

Collaborative Activity Recognition

George Kampis^{1,2} and Paul Lukowicz¹

¹ DFKI, German Research Institute for Artificial Intelligence, Kaiserslautern,
Germany

george.kampis@dfki.de, paul.lukowicz@dfki.de,
WWW home page: <http://ei.dfk.de>

² ITMO University, St. Petersburg, Russia

Abstract. We study simulation models of spreading on peer-to-peer communication networks where any peer (or agent) can be the source of information, be it sensory recognition or contextual knowledge. In such a situation the value or quality of information is of key relevance. Questions of trust, provenance and the problem of the interaction pattern arise and are approached by three different algorithms in our paper: (i) “quantitative democracy”, where knowledge is averaged on a meeting (ii) “experience takes all”, where the more experienced (the teacher) overwrites all prior knowledge of the less experienced (the “student”), and (iii) “transitive experience” where not only information but also experience is handed over. We compare these different regimes and identify their tradeoffs.

Keywords: trust, provenance, self-organization, emergence, collaborative information processing

1 Introduction

Peer-to-peer communication in smartphone networks is an important issue for crowd control and for smart society applications. Smartphone sensor data voluntarily contributed by users [6, 3] can be used for collaborative localization as well as monitoring and analysis (of crowd density for instance). A system was developed, deployed and tested by our group in various real life situations such as the Zurich festival of 2013 [3]. In recent work [1, 4] we followed a basic opportunistic approach to connect devices in a peer-to-peer fashion, based on the built-in WiFi hotspot functionality. The devices’ switching between access point and client modes allows for the propagation of messages on a multi-hop basis. We presented a large-scale simulation [5] based on a dataset consisting of movement traces from 28.000 people and fine-tuned parameters of the switching algorithm. We simulated and analyzed the conditions under which signal propagation from a dedicated source is possible to achieve in an efficient way in such systems.

Our proposal is motivated by problems arising from generalizations of the scheme. We are interested in smart society applications¹ where any peer (agent) can be the source of new information and a natural need arises as to forward

¹ <http://www.smart-society-project.eu>

and evaluate these information segments (to know, for instance, the activity of a given agent). Two seemingly contradicting criteria arise at this point: any locally generated useful information should be able to reach the entire community; harmful or misleading information should be filtered out. We may think of restaurant quality assessment at one extreme and messages about public hazards at the other. We should enable the system to self-organize in such a way as to facilitate but also to check spreading, done by the network itself, dynamically and in a distributed fashion. In other words, we deal with a “complex adaptive system” or CAS, where different types of entities collectively give rise to “emergent” functionalities (Fig. 1).




Fig. 1. A “complex adaptive system” (CAS). Which agents interact, who to trust, where has information its origin?

We focus on human-agent, or even more specifically human-agent-human relations where persons exchange messages about internal states via agent-to-agent interactions. Reliability, reputation, trustworthiness and other quality parameters of the information sources play a role. To understand them, questions related to trust, provenance, privacy as well as the role of different interaction patterns for the exchange of information need to be studied. Of these, the followings will be considered in this paper:

Trust. Which agents are trustworthy, or trustworthier than others - how much weight is allocated to each communication?

Provenance. What are sources of information, how well informed are the given agents?

Interaction pattern. Who can “talk” to whom, which agents’ information can be combined or brought in relation?

A preliminary analysis of these factors by way of simulation modeling is the target of the present paper. More specifically, we present and analyse 3 different algorithms that imply disparate answers to the above questions. Our algorithms can be directly deployed to active devices and the models can directly translate to a real environment (combination with GIS shapefiles has been realized but not discussed in this paper).

2 Algorithms for Collaborative Activity Recognition

The three algorithms. The algorithms considered here are informed by earlier studies in opinion dynamics [2] or p2p “gossip” systems [7], the connections to which will not be discussed here. The three algorithms to be presented are, in a nutshell: (i) “quantitative democracy” (QD), where knowledge is averaged on a meeting (ii) “experience takes all” (EXP), where the more experienced agent (such as a teacher) overwrites the prior knowledge of the less experienced (the “student”), and (iii) “transitive experience” (TRAN) where not only knowledge but also experience is handed over to the meeting partner (such that a “student” of a “guru” becomes a “guru” next).

Meetings in physical space, meeting pattern. We consider agents moving in physical space (such as on the streets of a city). Collaborative events occur upon the personal meeting of the agents. Agents can be persons, apps, machines, or persons with apps, etc. – a “meeting” is understood as a contact between two different agents so information exchange is possible within a radius of r , a model parameter. (Therefore, information aggregation will be space local, provoked by physical motion. We leave more general “network” studies to subsequent work - to define meetings more directly in a virtual space where connections could be established either in a uniform random fashion or by preferential attachment, etc.) Meeting patterns arise from Boolean motion in an unstructured field. From our earlier studies [5] we know that the choice of meeting regime (e.g. a street map) does not qualitatively change things.

