

On Gossip-based Information Dissemination in Pervasive Recommender Systems

Tobias Eichinger
tobias.eichinger@tu-berlin.de
Service-centric Networking
Technische Universität Berlin
Berlin, Germany

Felix Beierle
beierle@tu-berlin.de
Service-centric Networking
Technische Universität Berlin
Berlin, Germany

Robin Papke, Lucas Rebscher,
Hong Chinh Tran, Magdalena
Trzeciak
Technische Universität Berlin
Berlin, Germany

ABSTRACT

Pervasive computing systems employ distributed and embedded devices in order to raise, communicate, and process data in an anytime-anywhere fashion. Certainly, its most prominent device is the smartphone due to its wide proliferation, growing computation power, and wireless networking capabilities. In this context, we revisit the implementation of digitalized word-of-mouth that suggests exchanging item preferences between smartphones offline and directly in immediate proximity. Collaboratively and decentrally collecting data in this way has two benefits. First, it allows to attach for instance location-sensitive context information in order to enrich collected item preferences. Second, model building does not require network connectivity. Despite the benefits, the approach naturally raises data privacy and data scarcity issues. In order to address both, we propose Propagate and Filter, a method that translates the traditional approach of finding similar peers and exchanging item preferences among each other from the field of decentralized to that of pervasive recommender systems. Additionally, we present preliminary results on a prototype mobile application that implements the proposed device-to-device information exchange. Average ad-hoc connection delays of 25.9 seconds and reliable connection success rates within 6 meters underpin the approach's technical feasibility.

CCS CONCEPTS

• **Information systems** → **Collaborative search; Nearest-neighbor search; Information retrieval.**

KEYWORDS

Opportunistic Network; Mobile Network; Decentralized Recommender System; Pervasive Recommender System; Data Scarcity; Data Privacy; Information Dissemination

ACM Reference Format:

Tobias Eichinger, Felix Beierle, and Robin Papke, Lucas Rebscher, Hong Chinh Tran, Magdalena Trzeciak. 2019. On Gossip-based Information Dissemination in Pervasive Recommender Systems. In *Thirteenth ACM Conference on Recommender Systems (RecSys '19)*, September 16–20, 2019, Copenhagen, Denmark. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3298689.3347067>

1 INTRODUCTION

Nowadays, the possibility to collect, store, and process very large amounts of data in combination with powerful data transformation and analysis techniques have raised privacy concerns such as when in early 2018, Cambridge Analytica had been granted access to millions of Facebook user profiles for political campaigning without users being aware of it. Statutory counter measures to curb data misuse have in particular been undertaken by the European Union in May 2018 by adopting the General Data Protection Regulation (GDPR). The GDPR imposes strict rules on data processing, ownership, information obligation (including raising consent), transparency, and collection of user-related data. Since data protection laws penalize misuse of personal data reactively, they do not actively prohibit misuse on a technical level. In the literature, three major recommender system architectures have been proposed to address privacy issues arising from collecting large authoritative data pools.

Federated learning produces personalized recommendation models by communicating model building between a central server and mutually disconnected peers holding personal information [10, 22]. Thus, federated learning enacts model building centrally on data that is distributed across personal data owners.

In contrast, decentralized recommender systems feature direct interactions between distributed peers without a central server. They are commonly built on top of file-sharing peer-to-peer networks [3, 21, 28, 31] that use gossip mechanisms [19] in order to establish a logic overlay network for fast network search and network resilience in view of peers joining or churning the network. In short, dissimilar peers are dropped from and similar peers are added to views (lists of visible peers in the network) iteratively. In so doing, the network establishes homogenous interest groups, which share recommendations explicitly among each other.

Pervasive (or ubiquitous) recommender systems [23, 27] are systems that often revolve around location-aware recommendations of for instance items in a nearby shop, restaurants, or events in proximity. The location-based approach naturally circumvents transmitting personal data to a central remote authority for recommendation by instead requesting closeby profile as well as context information.

Permission to make digital or hard copies of all or part 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 bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

RecSys '19, September 16–20, 2019, Copenhagen, Denmark

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6243-6/19/09...\$15.00

<https://doi.org/10.1145/3298689.3347067>

Pervasive recommender systems are naturally confronted with limited profile data, for usually only a small subpopulation is available in proximity. Fortunately, there are indications that integrating context data into the recommendation process allows to outweigh profile data scarcity [1].

