S2, T2, P2, and I3 systems. In such systems the goal is exactly to combine data from multiple sources (multiple people locally interacting both in space and time) to acquire a more comprehensive view of the situation.

# 5.3. Unified Data Representation and Interpretation

An important challenge is the construction of a shared vocabulary: a description of the type of objects and/or concepts that exist in the application domain, and their properties and relations. The vocabulary should be able to cover a wide range of concepts (events, locations, dates, etc.). This representation should avoid complex and highly structured formats that would be difficult to be encoded and maintained.

An interesting research direction consists of the use of pragmatic (i.e., tag-based) ontologies to encode such diverse information. This comprises both an effective creation of such descriptions and an effective use by applications [Robu et al. 2009]. Their integration with shared vocabulary represents a challenge for future research. Ultimately, this problem boils down to natural language processing and it is even more exacerbated by the peculiar (and evolving) languages used in social network sites.

In addition, to better interpret such complex data, visualization techniques and tools should be developed. Data visualization tools can be the user interface to certain applications, and they could become a core asset to see and understand the data produced by social networks and pervasive sensors at multiple levels of granularity. Works like Situvis [Clear et al. 2010] represent an important step in this direction.

This kind of challenge mainly applies to I2 systems. When dealing with information coming from social networks, the challenge is to construct representation formats to understand and classify information correctly.

#### 5.4. Privacy Management

Privacy-related challenges are cross-cutting concerns that impact on all the aforesaid challenges (for example, one will not have data-critical mass if users are uncomfortable in sharing it). This is a general problem related to pervasive computing that is even more important when using it to leverage social interactions in a pervasive social context. Despite the fact that people consciously agree to both providing their personal data and to being tracked by social network services (e.g., many social networks utilize localization services), new rules for respecting and preserving overall user privacy have to be formulated. In Chen and Rahman [2008] the authors investigate current mobile social networks and identify their weaknesses and strengths. From another side, aggregating and anonymizing data can provide a useful rough solution [Shi et al. 2010], but more investigation is needed to sort out privacy intricacies of future pervasive applications.

Our idea is that privacy management should apply at different levels of content information, from single personal contextual information, to aggregated one, to inferred before. At the level of raw data, privacy management requires a user to be conscious of what data he really wants to share, in the presence of multiple and dynamic data sources, and reaching such a level of consciousness is far from being easy. In addition, since the pervasive application will be able to infer and aggregate information, the issue of understanding what information can be produced at the various levels of the privacy, and controlling how to share it, arises. For instance, while sharing blurred GPS



Fig. 6. Challenges in pervasive social context and their relationship with the STiPI taxonomy.

coordinates may be fine (Marco was near latX,lonY at 10.00am) sharing the inferred place may be not (e.g., Marco was at "The Fox Pub" at 10.00 a.m.). Furthermore, even when some inferred information can be shared (e.g., Marco was at "The Fox Pub" at 9.30 p.m.), some aggregated information should not be (Marco was at "The Fox" at 9.30pm dancing with Anna). Vice versa, it is possible to think that raw data should not be shared (a biosensor capable of identifying health problems and sensing bio parameters) and that can be nevertheless shared in some inferred forms (if the same biosensor is used only to infer happiness levels). Summarizing, these examples, and many others one can think of, show that the issue of defining proper privacy policies for information in the pervasive social context is a complex open issue, involving: (i) the different levels at which data is processed, (ii) the definition of privacy boundaries reflecting tension between the conflicting goals of enabling services while keeping confidentiality, (iii) human values to avoid unethical situations and bias.

This kind of challenge mainly applies to P1 systems. When dealing with information about individuals, privacy issues are extremely relevant. When dealing with groups and anonymous communities, privacy concerns are less severe.

Figure 6 summarizes these challenges in relationship with the STiPI taxonomy. Of course, we do not define sharp boundaries but rather zones of influence.

# 6. RELATED WORK AND SURVEYS

Our work is not the first attempt to classify systems combining a mobile device's sensor data and social information to support new means of social interactions. In the following, we review related taxonomies and surveys in the areas of social signal processing, social networking, pervasive sensing and social mining, and pervasive social computing.

