

SELECTING TRUSTWORTHY CONTENT USING TAGS

Daniele Quercia, Licia Capra and Valentina Zanardi

Department of Computer Science, University College London, London, WC1E 6BT, U.K.

Keywords: Mobile Computing, Reputation Models, Tagging.

Abstract: Networked portable devices enable their users to easily create and share digital content (e.g., photos, videos). Hitherto, this serendipitous form of sharing has not happened. That may be because, for sharing content, mobile users have no choice but to go through the Internet. Users are thus in need of decentralised mechanisms for browsing location-based content. To realize such mechanisms, the following two questions must be answered first: how to select “relevant content”, by semantically matching user queries, and how to select “quality content” from the clutter generated by a potentially huge number of producers. We explore ways to answer these questions. We propose a combined approach that infers “relevance” by reasoning about the semantics emerging from the tags that users associate to content, and “quality” by running distributed trust models that recognize trustworthy producers.

1 INTRODUCTION

In recent years, two separate trends have been observed: first, the rapid evolution of mobile technology, with current portable devices having increased computing capabilities (e.g., processing power and memory availability) and richer sets of functionalities (e.g., digital cameras, MP3 players, GPS receivers); second, the transformation of the Internet user from consumer to producer of content. It will not be long before these two trends will converge, thus leading to people generating and sharing location-based content using their portable devices. They, for example, will attach texts or audio clips to a point of interest, to be played back by others who come along later.

Currently, websites offer location-based services by collecting and adding “geotags” (encoding spatial co-ordinates) into content collected on the spot. However, being fully centralized, current location-based services do not scale and are not open to innovation, as we shall discuss in Section 2.2.

We argue that, in order to enable the sharing of massive amounts of location-dependent information, that will be increasingly produced and carried by mobile devices, a decentralised content sharing platform will become necessary (Section 2). In order to make such platform an enabling technology for pervasive computing, the following challenges will have to be addressed first:

- *Finding Relevant Content.* Mobile users will need to be assisted when browsing location-based data, in order to filter out irrelevant information, and be presented only with content they are interested in. In this domain, users typically describe content using a folksonomy, rather than a pre-defined taxonomy. As a result, mechanisms that will retrieve content of interest, based on the dynamically learned tags semantics, will be called upon (Section 3);
- *Finding Quality Content.* The amount of information that matches a user’s query may be overwhelming. In order to give end users a good pervasive experience, content should be ranked so that, the more reputable the source that produced it, the higher up its ranking. Mechanisms to dynamically assess a user’s reputation in highly decentralised systems are thus required (Sections 4).

These mechanisms will have to be evaluated in terms of accuracy (i.e., do they give end users content they like?), coverage (i.e., are they capable of digging out relevant content from the clutter produced?) and robustness (i.e., do they protect users from malicious manipulations of the system?). Evaluating the effectiveness of algorithms is a fundamental step to drive future innovation, but it also represents a major challenge for pervasive computing, as we shall describe in Section 5.

2 A DIGITAL TAPESTRY

Simply moving can be tantamount to browsing and generating content. People move and leave their digital traces and, by doing so, they create an invisible tapestry of location-based content. “As individuals traverse an urban landscape, they simply infuse their path with a unique and detectable digital redolence. Similarly, fixed places or objects can also emit unique scents once they are digitally tagged” (Paulos and Goodman, 2004).

Mobile users collaboratively contribute to the creation of the tapestry by (in descending order of user intervention):

- Attaching notes (e.g., texts, audio clips, pictures) to a place (e.g., park, plaza, bus stop) or to an object (e.g., bench, bridge, parking slot) using their mobile phones. Those notes are read by others who come along later (Sharon, 2006).
- Wearing cyber goggles that tag everything they see in the course of a day (Harada et al., 2007). Researchers of Tokyo University have been studying how a pair of glasses that mount a tiny camera and LCD screen helps elderly’s memory. This pair of glasses records what the wearer sees and names objects in the field of view in real time. The wearer can then type in a keyword later on (e.g., ‘butterfly’), and the screen will playback the clip from the moment he saw the insect.
- Carrying their mobile phones. For example, the Dutch GPS-maker TomTom recently launched a new service, dubbed High Definition Traffic, that exploits the fact that drivers carry their mobile phones. More specifically, the service “tracks the paths of about 4 million mobile phone users to expand the amount of traffic information available” (Steen, 2007). That is a striking example of how a simple act of movement becomes, in the tapestry, an act of content creation.