Among these three approaches we believe that pervasive recommender systems yield the highest potential for data privacy for two reasons. First, the availability of context data as well as mobile compute power on smartphones increases rapidly rendering more and more complex recommendation algorithms feasible on smartphones holding personal data. Second, the model building process happens on-device and does not require connectivity to the network. Consequently, peers become invisible in the network when they do not interact with other peers thus adding privacy on the level of model building. In spite of their potential for recommendation, we see two main burdens. First, pervasive recommender systems are susceptible to data privacy issues since the recommendation mechanism usually builds on the exchange of raw profile data with nearby peers. Second, local scarcity of profile data renders item recommendations on location-independent items taxing. We believe that introducing gossip-based mechanisms and data sampling strategies to the field of pervasive recommender systems can alleviate both issues. We present the following preliminary results:

- The design of *Propagate and Filter*, a gossip-based method that addresses the problem of data privacy and data scarcity in pervasive recommender systems.
- An implementation of the propagation part in the form of an Android mobile application utilizing Google's Nearby Connections API and its evaluation.

2 RELATED WORK

In the present work, we propose a method for disseminating recommendations epidemically in combination with an on-device filtering process for which we proposed a mobile software architecture in [5]. The most similar work is that of Barbosa et al. [25] that propose device-to-device raw profile exchanges in an opportunistic networking scenario. Yet, it differs in that they neither address scalability nor privacy issues. Other related work subsumes:

Information Dissemination: Technical works on device-to-device information exchange in proximity are key enablers for collaboratively built recommender systems on mobile devices. The *Haggle* API [24] implements data-centric message forwarding using Bluetooth, Ethernet, and WiFi¹. In [26], the authors implement a Bluetooth-based middleware to create opportunistic networks between passing users. On top of such enablers, the authors in [7] leverage real-world social phenomena such as large crowds at sports events for event-based mobile communication.

Information Filtering: Many approaches to collaboratively and decentrally build recommender systems on mobile devices in pervasive computing environments disseminate information epidemically and generate recommendations on top of user similarity. User similarities have been found to be reliably estimated by contextual information in the form of proximity at a music festival [11]

or spatio-temporal information in the tourism domain [16]. Estimating user similarity on rating vectors has been performed on subcommunities [13] or affinity networks [12, 29].

3 DESIGN

At the proposed information dissemination method's core is the social movement of people carrying mobile devices. On top of that core movement, we conceive a background data exchange between devices whenever they are geographically close to each other followed by an on-device customizable filtering process. We call this method *Propagate and Filter*.

3.1 Similarity Data, Peer Preference List, and Neighborhood Preference List

Each device carries four types of data (see Figure 1), where the entire stack of data² is used for the derivation of local personalized recommendations.

- *Peer Preference List:* A list of items rated by the peer³. It can contain binary or scalar ratings. In our prototype implementation, the peer can rate movies, where each movie is identified with a unique identifier provided by the publicly available Internet Movie Database⁴ (IMDb). Peer preference lists are kept on the device.
- *Neighborhood Preference List:* Every peer mixes⁵ previously collected neighborhood preference lists received from the k most similar peers into a single list of item ratings. It is thus an aggregated preference list of an unknown subset of peers. Neighborhood preference lists get propagated to other peers. Note that every peer controls the amount of his/her own peer preference list that gets propagated to nearby peers.
- *Similarity Data:* Any kind of data that can be used for peer similarity comparison. Similarity data gets propagated to nearby peers and therefore has to be privacy-preserving. Privacy-preserving similarity comparison can among others be performed on item vectors [4] as well as texting data [15].
- *Context Data:* Data that characterizes the encounter such as location, time, weather, or peer activity (running, eating, commuting) that can be *sensed* (for example via sensors) or *retrieved* (for example from the web) [6, 30].

3.2 Propagate

When two or more peers are geographically close to each other, their smartphones establish pairwise fast and secure connections and exchange their neighborhood preference lists and similarity data (propagation). At propagation time the received data is enriched with context data such as time or location characterizing the encounter.