# 6.1. Social Signal Processing

Earlier attempts started within the area of Social Signal Processing (SSP), thus analyzing human behavior based on lower-level verbal and nonverbal cues like blinks, smiles, and voice with the goal to detect higher-level behavior context like activity level, engagement, or stress. Pentland [2005] provides an overview of SSP systems developed at MIT, defines different social signals, and already uses the term social context for this type of information. A taxonomy and survey of SSP [Vinciarelli et al. 2009] defines the most important social behaviors and the technologies needed to detect them. Zhou et al. [2010] add a generic architecture for SSP as part of their five facets of pervasive social computing. These works on SSP focus on the problem of how to transfer one part of what we call pervasive social context into another more meaningful part. They are thus complementing our survey more specifically defining special types of context and the concrete techniques to infer them.

# 6.2. Social Networking

The second group of related surveys can be found in the research area of social networking. In Jones and Grandhi [2005] the authors propose and describe a framework for classifying people-to-people-to-geographical places (P3) systems according to two dimensions, considering if they are synchronous or asynchronous, and people- or placecentered. While the former can be mapped to our time dimension, the latter rather concerns two different design techniques each involving both space and people dimensions. Thus place versus people-centered is a classification mean which can optionally be used in addition to our taxonomy.

Quan [2011] categorizes conventional social network services on a more technical level concerning services for identification and profile, social graph, social presence, and social interactions. Our work does not provide any such technical issues thus focusing on the characteristics of the wealth of context information that arises out of the combination of social networking and pervasive computing. This is complementary to Quan [2011] and other surveys providing a technical overview like the surveys on SSP mentioned in Section 6.1.

# 6.3. Pervasive Sensing and Social Mining

Regarding the field of pervasive sensing and social mining, in Aggrawal and Abdelzaher [2011] the authors provide a broad survey of recent research in the field of integrating sensors and social networks, presenting the main technological enablers, some modeling approaches, and several key challenges of this field, as well as some solutions. Unlike our work, they do not propose a taxonomy, nor attempt to classify current systems.

Building on this work, in our earlier work [Rosi et al. 2011] we classified and surveyed possible approaches of integrating pervasive sensing and social mining. This was further developed into the I dimension of the STiPI taxonomy as presented here.

# 6.4. Pervasive Social Computing

The more recent research trend of pervasive social computing already yielded a few surveys, focusing on different aspects of this area. A quite comprehensive survey is Zhou et al. [2010] defining five facets of pervasive social computing: physical environment awareness, behavior awareness, community awareness, interaction awareness, and content awareness. A generic architecture is presented for each of the five facets showing its technical realization. If we replace "awareness" with "context", this can be used as a good taxonomy to classify the *what* of pervasive social context information. As already mentioned in Section 2.2 using such type classification involves the danger of being obsoleted by new developments.

More targeted towards what we call pervasive social context, Groh et al. [2010] propose a general model of individual social situations as comprised of a four-tuple (T, S, P, C) consisting of a time span (start time and end time), a subset of threedimensional space (a polygon on a map), a set of participants, and tags that freely describe semantics (mainly activity) of the situation. This is related to our approach as it also identifies the *when*, *where*, and *who* to be part of the social context while it does not define any subcategories for these dimensions like in our work. Thus the T, S, and P aspects each concern specific time spans, areas, and people and not classes of timeliness, space locality, and users as in our taxonomy. The content tags further describe the type of context event to be logged. The situation model of Groh et al. [2010] adds quite well to our work, as it shows a way how to use and benefit from at least part of the dimensions defined in the STiPI taxonomy.

#### 46:18

In our earlier work [Endler et al. 2011] we define a taxonomy for charaterizing timeand space-constrained social context, which we call situated social context. Based on the introduced spatial, temporal, inference, and people dimensions the taxonomy was called STIP taxonomy. This taxonomy is the foundation of the STiPI taxonomy introduced in this article. The classes introduced in each dimension were slightly adjusted to the perspective on pervasive social context introduced in this article. The I dimension completely changed its meaning to information source dimension.

### 7. CONCLUSIONS

In this survey, we introduced pervasive social context as a new information-centered view on pervasive social computing. While in line with traditional definitions of Context as used in the pervasive computing community for years, it is a useful term to precisely name the plethora of information available to pervasive social applications.

Pervasive social context provides manifold information to stimulate, trigger, and maintain rich social interactions in pervasive social applications. Many systems already use pervasive sensing and/or social mining, may it be to enable interactions between people in close proximity, collect data collectively, or to feed social information to existing pervasive services.