Information as binary classification. The concept with which we approach information is that of a *binary classifier*. A binary classifier (BC) is characterised with a decision probability p that yields decision $D(p)$ (0 or 1). This is a general formalism so we can assume every agent stores knowledge and contextual (sensory) information in the form of a binary classifier. Then our task is to examine how classifications change upon interactions locally, and yield meaningful patterns globally; from this simple starting point we can later move towards more general classifications and other forms (e.g. procedural) of knowledge.

Decision error. The advantage of the classifier representation is that recognition information is stored in the form of a probabilistic *decision*. For study purposes it will be important that the experimenter knows the “right” decision which is kept constant during a run. That way we can conveniently speak of information quality and define *decision error* in the population as the ratio of right to wrong decisions of the agents.

Well Informed Agents (WIAs). We also introduce well informed agents (WIAs), as entities that always have the right decision information with $p = 1$, and never update p upon meetings. The rationale is that WIAs have the right information not by chance but from a reliable source (see the notion of provenance above) so “they know that they know”. WIAs can facilitate the spreading of true information e.g. inform other agents of external events (to approximate epidemiological spreading where a message “infects” a population) and hinder the spreading of false information e.g. to prevent that false alarms cause panic. Thus we hypothesize that WIAs will play a key role in the collaborative recog-

dition dynamic. In fact without WIAs the 3 algorithms under investigation can only be of a very limited interest for the current study: as can be readily seen, QD ends up in a dull “average” global state, EXP (and therefore TRAN wich is a variant thereof) end up in random drift around emergent centers having arbitrary information quality.

Simulation models. We test and analyze computational simulation models for the above algorithms using agent based modeling [8]. We fix many important parameters (as listed in Table 1 in [4]) and now concentrate on (i) the quantitative effect of WIAs on information spreading in the different cases and (ii) the qualitative form of p distributions (characteristically different for the chosen algorithms and reflecting their different ways of handling information).

3 Quantitative Democracy (QD)

This is the first rule and the name tells the story. Under this rule, when two agents meet, they compute and share the average of their p values: $p_{new} = (p_1 + p_2)/2$. Each agent has an equal trust and importance, except for WIAs where the symmetry is broken: whereas the QD rule acts on the meeting partner A by updating its $p(A)$, the same is not true for the WIA where (by definition) no updating takes place. Here $p_{new}(A) = (p_{old}(A) + 1)/2$ which shows that upon repeated WIA encounters $p(A)$ should monotonically increase towards 1.




Fig. 2. QD rule. Left: the effect of WIA proportions (horizontal) on convergence speed (vertical). Right: the histogram of p values is a unimodal function (cf. Figs. 3 & 4).

Thus in QD the role of WIAs is seen in the cleanest form. Without WIAs there will be no interesting dynamics - as discussed above -, but from the above argument one single WIA is enough to make the system converge to the decision with $p = 1$. The time necessary for such a convergence (called zero crossing time

by a physics analogy) is variable however and expected to depend on factors including meeting density, the geometry of motion and others, among them the number of WIAs. Studying the effect of the latter we get the results in Figure 2.

It is seen that, not surprisingly, an increasing proportion of WIAs decreases the time necessary to reach convergence; the value changes rapidly with more WIAs first but changes only a little above 6-7%. The frequency distribution (histogram) of p values is a potentially Gaussian or lognormal single-mode distribution that dynamically changes with the p values. WIAs are on the right side and the distribution shifts towards them as time progresses.

4 Experience Takes All (EXP)

In QD even WIAs are not recognised as trustable information sources and their effect, however dramatic, is due to the fact that WIAs don't update and hence there is no equilibrium point at any other p than theirs; whenever an agent meets a WIA, its p changes towards the WIA's. But this can be a slow process and is based on the effect of repeated meetings with the WIAs (or with other agents already informed by the WIAs on such repeated meetings).

The EXP rule changes the situation by introducing a simple reputation system based on experience (therefore the name). Under EXP the less experienced agent takes information from the more experienced upon a meeting.




Fig. 3. EXP rule. Left: average behavior without convergence. Right: the p histogram shows bars of values where fluctuations give rise to spreading centers besides WIAs (which are shown on the right).

Experience is defined here simply as the number of prior meetings, and the idea is that experienced agents integrate information and are more reliable sources. We can thus think of this as a teacher-student situation where the

teacher is trusted for his experience. It is natural then to conceptualize WIAs as having a *bias*: an initial nonzero experience that defines them as teachers from the beginning on. Thus WIAs have a head start; can they retain the advantage?

Analysis results (Figure 3) show this not to be the case. It is easy to see that an arbitrary high bias would put WIAs in the same position as in QD, just better: for an agent A in this case $p_{new}(A) = 1$ which, other things being equal, ensures a similar yet even faster convergence than seen in the QD case. At the other extreme, with no WIAs (or when the *WIAs have no bias*) the local agents develop experience subject to various random fluctuations and some of them can get high reputation scores for no reason with respect to information quality (or decision error). So the interesting dynamics must occur in the range where bias is present but it is bounded. On Fig. 3 (left) we see that for a fixed value of bias the average behavior shows no convergence within the studied time window (hence maximal time values occur) yet with increasing WIA proportions an ever higher number of outliers move the error bars downwards.