²The integration of retrievable content information on item references received from nearby peers is omitted here since it is independent of pervasive recommender systems.

³Commonly referred to as *item vector*, yet omitted due to a conceptual overlap with neighborhood preference list.

⁴<https://www.imdb.com/>

⁵We envision a bootstrapping approach for mixing [8, 14].

¹<http://user.it.uu.se/~erikn/papers/haggle-arch.pdf>

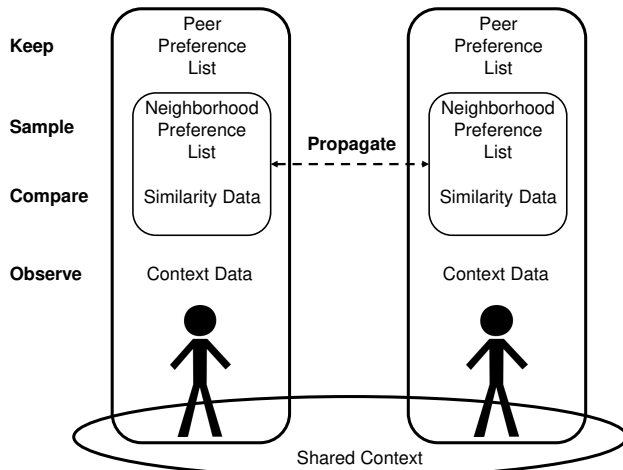


Figure 1: Four types of data proposed by Propagate and Filter and their respective use at propagation. Recommendations are calculated on the mobile device upon all four data pools.

3.3 Filter

Data collection in the propagation step includes unfiltered data from every encounter. It is necessary to filter by similarity in order to arrive at relevant information. Upon receiving data from another peer, the filtering process starts on the device. Three steps happen:

- (1) *Similarity Comparison*: Similarity data is used in order to compare peer similarity between sender and receiver.
- (2) *Sample Neighborhood Preference List*: If the peer similarity is above the k -th highest, resample the neighborhood preference list on the basis of the peer's preference list and the neighborhood preference lists by the k most similar peers.
- (3) *Update Personal Recommendations*: Run a recommendation algorithm⁶ on the locally available data (see Figure 1) in order to derive new recommendations or update ratings of previously generated recommendations.

3.4 Privacy by Disconnection

Propagate and Filter's propagation step establishes wireless connections between smartphones in proximity in an ad-hoc fashion, exchanges similarity data and neighborhood preference lists, and thereafter terminates the connection. Consequently, the network topology is essentially disconnected and no information on inter-peer relationships – such as is the case with Peer-to-Peer network overlays in social networks [20], recommender systems [3], or vehicular networks [17] – are exploitable. The entire peers' local databases form a geographically distributed and essentially disconnected database, where every peer holds only a very limited portion of data. Data query and data search are thus not possible at will. Access to other peers' data is limited to the time of contact and amount of data individually made available by nearby peers.

⁶Propagate and Filter is independent of any specific recommendation algorithm. The filtering techniques described in [2] are viable approaches leveraging contextual data.

We call this property *Privacy by Disconnection*. It is Propagate and Filter's contribution to data privacy.

3.5 Enhance Locally Available Profile Data

Gossip protocols require a connected peer-to-peer network in order to converge similar peers toward each other, where network connectivity is retained by peer sampling [18]. In traditional decentralized recommender systems, peer sampling that requires network connectivity is applicable since neither items nor peers are *spatial*, that is both can be moved in the network at will. In the Propagate and Filter scenario, peers *are* spatial and cannot be moved in the network at will, though items can. Therefore, Propagate and Filter proposes to converge the recommendation list of latent interest communities in the following way: Peer sampling does not have to be administered by any protocol since it is performed by the global movement of self-organized agents carrying smartphones. Propagate and Filter unfolds its dissemination potential at locations of high degrees of peer mobility such as urban areas [9]. When a peer receives data from a similar peer, he/she resamples his/her neighborhood preference list and else does nothing. As a consequence, Propagate and Filter creates a constant flow of item recommendations that flows between similar peers and dries out between dissimilar peers. This property allows to relay recommendations between peers that have never been geographically close to each other and avoids having to transmit any information on dissimilar peers. In that sense Propagate and Filter addresses the profile data scarcity problem prone to pervasive recommender systems.