The main contribution of this article is the STiPI taxonomy based on the four dimensions space, time, people, and information source. We used it to classify recent research in the area of pervasive social computing thus providing a broad overview of this fast emerging field. We presented main research challenges for providing pervasive social context thus showing the need for more research activities, especially targeted at providing and utilizing context information.

The taxonomy and survey presented here are useful tools to keep pace with the fast development in the area of pervasive social computing. What is still missing is a comprehensive technical survey or book analyzing, describing and clustering techniques used to create pervasive social context. As mentioned earlier, there are already taxonomies and surveys for certain aspects like social signal processing, pervasive sensing, or social mining. Our work, as well as the works mentioned in Section 6, could form the basis for a more comprehensive view on pervasive social computing.

The final goal would be to identify a general infrastructure to support applications producing and consuming pervasive social context. This infrastructure should be capable of efficiently gathering, storing, securing, and finding context information in a generic way, thus enabling any application to use it transparently.

#### ACKNOWLEDGMENTS

The very detailed comments of the anonymous reviewers helped a lot to improve the article and clarify terms and definitions.

#### REFERENCES

- ADAMS, B., PHUNG, D., AND VENKATESH, S. 2008. Sensing and using social context. ACM Trans. Multimedia Comput. Comm. Appl. 5, 11:1–11:27.
- AGGRAWAL, C. C. AND ABDELZAHER, T. 2011. Integrating sensors and social networks. In *Social Network Data Analytics*. Springer.
- AKA-AKI NETWORKS GMB H. 2010. aka-aki. http://www.aka-aki.com/.
- ANKOLEKAR, A., SZABO, G., LUON, Y., HUBERMAN, B. A., WILKINSON, D., AND WU, F. 2009. Friendlee: A mobile application for your social life. In Proceedings of the 11<sup>th</sup> International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI'09).
- ANWAR, Z., YURCIK, W., PANDEY, V., SHANKAR, A., GUPTA, I., AND CAMPBELL, R. H. 2005. Leveraging socialnetwork infrastructure to improve peer-to-peer overlay performance: Results from Orkut. The Computing Research Repository (CoRR) abs/cs/0509095.

ACM Transactions on Intelligent Systems and Technology, Vol. 4, No. 3, Article 46, Publication date: June 2013.