2.1 Browsing the Tapestry: For What?

Apart from creating the tapestry, mobile users can also browse it, and they usually do so by issuing a query. More specifically, by either:

- *Specifying their likes and dislikes beforehand.* Their devices will then search for things they might find interesting on the way (e.g., old movies they have been willing to see, or popular hangouts for folks with their own inclinations).
- *Performing custom searches.* They do so whenever they are looking for something in particular at a certain time. For example, whenever drivers

are hungry, they can search for cheap and nearby restaurants.

More generally, mobile users can find several things of personal interest:

- *Songs of emerging musicians.* To get some free publicity, emerging artists upload their latest tracks into publicly-available WiFi hotspots and add the date of their next gig as a note to the track (Bassoli et al., 2007; L. McNamara, C. Mascolo and L. Capra, 2007).
- *Prices of outlets.* Instead of showing generic icons for restaurants and petrol stations, mobile maps can be fed with specific information - for example, outlets can embed their latest offerings or discounts or seasonal menus within their clickable logos displayed on the map. By simply looking up their maps, drivers can plan fill-ups or find cheap places to have lunch.
- *Street performances.* Whenever musicians put on impromptu street performances, they can inform people in their proximity by disseminating electronic flyers. By receiving flyers, people can make the most out of the leisure zones of their chaotic cities - what Foucault calls “sites of temporary relaxation” (Foucault, 1998).
- *Local protests.* To galvanize their neighborhood in opposition to a nearby logging project, mobile users could attach notes (e.g., texts, audio clips) to local buildings, to be read by others who come along later. Mobile phones have been already used to summoning people to demonstrations. In China, the biggest middle-class protests of recent years (against the use of abducted boys to perform dangerous work) has been organized by exchanging text messages. Empowering more people to become involved in their communities can improve public sector governance and enrich democracy.
- *Neighbors’ likes and dislikes.* Using their Bluetooth-enabled phones, people can share information about their personal interests with others (friends or strangers) in their proximity. Sharing metadata (not content) is old hat - it is what people do in Web 2.0 applications: they mostly share information about themselves and their personal interests.

2.2 Unlocking the Tapestry

All of the above location-based services are already offered on the Internet. Websites collect content generated by registered users and add “geotags” to that content (i.e., encode spatial co-ordinates).

Ironically, *location-based* content that is collected in such a *distributed* way finds itself “enclosed” on the Internet - a centralized and location-independent infrastructure. One may well ask why. Here is a possible explanation: by channeling user-generated content into their web sites, companies attempt to make money. Take Google: it “is often compared to Microsoft; but its evolution is actually closer to that of the banking industry” (TheEconomist, 2007). According to this widely shared view, Google is similar to a bank that capitalizes not on our money but on our personal data. Consequently, giving up data for Google would be tantamount to giving up profits - money coming from advertisers who exploit personal information to promote their wares in a targeted way.

However, most Web 2.0 companies are struggling to find viable business models, and they are not making any profit because they are pursuing Starbucks’ business model. Starbucks offers comfy chairs and does not charge people for sitting on them; people will buy overpriced coffee instead. “By offering a setting for free interaction, such sites provide the online equivalent of comfy chairs. The trouble is that, so far, there is no equivalent of the overpriced coffee that brings in the money and pays the bills” (TheEconomist, 2006). In theory, advertisements may generate profits. In practice, they have been found to annoy and drive people away.

Since Web 2.0 companies do not know how to make money, they are trying to get ideas from (the crowd of) external programmers. They let programmers access part of their user-generated data through APIs. Unfortunately, most of those companies may be doomed to failure because they:

- *Offer unscalable services.* The urban tapestry will be measured in petabytes of data, and Internet services will not scale simply because processing and exchanging data at this scale requires an infrastructure well beyond the means of the Internet.
- *Need to keep switching costs high.* As users are free to switch from one service to another, companies have little financial incentive to reduce switching costs. So data is often stored in proprietary file formats (protected by patents) and protected by service vendors. Giving access to their data via APIs is a first good step towards more open and innovative solutions. However, with company-defined APIs, the amount of accessible data is typically only a tiny part of the company’s knowledge base, so that the “wisdom of the (programming) crowds” is only partially exploited: unplanned innovation is serendipitous in nature and APIs are not open enough to accommodate it.