4 IMPLEMENTATION

We present an Android mobile application that implements a prototype of the propagation step described in Section 3.2. We restrict to the transmission of peer preference lists for reasons of simplicity. The application is available in the Google Play Store⁷.

4.1 Prepare a Peer Preference List

The prototype allows to search and rate movies locally which are registered and uniquely identified in the IMDb. Movie ratings are scalar, ranging from 1 to 5 stars, and stored in the format (userID, movieID, scalarRating). A list of movie ratings implements the peer preference list introduced in Section 3.1. Once ratings have been specified by the user, he/she can activate sharing.

4.2 Propagate via Google Connections API

The application's active sharing mode applies *advertise* – the device broadcasts its existence to other devices in proximity – and *discover* – the device listens on other device's broadcasting. The sharing process is handled by the Google Nearby Connections API⁸ and entirely happens in the background. The Google Nearby Connections API connects devices using one of the three wireless technologies Bluetooth, Bluetooth Low Energy (BLE), or WiFi, automatically selecting the most efficient one in each scenario.

When two smartphones establish a connection, ratings are exchanged immediately and stored locally. Note that ratings can be exchanged without a connection to the internet as it is commonly

⁷<https://play.google.com/store/apps/details?id=de.tub.affinity.android>

⁸<https://developers.google.com/nearby/connections/overview>

Table 1: Connection success rates subject to transmission range

distance	success rate	
	without obstacles	with obstacles
3 m	100%	100%
6 m	80%	70%
10 m	20%	0%
12 m	0%	0%

the case in underground trains. As soon as the internet connection is re-established, detailed movie information such as the movie genre, cast or, trailer can be queried via movieIDs.

5 EVALUATION

We evaluate the prototype implementation in view of its potential to facilitate the proposed pervasive recommender system sketched in Section 3 in particular in urban areas. We conducted six distinct experiments, where every experiment was conducted with Nexus 5 smartphones that support Bluetooth up to version 4 (including BLE) and Wifi direct. We re-ran experiments 10 times by default, unless stated otherwise, with distinct experiment setups in order to have empirical evidence of the results' soundness.

- (1) *Share Large Amounts of Ratings*: We share a bulk of 1000 ratings, accounting to roughly 100 kB uncompressed. As soon as the connection is established, all 1000 ratings get transmitted instantly and reliably without any data loss.
- (2) *Share Data Between Multiple Devices*: We share ratings between four smartphones simultaneously holding mutually disjoint ratings. Ratings get transmitted correctly and losslessly.
- (3) *Share Data in Public Transportation*: We successfully share ratings between three devices in the bus and the underground in the urban city of Berlin being exposed to many WiFi and Bluetooth disturb signals. The loss of internet connectivity does not impair the data transfer. As soon as internet connectivity is re-established, the shared placeholder recommendations get filled with movie metadata fetched from the Open Movie Database⁹ (OMDb) API.
- (4) *Effective Transmission Range*: Recall that the Nearby Connections API comprises Bluetooth, BLE, and Wifi direct under its hood. All three allegedly provide an effective transmission radius of 10 meters for class 2 devices such as smartphones. We tested transmission ranges between 3 and 12 meters, both outdoors and indoors, and with and without obstacles. The results are shown in Table 1. We conclude that the effective radius of ratings propagation is between 3 and 6 meters.
- (5) *Average Initial Connection Delay*: We measure the initial connection delay for two devices, that is the elapsed time before a connection gets established. We place the devices at a distance of 1 meter in order to guarantee connectivity, where the connections were made with distinct app sessions (partner device's id not in cache). The average connection delay

⁹<http://www.omdbapi.com/>

Table 2: Average battery drain of pre-connection (advertising and discovery)

scenario	Battery Drain	Battery Logger	Average
sharing on	5.94 %/h	5.59 %/h	5.77 %/h
sharing off	0.34 %/h	0.66 %/h	0.50 %/h
factory settings	0.51 %/h	0.49 %/h	0.50 %/h
sharing only	5.6 %/h	4.93 %/h	5.27 %/h

is 25.9 seconds, with a minimum and maximum connection delay of 11 and 41 seconds respectively.