- ARIKAWA, M., KONOMI, S., AND OHNISHI, K. 2007. Navitime: Supporting pedestrian navigation in the real world. *IEEE Pervas. Comput.* 6, 3, 21–29.
- ASUR, S. AND HUBERMAN, B. A. 2010. Predicting the future with social media. In Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT'10). IEEE Computer Society, 492–499.
- BAQER, M. AND KAMAL, A. 2009. S-Sensors: Integrating physical world inputs with social networks using wireless sensor networks. In Proceedings of the 5<sup>th</sup> International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP'09). 213–218.
- BEACH, A., GARTRELL, M., AKKALA, S., ELSTON, J., KELLEY, J., NISHIMOTO, K., RAY, B., RAZGULIN, S., S UNDARESAN, K., SURENDAR, B., TERADA, M., AND HAN, R. 2008. WhozThat? Evolving an ecosystem for context-aware mobile social networks. *IEEE Netw.* 22, 4, 50–55.
- BELL, G. AND GEMMELL, J. 2009. Total Recall: How the E-Memory Revolution Will Change Everything. Dutton.
- BONCHI, F., CASTILLO, C., GIONIS, A., AND JAIMES, A. 2011. Social network analysis and mining for business applications. ACM Trans. Intell. Syst. Technol. 2, 3.
- BOTTAZZI, D., MONTANARI, R., AND TONINELLI, A. 2007. Context-aware middleware for anytime, any-where social networks. *IEEE Intell. Syst. 22*, 23–32.
- CHEN, G. AND RAHMAN, F. 2008. Analyzing privacy designs of mobile social networking applications. In Proceedings of the IEEE/IFIP International Conference on Embedded and Ubiquitous Computing (EUC'08).
- CLEAR, A. K., HOLLAND, T., DOBSON, S., QUIGLEY, A., SHANNON, R., AND NIXON, P. 2010. Situvis: A sensor data analysis and abstraction tool for pervasive computing systems. *Pervas. Mob. Comput.* 6, 575–589.
- CRANDALL, D. J., BACKSTROM, L., H UTTENLOCHER, D., AND KLEINBERG, J. 2009. Mapping the world's photos. In Proceedings of the 18<sup>th</sup> International Conference on World Wide Web (WWW'09).
- DEMIRBAS, M., BAYIR, M. A., AKCORA, C. G., YILMAZ, Y. S., AND FERHATOSMANOGLU, H. 2010. Crowd-sourced sensing and collaboration using twitter. In *Proceedings of the IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM'10).*
- DEY, A. K. 2001. Understanding and using context. Personal Ubiquit. Comput. J. 5, 1, 4-7.
- EAGLE, N., PENTLAND, A. S., AND LAZER, D. 2009. Inferring friendship network structure by using mobile phone data. Proc. Nat. Acad. Sci. Unit. States Amer. 106, 36, 15274–15278.
- ENDLER, M., SKYRME, A., SCHUSTER, D., AND SPRINGER, T. 2011. Defining situated social context for pervasive social computing. In Proceedings of the 2<sup>nd</sup> IEEE Workshop on Pervasive Collaboration and Social Networking (PerCol'11). IEEE.
- FERRARI, L., MAMEI, M., AND ZAMBONELLI, F. 2011. All-about digital diaries: Opportunities and challenges. IEEE IT Prof. 13, 5, 37–43.
- FREYNE, J., BRENNAN, A. J., SMYTH, B., BYRNE, D., SMEATON, A. F., AND JONES, G. J. F. 2009. Automated murmurs: The social mobile tourist application. In Proceedings of the International Conference on Computational Science and Engineering (CSE'09).
- FUJISAKA, T., LEE, R., AND SUMIYA, K. 2010. Detection of unusually crowded places through micro-blogging sites. In Proceedings of the IEEE 24<sup>th</sup> International Conference on Advanced Information Networking and Applications Workshops (WAINA'10).
- GARTRELL, C. M. 2008. SocialAware: Context-aware multimedia presentation via mobile social networks. Masters thesis, Department of Computer Science, University of Colorado.
- GOOGLE. 2009. Google latitude. https://www.google.com/latitude/.
- GROB, R., KUHN, M., WATTENHOFER, R., AND WIRZ, M. 2009. Cluestr: Mobile social networking for enhanced group communication. In Proceedings of the ACM International Conference on Supporting Group Work (GROUP'09).
- GROH, G., LEHMANN, A., WANG, T., HUBER, S., AND HAMMER, F. 2010. Applications for social situation models. In Proceedings of the IADIS International Conference on Wireless Applications and Computing.
- GUPTA, A., KALRA, A., BOSTON, D., AND BORCEA, C. 2009. MobiSoC: A middleware for mobile social computing applications. *Mobile Netw. Appl.* 14, 35–52.
- HUIVI, H. AND HUNG-YUAN, P. 2010. Use behaviors and website experiences of facebook community. In *Proceedings of the International Conference on Electronics and Information Engineering (ICEIE)*.
- JI, R., XIE, X., YAO, H., AND MA, W.-Y. 2009. Mining city landmarks from blogs by graph modeling. In Proceedings of the 17<sup>th</sup> ACM International Conference on Multimedia (MM'09).
- JONES, Q. AND GRANDHI, S. A. 2005. P3 Systems: Putting the place back into social networks. IEEE Internet Comput. 9, 38–46.
- KIM, S., KO, D., AND AN, S. 2008. Geographical location based rfid tracking system. In Proceedings of the 9<sup>th</sup> IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM'08).