To sort out this current impasse, one may turn to managing location-based content using highly decentralised and open solutions which are more likely to:

- *Eliminate switching costs* - Users may be empowered to retain control of their data by simply storing it on their devices. To make that happen, MIT have recently put forward “A World Wide Web Without Walls” (W5) proposal: a project “that imagines a very different Web ecosystem, in which users retain control of their data and developers can justify their existence without hoarding that data”. In so doing, one eliminates switching costs - users do not need to share their data with each service provider. Plus, this approach comes with a pleasant by-product for privacy-conscious users: they would have control over what data they are willing to disclose.
- *Scale* - While existing companies fight over their “one size fits all” search engines, new companies may offer customized search solutions for communities in particular locations. That is made possible by two recent communication technologies: the first is Bluetooth, which connects only people who are in proximity; the second is WiFi, which connects mobile users to the Internet and enables the storage of location-relevant content on hotspots. These two technologies can assure dissemination and availability of location-dependent information. Assuring the availability of electronic data is a problem of scientific importance, and Ross Anderson has masterfully explored it in “The Eternity Service” (Anderson, 1996).

That is not to say that we stand at a crossroads. We do not need to decide whether to either lock the digital tapestry on the Internet or fully distribute it across portable devices. The future may well reside somewhere in the middle, and that “somewhere” will change depending on what technologies will be available. The introduction of new technologies largely depends on research. Since past research has focused on Internet solutions, it is time to study solutions that are distributed, and potentially mobile.

2.3 Problem Statement: Bringing Order to the Tapestry

Imagine that a decentralised infrastructure for storing user-generated, location-dependent content were available. Mobile users could then run software on their portable devices so that, when willing to consume content, such content would be retrieved from the tapestry and displayed on their devices. What challenges would such a software face? The two

problems to which the rest of the paper is devoted are the following:

1. How to select “relevant content” (Section 3). By relevant, we mean content that semantically matches a user query. For example, given the query “Japanese restaurant tempura”, relevant content could be user reviews of Japanese restaurants that serve dishes of deep-fried seafood and vegetables in tempura batter.
2. How to select “quality content” (Section 4). By quality, we mean content that has been produced by reputable sources. After receiving user reviews of Japanese restaurants, a device can rank them by reviewer’s reputation.

3 SELECTING RELEVANT CONTENT

The first problem is to select relevant content. Social (or folksonomic) tagging has become a very popular way to describe, categorise, search, discover and navigate content. This is done either by people, who associate keywords to some content, or even automatically by means of some tagging mechanism (e.g., by GPS-enabled cameras that tag pictures depending on location of capture (Rattenbury et al., 2007)). Unlike taxonomy, which overimposes a hierarchical categorisation of content, folksonomy are informally defined, continually changing, and ungoverned. In order to retrieve relevant content in this domain, the emergent semantics of tags must thus be learned and used to quantify the similarity between a query and (the tags associated to) an item.

Studies have been conducted both to understand tag usage and evolution (e.g., (Sen et al., 2006; Halpin et al., 2007; Heymann et al., 2007)), and to learn and exploit their hidden semantics. For example, in (Wu et al., 2006) a probabilistic generative model is proposed to describe users’ annotation behavior, and to automatically derive tags emergent semantics; during searches, the approach is capable of grouping together synonymous tags, while it calls for user’s intervention when highly ambiguous tags are found. Research has been very active also in relating tag activity to users, in order to discover their interests and consequently users’ communities, either by exploiting users’ explicitly stated profile (Hsu et al., 2007), or by using a probabilistic model which takes into account users’ interest to topics (Zhou et al., 2006), or based on their level of tagging activity and breadth of interests (Kelkar et al., 2007). All these works are based on the observation that real world networks exhibit a so-

called community structure (Ruan and Zhang, 2008); defining the set of characteristics that would enable the best fitting and natural clustering of taggers is an open research question.

Our Proposal: Social Filtering. In order to automatically filter content, we argue that the two research streams highlighted above (i.e., automatic learning of tag semantics and users’ interests) have to be combined (Zanardi et al., 2008). More precisely, for each query-item pair, we first compute the “relevance” of the item with respect to the query, based on the semantic distance between query tags and item tags; we then compute the similarity between “who has issued the query” and “who has tagged the item” based on their past tag activity, and use this weight as a multiplying factor to rank relevant content. Preliminary results on the CiteULike dataset demonstrate that users’ similarity improves accuracy of the results, while tags’ similarity improves coverage.