- (6) *Battery Drainage*: Since the advertising and discovery step (pre-connection) has to be performed continuously, we would like to know whether or not advertising, discovering, and sharing information in the background for extended periods of time is feasible or not. Since we cannot experiment in a real-life scenario with at times no devices to connect to, and then multiple devices to connect to, we only measure the application's pre-connection battery drainage, which therefore presents a lower bound for battery drainage. We use two devices reset to factory settings. We track the battery levels for three distinct scenarios (a) application running in the background with sharing on, (b) with sharing off, and (c) factory settings, where in all three cases the displays were off. The results are shown in Table 2.

5.1 Conclusion

The experimental findings indicate that the Propagate and Filter's propagation step works reliably with larger amounts of ratings, in multi-device scenarios, and in areas without internet connectivity such as underground trains. Yet, the implementation has certain limitations. First of all, information can only be disseminated reliably within a radius of 6 meters, which on the one hand strengthens privacy, and on the other limits the number of potential peers to share information with. Furthermore, the average initial connection delay of 25.9 seconds is considerable and does not allow to share information between for instance passing pedestrians, yet includes scenarios such as waiting at the traffic lights, sitting next to each other in a café or restaurant, or taking public transportation. Last but not least, battery drainage is relatively high at at least 5% per hour, which thus prohibits continuous advertising and discovery. Apart from ever improving transmission standards and hardware, we believe that it is possible to limit activity and inactivity of the sharing process efficiently on the logical level leveraging sensor information on smartphones.




6 FUTURE WORK

Future work includes the launch of the mobile application for the collection of usage data. Recall that we left the sampling process and the recommendation algorithm as placeholders (see Section 3.3 (2) and (3) respectively). Once usage data is collected it will be possible to test distinct combinations of sampling and recommendation strategies. Furthermore, the evaluation of feedback on the application including the user interface and user experience is due.