- KOKKU, R., SUNDARESAN, K., AND JIANG, G. 2008. Enabling location specific real-time mobile applications. In Proceedings of the 3<sup>rd</sup> International Workshop on Mobility in the Evolving Internet Architecture (MobiArch'08).
- LEE, H.-C., LIU, C.-J., YANG, J., HUANG, J.-T., FANG, Y.-M., LEE, B.-J., AND KING, C.-T. 2008. Using mobile wireless sensors for in-situ tracking of debris flows. In Proceedings of the 6<sup>th</sup> ACM Conference on Embedded Network Sensor Systems (SenSys'08).
- LI, K. A., SOHN, T. Y., HUANG, S., AND GRISWOLD, W. G. 2008. Peopletones: A system for the detection and notification of buddy proximity on mobile phones. In Proceedings of the 6<sup>th</sup> International Conference on Mobile Systems, Applications, and Services (MobiSys'08).
- LINDEN, G., SMITH, B., AND YORK, J. 2003. Amazon.com recommendations: Item-to-item collaborative filtering. IEEE Internet Comput. 7, 1, 76–80.
- LOVETT, T., O'NEILL, E., IRWIN, J., AND POLLINGTON, D. 2010. The calendar as a sensor: Analysis and improvement using data fusion with social networks and location. In *Proceedings of the 12<sup>th</sup>ACM International Conference on Ubiquitous Computing (Ubicomp'10).*
- MAMEI, M., ROSI, A., AND ZAMBONELLI, F. 2010. Automatic analysis of geotagged photos for intelligent tourist services. In Proceedings of the 6<sup>th</sup> International Conference on Intelligent Environments (IE'10).
- MEHMOOD, R., CERQUEIRA, E., PIESIEWICZ, R., AND CHLAMTAC, I. 2009. Challenges for mobile social networking applications. In *Lecture Notes of the Institute for Computer Sciences*, Social Informatics and Telecommunications Engineering Series, vol., 16. Springer, 275–285.
- MEMON, N., XU, J. J. J., HICKS, D. L. L., CHEN, H., XU, J. J., AND HICKS, D. L. 2010. Social network data mining: Research questions, techniques, and applications. Ann. Inf. Syst. 12, 1–7.
- MENDOZA, M., POBLETE, B., AND CASTILLO, C. 2010. Twitter under crisis: Can we trust what we RT? In Workshop on Social Media Analytics (SOMA'10).
- MILUZZO, E., LANE, N. D., FODOR, K., PETERSON, R., LU, H., MUSOLESI, M., EISENMAN, S. B., ZHENG, X., AND CAMPBELL, A. T. 2008. Sensing meets mobile social networks: The design, implementation and evaluation of the CenceMe application. In Proceedings of the 6<sup>th</sup> ACM Conference on Embedded Network Sensor Systems (SenSys'08).
- MIN, Z., WENFENG, L., ZHONGYUN, W., BIN, L., AND XIA, R. 2007. A rfid-based material tracking information system. In *Proceedings of the IEEE International Conference on Automation and Logistics (ICAL'07)*.
- MOKHTAR, S. B. AND CAPRA, L. 2009. From pervasive to social computing: Algorithms and deployments. In *Proceedings of the International Conference on Pervasive Services (ICPS'09).*
- MOTANI, M., SRINIVASAN, V., AND NUGGEHALLI, P. S. 2005. PeopleNet: Engineering a wireless virtual social network. In Proceedings of the 11<sup>th</sup> Annual International Conference on Mobile Computing and Networking (MobiCom'05).
- NEISSE, R., WEGDAM, M., AND VANSINDEREN, M. 2008. Trustworthiness and quality of context information. In Proceedings of the 9<sup>th</sup> International Conference for Young Computer Scientists (ICYCS'08).
- OLGUIN, D. O. AND PENTLAND, A. S. 2008. Social sensors for automatic data collection. In Proceedings of the 14<sup>th</sup>Americas Conference on Information Systems (AMCIS'08).
- PATTERSON, D. J., DING, X., KAUFMAN, S. J., LIU, K., AND ZALDIVAR, A. 2009. An ecosystem for learning and using sensor-driven IM status messages. *IEEE Pervas. Comput.* 8, 4, 42–49.
- PENTLAND, A. S. 2005. Socially aware computation and communication. IEEE Comput. 38, 3, 33-40.
- PIETILAINEN, A.-K., OLIVER, E., LEBRUN, J., VARGHESE, G., AND DIOT, C. 2009. MobiClique: Middleware for mobile social networking. In Proceedings of the 2<sup>nd</sup> ACM Workshop on Online Social Networks (WOSN'09).
- QUAN, H. 2011. Online social networks & amp; social network services: A technical survey. In Pervasive Communication Handbook, CRC Press, Boca Raton, FL. 4.
- RAHMAN, M., EL SADDIK, A., AND GUEAIEB, W. 2010. Building dynamic social network from sensory data feed. IEEE Trans. Instrument. Measur. 59, 5, 1327–1341.
- RASHID, O., COULTON, P., AND EDWARDS, R. 2005. Implementing location based information/advertising for existing mobile phone users in indoor/urban environments. In Proceedings of the International Conference on Mobile Business (ICMB'05).
- ROBU, V., HALPIN, H., AND SHEPHERD, H. 2009. Emergence of consensus and shared vocabularies in collaborative tagging systems. ACM Trans. Web 3, 14:1–14:34.
- ROSI, A., MAMEI, M., ZAMBONELLI, F., DOBSON, S., STEVENSON, G., AND YE, J. 2011. Social sensors and pervasive services: Approaches and perspectives. In Proceedings of the 2<sup>nd</sup> IEEE Workshop on Pervasive Collaboration and Social Networking (PerCol'11).
- SAKAKI, T., OKAZAKI, M., AND MATSUO, Y. 2010. Earthquake shakes twitter users: Real-time event detection by social sensors. In *Proceedings of the 19<sup>th</sup> International Conference on World Wide Web (WWW'10)*.