Equally interesting is the p histogram, Fig. 3 (right). It shows many emergent centers represented with a high number of copies with arbitrary p values. The agents that become teachers know no better than any other; but fluctuations can make them experienced even beyond that of WIAs at any bias level. With a decreasing but nonzero probability such random agents can thus outperform WIAs (or perform just as well). The histogram shows such “false prophets” that spread unreliable information locally. The two figures combined together, it is the endurance of ‘false prophets’ that prevents convergence to WIAs’ p in a reasonable time. (Convergence occurs at the time limit, yet not necessarily to the WIA values.)

5 Transitive Experience (TRAN)

Despite the intuitive advantage, EXP is thus better at maintaining variety than spreading quality information collectively. A possible remedy can be suggested rethinking the role of experience as a key determinant of interactions in EXP. The new rule TRAN is like EXP but it makes experience transitive: on meeting a teacher, not only does the p value change but also teacher experience is taken over. Then we expect information to spread more reliably: those possessing quality information will themselves become instant spreaders (meetings with a WIA thus turns every agent into a new WIA except that theoretically it could still update.) The expectation is that the head start of the WIAs thereby leads to a fast initial spreading that gives enough advantage to the trustworthy to prevail.

Analysis shows this indeed to be the case – but there is a twist. In most cases there is a fast convergence as expected (Fig. 4) but the error term is large and there are systematic outliers even beyond that. That is, the typical process (and thus the average behavior) is fast convergence but the large error terms and many remote outliers color the picture: it is thus with a high residual probability that a given concrete process won’t finish early. The reason is similar as in EXP: local centres can survive for a long time (yet not outperforming the WIAs here),

and only ultimately giving way. In other words, in this rule we see the power of contingencies - physically remote places can nurse false information for a long time, depending on whether and when meetings with WIAs or their kin occur.




Fig. 4. TRAN rule. Left: on average, rapid convergence to true information but with a high number of remote outliers. Right: the p histogram as in EXP, except that the WIA bar grows and finally wins, albeit sometimes slowly.

6 Comparison and Summary




Fig. 5. Comparison of QD (dark blue), EXP (orange) and TRAN (turquoise). A composite of Figs. 2-4. with WIA proportions left to right, and zero crossing time vertical.

The three algorithms examined handle trust, provenance and interactions differently. QD equalizes roles (and thus information), making WIAs move forward slowly but steadily. EXP on the other hand amplifies prior events and permits the experienced, called “teachers” spread information rapidly; however any teacher advantage is bound to disappear at one point due to fluctuations throwing up random agents with a randomly accumulated high experience. Finally, TRAN makes sure that teachers preserve their advantage and this leads to a fast process typically, but with a frequent slowdown caused by similar contingencies that can entirely ruin EXP. The summary in Fig. 5 compares the three to scale.

In this paper we presented and analyzed basic algorithms for local information fusion for collaborative activity recognition. The aim was to test principles; exploitation and general discussion is left to subsequent papers.

7 Acknowledgements

This work is supported by the European Community’s Seventh Framework Program (FP7/2007-2013) under grant agreement #600854 “Smart Society” as well as the European Commission’s H2020 Program under the funding scheme “FET Proactive: Global Systems Sciences” (GSS) as a Research and Innovation Action, Grant agreement n. 641191 (CIMPLEX).

References

1. T. Franke, G. Kampis, and P. Lukowicz. Leveraging human mobility in smartphone based ad-hoc information distribution in crowd management scenarios. Submitted to MobiSys 2015. <http://www.sigmobile.org/mobisys/2015/>.
2. R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3), 2002.
3. G. Kampis, J.W. Kantelhardt, K. Kloch, and P. Lukowicz. Analytical and simulation models for collaborative localization. *J. Computational Science*, 6(1):110, 2015.
4. G. Kampis and P. Lukowicz. Collaborative knowledge fusion by ad-hoc information distribution in crowds. Submitted to ICCS 2015, International Conference on Computational Science, Reyjavik.
5. G. Kampis and P. Lukowicz. Collaborative localization as a paradigm for incremental knowledge fusion. 5th IEEE CogInfoCom 2014 Conference, 2014. http://coginfocom.hu/conference/CogInfoCom14/downloads/Program_CogInfoCom_2014_final.pdf.
6. K. Kloch, P. Lukowicz, and C. Fischer. Collaborative PDR localisation with mobile phones. In *Proceedings of the 2011 15th Annual International Symposium on Wearable Computers, ISWC 11*, pages 37–40. IEEE Computer Society, Washington, DC, USA, 2011.
7. R. Ormandi, I. Hegedus, , and M. Jelasity. Asynchronous peer-to-peer data mining with stochastic gradient descent. In *In Emmanuel Jeannot, Raymond Namyst, and Jean Roman, editors, Euro-Par 2011, volume 6852 of Lecture Notes in Computer Science*,, pages 363–368, New York, NY, USA, 2011. Springer-Verlag.
8. S. Tisue and U. Wilensky. Netlogo: A simple environment for modeling complexity. In *International conference on complex systems*, pages 16–21, 2004.