REFERENCES

- [1] Gediminas Adomavicius, Ramesh Sankaranarayanan, Shahana Sen, and Alexander Tuzhilin. 2005. Incorporating Contextual Information in Recommender Systems Using a Multidimensional Approach. *ACM Trans. Inf. Syst.* 23, 1 (Jan. 2005), 103–145. <https://doi.org/10.1145/1055709.1055714>
- [2] Gediminas Adomavicius and Alexander Tuzhilin. 2011. *Context-Aware Recommender Systems*. Springer US, Boston, MA, 217–253. https://doi.org/10.1007/978-0-387-85820-3_7
- [3] Ranieri Baraglia, Patrizio Dazzi, Matteo Mordacchini, and Laura Ricci. 2013. A peer-to-peer recommender system for self-emerging user communities based on gossip overlays. *J. Comput. System Sci.* 79, 2 (2013), 291–308. <https://doi.org/10.1016/j.jcss.2012.05.011> 10th IEEE International Conference on Computer and Information Technology, 2010.
- [4] Felix Beierle. 2018. Do You Like What I Like? Similarity Estimation in Proximity-based Mobile Social Networks. In *Proc. 2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications (TrustCom)*. IEEE, 1040–1047. <https://doi.org/10.1109/TrustCom/BigDataSE.2018.00146>
- [5] Felix Beierle and Tobias Eichinger. 2019. Collaborating with Users in Proximity for Decentralized Mobile Recommender Systems. In *Proceedings of the IEEE 16th International Conference on Ubiquitous Intelligence and Computing (UIC)*. IEEE (in press).
- [6] Felix Beierle, Vinh Thuy Tran, Mathias Allemand, Patrick Neff, Winfried Schlee, Thomas Probst, Rüdiger Pruss, and Johannes Zimmermann. 2018. Context Data Categories and Privacy Model for Mobile Data Collection Apps. *Procedia Computer Science* 134 (2018), 18–25. <https://doi.org/10.1016/j.procs.2018.07.139>
- [7] Antoine Boutet, Stephane Frenot, Frederique Laforest, Pascale Launay, Nicolas Le Sommer, Yves Maheo, and Damien Reimert. 2015. C3PO: A Network and Application Framework for Spontaneous and Ephemeral Social Networks. In *Web Information Systems Engineering – WISE 2015 (Lecture Notes in Computer Science)*. Springer International Publishing, 348–358. https://doi.org/10.1007/978-3-319-26187-4_33
- [8] Leo Breiman. 1996. Bagging Predictors. *Machine Learning* 24, 2 (01 Aug 1996), 123–140. <https://doi.org/10.1023/A:1018054314350>
- [9] L. Chancay-García, E. Hernández-Orallo, P. Manzoni, C. T. Calafate, and J. Cano. 2018. Evaluating and Enhancing Information Dissemination in Urban Areas of Interest Using Opportunistic Networks. *IEEE Access* 6 (2018), 32514–32531. <https://doi.org/10.1109/ACCESS.2018.2846201>
- [10] Fei Chen, Zhenhua Dong, Zhenguo Li, and Xiuqiang He. 2018. Federated Meta-Learning for Recommendation. *CoRR abs/1802.07876* (2018). [arXiv:1802.07876](http://arxiv.org/abs/1802.07876)
- [11] Alexandre de Spindler, Moira C. Norrie, and Michael Grossniklaus. 2007. Collaborative filtering based on opportunistic information sharing in mobile ad-hoc networks. In *On the Move to Meaningful Internet Systems 2007: CoopIS, DOA, ODBASE, GADA, and IS (Lecture Notes in Computer Science)*. Springer Berlin Heidelberg, Berlin, Heidelberg, 408–416. https://doi.org/10.1007/978-3-540-76848-7_28
- [12] Juan Vera del Campo, Josep Pegueroles, Juan Hernández-Serrano, and Miguel Soriano. 2012. Design of a P2P content recommendation system using affinity networks. *Computer Communications* 36, 1 (2012), 90–104. <https://doi.org/10.1016/j.comcom.2012.07.018>
- [13] L. Del Prete and L. Capra. 2010. diffERS: A Mobile Recommender Service. In *2010 Eleventh International Conference on Mobile Data Management*. 21–26. <https://doi.org/10.1109/MDM.2010.22>
- [14] B. Efron. 1979. Bootstrap Methods: Another Look at the Jackknife. *Ann. Statist.* 7, 1 (01 1979), 1–26. <https://doi.org/10.1214/aos/1176344552>
- [15] Tobias Eichinger, Felix Beierle, Sumsam Ullah Khan, Robin Middelanis, Veer-araghavan Sekar, and Sam Tabibzadeh. 2019. affinity: A System for Latent User Similarity Comparison on Texting Data (accepted article) (*IEEE ICC 2019*). IEEE. <http://arxiv.org/abs/1904.01897>
- [16] Damianos Gavalas, Charalampos Konstantopoulos, Konstantinos Mastakas, and Grammati Pantziou. 2014. Mobile Recommender Systems in Tourism. *J. Netw. Comput. Appl.* 39, C (March 2014), 319–333. <https://doi.org/10.1016/j.jnca.2013.04.006>
- [17] B. Heep, M. Florian, J. Volz, and I. Baumgart. 2013. Overdrive an overlay-based geocast service for smart traffic applications. In *2013 10th Annual Conference on Wireless On-demand Network Systems and Services (WONS)*. 1–8. <https://doi.org/10.1109/WONS.2013.6578313>
- [18] Márk Jelasity and Ozalp Babaoglu. 2006. T-Man: Gossip-Based Overlay Topology Management. In *Engineering Self-Organising Systems*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1–15. https://doi.org/10.1007/11734697_1
- [19] Márk Jelasity, Spyros Voulgaris, Rachid Guerraoui, Anne-Marie Kermerrec, and Maarten van Steen. 2007. Gossip-based Peer Sampling. *ACM Trans. Comput. Syst.* 25, 3, Article 8 (Aug. 2007). <https://doi.org/10.1145/1275517.1275520>
- [20] Salma Ktari and Artur Hecker. 2011. A Peer-to-Peer Social Network Overlay for Efficient Information Retrieval and Diffusion. In *Future Information Technology*. Springer Berlin Heidelberg, Berlin, Heidelberg, 24–33. https://doi.org/10.1007/978-3-642-22309-9_4
- [21] S. Magureanu, N. Dokoohaki, S. Mokarizadeh, and M. Matskin. 2012. Epidemic Trust-Based Recommender Systems. In *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing*. 461–470. <https://doi.org/10.1109/SocialCom-PASSAT.2012.94>
- [22] H. Brendan McMahan, Eider Moore, Daniel Ramage, and Blaise Agüera y Arcas. 2016. Federated Learning of Deep Networks using Model Averaging. *CoRR abs/1602.05629* (2016).
- [23] Christos Mettouris and George Papadopoulos. 2014. Ubiquitous recommender systems. *Computing* 96 (03 2014), 223–257. <https://doi.org/10.1007/s00607-013-0351-z>
- [24] Erik Nordström, Per Gunningberg, and Christian Rohner. 2009. Haggie : A Data-centric Network Architecture for Mobile Devices. Invited talk.
- [25] L. Nunes Barbosa, J. Gemmel, M. Horvath, and T. Heimfarth. 2018. Distributed User-Based Collaborative Filtering on an Opportunistic Network. In *2018 IEEE 32nd International Conference on Advanced Information Networking and Applications (AINA)*. 266–273. <https://doi.org/10.1109/AINA.2018.00049>
- [26] Anna-Kaisa Pietiläinen, Earl Oliver, Jason LeBrun, George Varghese, and Christophe Diot. 2009. MobiClique: Middleware for Mobile Social Networking. In *Proceedings of the 2Nd ACM Workshop on Online Social Networks (WOSN '09)*. ACM, New York, NY, USA, 49–54. <https://doi.org/10.1145/1592665.1592678>
- [27] Nikolaos Polatidis and Christos K. Georgiadis. 2015. A Ubiquitous Recommender System Based on Collaborative Filtering and Social Networking Data. *Int. J. Intell. Eng. Inform.* 3, 2/3 (June 2015), 186–204. <https://doi.org/10.1504/IJIEL.2015.069895>
- [28] Giancarlo Ruffo and Rossano Schifanella. 2009. A peer-to-peer recommender system based on spontaneous affinities. *ACM Trans. Internet Techn.* 9, 1 (2009), 4:1–4:34. <https://doi.org/10.1145/1462159.1462163>
- [29] Rossano Schifanella, André Panisson, Cristina Gena, and Giancarlo Ruffo. 2008. MobHunter: Epidemic Collaborative Filtering and Self-organization in Mobile Ad-hoc Networks. In *Proceedings of the 2008 ACM Conference on Recommender Systems (RecSys '08)*. ACM, New York, NY, USA, 27–34. <https://doi.org/10.1145/1454008.1454014>
- [30] O. Yurur, C. Liu, Z. Sheng, V. Leung, W. Moreno, and K. Leung. 2014. Context-Awareness for Mobile Sensing: A Survey and Future Directions. *IEEE Communications Surveys Tutorials* 18, 1 (2014), 1–28. <https://doi.org/10.1109/COMST.2014.2381246>
- [31] Cai-Nicolas Ziegler. 2005. *Towards decentralized recommender systems*. Ph.D. Dissertation. Albert-Ludwigs-Universität Freiburg.