- SHI, J., ZHANG, R., LIU, Y., AND ZHANG, Y. 2010. Prisense: Privacy-preserving data aggregation in people-centric urban sensing systems. In Proceedings of the 29<sup>th</sup> IEEE Conference on Information Communications (INFOCOM'10).
- TANG, Y. 2009. Interactive pervasive wireless services over gprs and bluetooth. In Proceedings of the 8th International Conference on Dependable, Autonomic and Secure Computing (DASC'09).
- TONINELLI, A., PATHAK, A., SEYEDI, A., SPEICYS CARDOSO, R., AND ISSARNY, V. 2010. Middleware support for mobile social ecosystems. In Proceedings of the 2<sup>nd</sup> IEEE International Workshop on Middleware Engineering (ME'10).
- VAN DEN BROECK, W., CATTUTO, C., BARRAT, A., SZOMSZOR, M., CORRENDO, G., AND A LANI, H. 2010. The live social semantics application: A platform for integrating face-to-face presence with on-line social networking. In Proceedings of the 1<sup>st</sup> International Workshop on Communication, Collaboration and Social Networking in Pervasive Computing Environments (PerCol'10).
- VINCIARELLI, A., PANTIC, M., AND BOURLARD, H. 2009. Social signal processing: Survey of an emerging domain. Image Vis. Comput. 27, 12, 1743–1759.
- VISWANATH, B., MISLOVE, A., CHA, M., AND GUMMADI, K. P. 2009. On the evolution of user interaction in Facebook. In *Proceedings of the 2<sup>nd</sup> ACM Workshop on Online Social Networks (WOSN'09)*.
- VONARB, M., BADER, M., KUHN, M., AND WATTENHOFER, R. 2008. VENETA: Serverless friend-of-friend detection in mobile social networking. In Proceedings of the IEEE International Conference on Wireless and Mobile Computing, Networking and Communications (WiMOB'08).
- WERNER-ALLEN, G., LORINCZ, K., WELSH, M., MARCILLO, O., JOHNSON, J., RUIZ, M., AND LEES, J. 2006. Deploying a wireless sensor network on an active volcano. *IEEE Internet Comput.* 10, 18–25.
- ZHAO, Q., LIU, T.-Y., BHOWMICK, S. S., AND MA, W.-Y. 2006. Event detection from evolution of click-through data. In Proceedings of the 12<sup>th</sup> ACM International Conference on Knowledge Discovery and Data Mining (KDD'06).
- ZHDANOVA, A., DAVIES, M., JORNS, O., STAVROULAKI, V., DEMESTICHAS, P., G ONZALEZ, M., MOESSNER, K., CARREZ, F., RAJASEKARAN, H., AND IACONO, L. 2008. MOSS: Mobile social spaces. In Proceedings of the 1st International Workshop on Blending Physical and Digital Spaces on the Internet (OneSpace2008) at the 1<sup>st</sup> Future Internet Symposium (FIS'08).
- ZHOU, J., SUN, J., ATHUKORALA, K., AND WIJEKOON, D. 2010. Pervasive social computing: Augmenting five facets of human intelligence. In Proceedings of the Conference on Ubiquitous Intelligence & Computing and 7th International Conference on Autonomic & Trusted Computing (UIC/ATC). Xian, Shaanxi.
- ZIV, N. AND MULLOTH, B. 2006. An exploration on mobile social networking: Dodgeball as a case in point. In Proceedings of the International Conference on Mobile Business (ICMB'06).

Received June 2011; revised November 2011; accepted December 2011