Future. All algorithms developed to date to learn tags semantics and filter content have been evaluated on Internet-based datasets, where a huge collection of data is available, and thus amenable to intensive processing. One of the main challenges we will thus have to face is porting these algorithms to the distributed setting, without compromising on accuracy, coverage and performance. Various techniques for data clustering will be called for, in order to aggregate related information together, for example around hotspots. Moreover, tag systems are highly susceptible to tag spam, that is, malicious annotations generated to confuse users (Koutrika et al., 2007). Robust solutions to tag spamming require further investigation, both in the centralised and decentralised setting.

4 SELECTING QUALITY CONTENT

The second problem is to select quality content. Mobile users may do so by simply selecting content coming from reputable sources. Sources are reputable if people have found them to be so in the past. In practice, this translates into people rating the content they consume. Upon those ratings, one identifies reputable producers - those who have consistently created highly-rated content.

To decide whether a certain producer is reputable, a filtering software needs to implement three functions:

- Rate the producer (Section 4.1).

- Personalize that rating based on its user's interests (Section 4.2).
- Update ratings whenever its user consumes content (Section 4.3).

4.1 Rating Producers

Consider that mobile phone *A* needs to rate a certain producer. It may do so by collecting ratings and arranging them in a graph - dubbed "web of trust". This is a network of trust relationships: we trust (link to) only a handful of other people; these people, in turn, trust (link to) a limited number of other individuals; overall, these trust relationships form a network (a web of trust) of individuals linked by trust relationships. Based upon this web of trust, *A* can then form opinions of producers (in technical parlance, it *propagates trust* in producers) from whom it has never received content before.

Existing ways of propagating trust cannot be readily applied in mobile computing because they are usually designed to work on a centrally stored web of trust and to run on high-end machines. Most of the work on how *A* propagates its trust for *B* is based on a simple, yet effective mechanism: *A* finds all paths leading to *B*; for each path, *A* then concatenates the ratings along the path; *A* finally aggregates all path concatenations into a single trust rating for *B*. Algorithmically, this is equivalent to *A* arranging trust ratings into a matrix and, over a series of iterations, propagating trust by, for example, *direct* propagation: if *A* trusts *C* and *C* trusts *B*, then trust propagates from *A* to *B*. The resulting matrix values are then rounded into a single trust rating. Unfortunately, this way of propagating trust suffers from two main limitations:

- Literature has proved direct trust propagation to be extremely effective, but it has done so only on datasets of *binary* ratings. However, an individual may express whether she trusts another individual or not, and, if she does, she may then express the extent to which she trusts by a *discrete* value. There is no published work on how direct propagation would perform on a *large* dataset of *discrete* ratings, not necessarily binary.
- Direct trust propagation does not scale on mobile devices. Direct trust propagation is meant for Web applications in which centralized servers store full webs of trust upon which trust is then propagated by multiplying vectors and matrices whose dimensions are extremely high. As a consequence, it is computationally expensive and would not scale well on any existing portable device. Moreover, mobile devices would only know a very

small subset of the web of trust at any given time (it is unrealistic to assume complete knowledge) because of, for example, network partition, device (un)availability, and limited resources.

We need a way of propagating trust that works in *distributed* settings and runs on (*resource-constrained*) mobile phones.

Our Proposal: Distributed Trust Propagation. We have recently designed one such way (Quercia et al., 2007a) by carefully adapting a graph-based semi-supervised learning scheme (Herbster et al., 2005; Zhu et al., 2003). The key idea is that each mobile device stores a very limited subset of the web of trust; on that subset, it then applies a machine learning technique for propagating trust.

The model scales (it entails minimal storage and communication overhead) and is effective (its predictive accuracy on the Advogato dataset is as high as 82.9%). That accuracy remains unchanged even if most of the users were concerned about privacy and, as such, were not to make available their ratings. The model also runs on portable devices (a J2ME implementation spends at most 2.8ms for one propagation on a modern Nokia phone).

Future. Our distributed trust propagation model assumes that users' ratings are stored in distributed way. However, the lack of a centralised server storing ratings result in such ratings being susceptible to malicious manipulation. To this end, we are currently working on a mechanism with which mobile phones store ratings in (local) *tamper-evident* tables and check the integrity of those tables through a gossiping protocol.