Additional Information

Bibliographic Data	T. Eichinger, F. Beierle, R. Papke, L. Rebscher, H. Chinh Tran, and M. Trzeciak, "On Gossip-based Information Dissemination in Pervasive Recommender Systems," in <i>Proc. 13th ACM Conference on Recommender Systems (RecSys)</i> . ACM, 2019.
Pre-print from	https://beierle.de
Online at	(in press) https://doi.org/10.1145/3298689.3347067
Authors	<p>Tobias Eichinger </p> <p>Felix Beierle  </p> <p>Robin Papke Lucas Rebscher Hong Chinh Tran Magdalena Trzeciak</p>
BibTeX	<pre>@inproceedings{Eichinger2019RecSys, title = {{On Gossip-based Information Dissemination in Pervasive Recommender Systems}}, author = {Eichinger, Tobias and Beierle, Felix and Papke, Robin and Rebscher, Lucas and Chinh Tran, Hong and Trzeciak, Magdalena}, booktitle = {{Proc. 13th ACM Conference on Recommender Systems (RecSys)}}, publisher = {ACM}, doi = {10.1145/3298689.3347067}, year = {2019} }</pre>
Copyright Note	<p>ACM Copyright Notice Copyright (c) 2019 ACM</p> <p>Permission to make digital or hard copies of all or part 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 bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.</p>