4.2 Personalizing Ratings

Trust propagation techniques generate single ratings. However, *A* may well have more than one rating for each content producer. To see why, say that *A* received "financial" news from *B*, found them interesting, and, as such, highly rated *B*. *A* is now interested in "economic" news, and *B* happens to produce them. From its past rating on "financial" news, can *A* conclude that *B*'s "economic news" are also of good quality? *A* may well conclude so since "economics" and "finance" are (semantically) similar.

To automatically decide whether two categories are similar, existing algorithms typically use an ontology (e.g., a taxonomy of content categories). Let us take two common approaches. The first (Capra, 2005; Liu and Issarny, 2004) defines similarity between any two categories in an ontology as the dis-

tance between the two corresponding nodes. The second approach (Kinateder and Rothermel, 2003) draws category similarity based on a direct graph of categories (a less-constrained structure than a tree) whose weights have to be, however, manually set by device users. The researchers who proposed the first approach have acknowledged that the idea of a universally accepted ontology hardly belongs to reality; those of the second approach concede that, on poor usability grounds alone, their solution has to be automated. More generally, existing approaches require that the same ontology is shared by all users and that those users agree on that ontology for good (i.e., the ontology is not supposed to change over time).

Our Proposal: TRULLO. To do away with these two problems, we have recently proposed an algorithm dubbed TRULLO (Quercia et al., 2007b) that *automatically* personalize ratings across categories *without relying on an ontology* shared by all users. This algorithm gathers ratings of past experiences in a matrix, learns statistical “features” from that matrix by applying the “Singular Value Decomposition”, and combines those features to set initial trust values for new content categories. By features, we simply mean textual information that describes categories. In contrast to existing approaches, TRULLO relies only on *local* information (the ratings of its user’s past experiences) and, as such, does not need to collect recommendations, thus avoiding the need for a common ontology shared by all (recommending) users.

TRULLO works well in a simulated antique market (whose simulation parameters are partly based on eBay). It performs close to how exchanging recommendations would do in an ideal (though unrealistic) world, one in which recommenders are wholly truthful and, furthermore, share the same ontology. Also, its J2ME implementation is reasonably fast on a modern Nokia mobile phone.

Future. To personalize ratings, TRULLO processes only the ratings of its user. However, to discover general relationships among categories, one needs a larger fraction of user ratings. That would be possible if mobile phones upload their ratings on WiFi hotspots, which then run more computational-demanding techniques for discovering category relationships.

4.3 Updating Ratings

Using existing mobile reputation systems, *A* rates *B* on a binary scale (good or bad) and consequently updates its trust for *B* with hand-crafted formulae.

To do away with *hand-crafted formulae*, Mui *et al.* (Mui et al., 2001) proposed a Bayesian formalization for a distributed rating process. However, two issues remained unsolved: they considered only binary ratings and did not discount them over time. Buchegger and Le Boudec (Buchegger and Boudec, 2004) tackled the latter issue, but not the former: they proposed a Bayesian reputation mechanism in which each node isolates malicious nodes, ages its reputation data (i.e., weights past reputation less), but can only evaluate encounters with a *binary* value (i.e., encounters are either good or bad). So literature lacks a *formal* way of updating ratings on a *generic* scale (not necessarily binary).

Our Proposal: B-trust. We designed a new trust model (Quercia et al., 2006) that updates *n*-level ratings (generally, $n > 2$) according to a Bayesian process. After rating *B*’s content, *A* updates its trust for *B* using Bayes’ theorem. As an example of application of this theorem, assume that *A*’s rating is “good”. Given that, *A* updates the probability p_t that *B* is trustworthy by taking the old p_t and multiplying it by $l_{g|t}$ - the likelihood that good content comes from trustworthy devices. If we leave out a proportionality constant at the denominator, the updating looks like:

$$p_t \propto p_t \cdot l_{g|t}$$

Common sense would suggest that good content usually comes from trustworthy devices (i.e., $l_{g|t}$ is high), and that bad content does not usually come from trustworthy devices (i.e., $l_{b|t}$ is low).

However, *A* does not set those likelihoods according to common sense. Instead, it learns them while receiving content, that is, by counting the number of times what type of content comes from what type of device (e.g., counting the number of times good content comes from trustworthy producers).

In designing B-trust, we have extended this formulation to the case in which *A* rates on a generic *n*-scale (not necessarily binary – good/bad).

Future. Producers may excessively capitalize on their old ratings. So B-trust decreases confidence in its ratings over time. However, by doing so, B-trust may fail to identify communities of trustworthy producers that are *stable*. So researchers have started to study how ratings evolve over time, and how that affects the ability to identify stable communities (Lathia et al., 2008).

5 EVALUATING MOBILE SOLUTIONS

Our research agenda has been evolving around the theme of finding relevant content that will satisfy a user's query. To this extent, we have been proposing various algorithms to: select relevant content, based on dynamically inferred tags semantics; rank filtered content based on quality, by dynamically assessing content sources' reputation. Will these algorithms become enabling technologies for pervasive content sharing applications? In order to answer this question, we (and the research community working on these topics) is faced with a big challenge: how to evaluate these algorithms.

Data about content and content sharing abound on the Internet; however, conducting studies on such data inevitably fails to measure what would happen in a truly distributed setting. On the other hand, there exist plenty of experimental observations of how people move while carrying their portable devices; in this case, though, there is little or no information about what content people produce and consume.

As a short-term solution, researchers can "mimic" what would happen in a real pervasive system, by overlaying these different datasets; however, doing so in a meaningful way is a research question of its own. Simulation should be coupled with controlled experiments; the problem in so doing is that those studies are expensive, so one tends to trade off between (user) sample size, time requirements, and monetary costs; the generality of the results obtained thus becomes questionable. To help solve this problem, PARC researchers have recently proposed to collect user measurements from micro-task markets (such as Amazon's Mechanical Turk) (Kittur et al., 2008). In the long run, an actual large-scale system deployment will be needed.

6 CONCLUSIONS

In this paper, we have discussed distributed mechanisms with which mobile users can find content of interest and of high quality. Compared to existing (centralized) mechanisms, distributed mechanisms promise to scale and be fully open to innovation. However, to deliver on this this promise, we still need to study how effective those mechanisms are in practice. The lack of real datasets, combining mobility with user's interests and content, makes evaluating these mechanisms an open challenge.

REFERENCES

- Anderson, R. (1996). The Eternity Service. In *Proc. of Pragocrypt*.
- Bassoli, A., Brewer, J., Martin, K., Dourish, P., and Mainwaring, S. (2007). Underground Aesthetics: Rethinking Urban Computing. *IEEE Pervasive Computing*, 6(3):39–45.
- Buchegger, S. and Boudec, J.-Y. L. (2004). A robust reputation system for P2P and mobile ad-hoc networks. In *Proc. of the 2nd Workshop on the Economics of Peer-to-Peer Systems*.
- Capra, L. (2005). Reasoning about Trust Groups to Coordinate Mobile Ad-Hoc Systems. In *Proc. of the 1st IEEE Workshop on the Value of Security Through Collaboration*, Athens, Greece.
- Foucault, M. (1998). Of other space. *The visual culture reader*.
- Halpin, H., Robu, V., and Shepherd, H. (2007). The complex dynamics of collaborative tagging. In *Proc. of the 16th Intl. Conference on World Wide Web*, pages 211–220, NY, USA.
- Harada, T., Gyota, T., Kuniyoshi, Y., and Sato, T. (2007). Development of Wireless Networked Tiny Orientation Device for Wearable Motion Capture and Measurement of Walking Around, Walking Up and Down, and Jumping Tasks. In *Proceedings of the IEEE Conference of Intelligent Robots and Systems*, San Diego, US.
- Herbster, M., Pontil, M., and Wainer, L. (2005). Online learning over graphs. In *Proc. of the 22nd Int. Conference on Machine Learning*.
- Heymann, P., Koutrika, G., and Garcia-Molina, H. (2007). Can Social Bookmarking Improve Web Search? *Resource Shelf*.
- Hsu, W. H., Lancaster, J., Paradesi, M. S., and Weninger, T. (2007). Structural Link Analysis from User Profiles and Friends Networks: A Feature Construction Approach.
- Kelkar, S., John, A., and Seligmann, D. (2007). An Activity-based Perspective of Collaborative Tagging. *Intl. Conference on Weblogs and Social Media*.
- Kinatered, M. and Rothermel, K. (2003). Architecture and Algorithms for a Distributed Reputation System. In *Proc. of the 1st Intl. Conference on Trust Management*, pages 48–62, Crete.
- Kittur, A., Chi, E., and Suh, B. (2008). Crowdsourcing User Studies With Mechanical Turk. In *Proceedings of the ACM Conference on Human-factors in Computing Systems*, Florence, Italy.
- Koutrika, G., Effendi, F. A., Gyöngyi, Z., Heymann, P., and Garcia-Molina, H. (2007). Combating spam in tagging systems. In *Proc. of the 3rd Intl. Workshop on Adversarial Information Retrieval on the Web*, pages 57–64, NY, USA.
- L. McNamara, C. Mascolo and L. Capra (2007). Content Source Selection in Bluetooth Networks. In *Proc. of*

- the 4th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, Philadelphia, USA.
- Lathia, N., Hailes, S., and Capra, L. (2008). Evolving communities of recommenders: A temporal evaluation. In *Research Note RN/08/01, Dept. of Computer Science, University College London*.
- Liu, J. and Issarny, V. (2004). Enhanced Reputation Mechanism for Mobile Ad Hoc Networks. In *Proc. of the 2nd Intl. Conference on Trust Management*, volume 2995, pages 48–62, Oxford.
- Mui, L., Mohtsahemi, M., Ang, C., Szolovits, P., and Halberstadt, A. (2001). Ratings in Distributed Systems: A Bayesian Approach. In *Proc. of the 11th Workshop on Information Technologies and Systems*, New Orleans, USA.
- Paulos, E. and Goodman, E. (2004). The familiar stranger: anxiety, comfort, and play in public places. In *Proc. of ACM Conference on Human Factors in Computing Systems*, pages 223–230.
- Quercia, D., Hailes, S., and Capra, L. (2006). B-trust: Bayesian Trust Framework for Pervasive Computing. In *Proc. of the 4th International Conference on Trust Management*, pages 298–312, Pisa, Italy. LNCS.
- Quercia, D., Hailes, S., and Capra, L. (2007a). Lightweight Distributed Trust Propagation. In *Proc. of the 7th IEEE International Conference on Data Mining*, Omaha, US.
- Quercia, D., Hailes, S., and Capra, L. (2007b). TRULLO - local trust bootstrapping for ubiquitous devices. In *Proc. of the 4th IEEE Intl. Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*.
- Rattenbury, T., Good, N., and Naaman, M. (2007). Towards automatic extraction of event and place semantics from flickr tags. In *Proc. of the 30th ACM Conference on Research and Development in Information Retrieval*, pages 103–110, Amsterdam, The Netherlands.
- Ruan, J. and Zhang, W. (2008). Identifying network communities with a high resolution. *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)*, 77(1).
- Sen, S., Lam, S. K., Rashid, A. M., Cosley, D., Frankowski, D., Osterhouse, J., Harper, M. F., and Riedl, J. (2006). Tagging, Communities, Vocabulary, Evolution. In *Proc. of the 20th Conference on Computer Supported Cooperative Work*, pages 181–190, NY, USA.
- Sharon, M. (2006). Mobile Mappa Mundi: using cell phones as associative mapping tools. In *Sociallight White Paper*.
- Steen, M. (2007). TomTom and Vodafone crowdsource traffic information. *Financial Times*, November 12th.
- TheEconomist (2006). The trouble with YouTube. August 31st.
- TheEconomist (2007). Who's afraid of Google? August 30th.
- Wu, X., Zhang, L., and Yu, Y. (2006). Exploring social annotations for the semantic web. In *Proceedings of the 15th ACM Conference on World Wide Web*, Edinburgh, UK.
- Zanardi, V., , and Capra, L. (2008). Social Ranking: Finding Relevant Content in Web 2.0. In *Proceedings of International Workshop on Recommender Systems*, Patras, Greece.
- Zhou, D., Manavoglu, E., Li, J., Giles, L. C., and Zha, H. (2006). Probabilistic models for discovering e-communities. In *Proceedings of the 15th International Conference on World Wide Web*, pages 173–182, New York, NY, USA. ACM Press.
- Zhu, X., Ghahramani, Z., and Lafferty, J. (2003). Semi-supervised learning using Gaussian fields and harmonic functions. In *Proc. of the 20th International Conference on Machine Learning*, Washington, USA.