IRIS AperTO

UNIVERSITÀ DEGLI STUDI DI TORINO

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

A Compositional Model of Multi-faceted Trust for Personalized Item Recommendation

This is the author's manuscript

Original Citation:

Availability:

This version is available http://hdl.handle.net/2318/1710621 since 2022-07-27T10:30:35Z

Published version:

DOI:10.1016/j.eswa.2019.112880

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

12 October 2024

Published in Expert Systems with Applications, Elsevier. DOI: https://doi.org/10.1016/j.eswa.2019.112880.
Link to the page of the paper on Elsevier web site:
http://www.sciencedirect.com/science/article/pii/S0957417419305901

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit http://creativecommons.org/licenses/by-nc-nd/4.0/ or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

Highlights

A Compositional Model of Multi-faceted Trust for Personalized Item Recommendation

Liliana Ardissono[1] and Noemi Mauro

- We propose a compositional recommender system based on multi-faceted trust.
- The trust model is based on social links and global feedback about users.
- The recommender can work with or without using information about social relations.
- We validate our recommendation model on two public datasets of item reviews.
- Our recommender outperforms state-of-the-art trust-based recommender systems.

[1]Corresponding author.
E-mail addresses: liliana.ardissono@unito.it (L. Ardissono), noemi.mauro@unito.it (N. Mauro).

Published in Expert Systems with Applications, Elsevier. DOI: https://doi.org/10.1016/j.eswa.2019.112880.
Link to the page of the paper on Elsevier web site:
http://www.sciencedirect.com/science/article/pii/S0957417419305901

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit http://creativecommons.org/licenses/by-nc-nd/4.0/ or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

A Compositional Model of Multi-faceted Trust for Personalized Item Recommendation

Liliana Ardissono[1] and Noemi Mauro

Computer Science Department, University of Torino
Corso Svizzera, 185, I-10149, Torino, Italy

Abstract

Trust-based recommender systems improve rating prediction with respect to Collaborative Filtering by leveraging the additional information provided by a trust network among users to deal with the cold start problem. However, they are challenged by recent studies according to which people generally perceive the usage of data about social relations as a violation of their own privacy. In order to address this issue, we extend trust-based recommender systems with additional evidence about trust, based on public anonymous information, and we make them configurable with respect to the data that can be used in the given application domain:

1. We propose the Multi-faceted Trust Model (MTM) to define trust among users in a compositional way, possibly including or excluding the types of information it contains. MTM flexibly integrates social links with public anonymous feedback received by user profiles and user contributions in social networks.

[1]Corresponding author.
E-mail addresses: liliana.ardissono@unito.it (L. Ardissono), noemi.mauro@unito.it (N. Mauro).

2. We propose LOCABAL+, based on MTM, which extends the LOCABAL trust-based recommender system with multi-faceted trust and trust-based social regularization.

Experiments carried out on two public datasets of item reviews show that, with a minor loss of user coverage, LOCABAL+ outperforms state-of-the art trust-based recommender systems and Collaborative Filtering in accuracy, ranking of items and error minimization both when it uses complete information about trust and when it ignores social relations. The combination of MTM with LOCABAL+ thus represents a promising alternative to state-of-the-art trust-based recommender systems.

Keywords:
Multi-faceted trust, Trust-based Recommender Systems, Compositional trust model, Matrix Factorization

---

Declarations of interest: none.

1. Introduction

Trust-based recommender systems improve rating prediction with respect to Collaborative Filtering (Desrosiers and Karypis, 2011) by combining rating similarity with the additional information provided by a trust network among users to deal with the cold start problem. Most of these systems predict the rating scores that a person would attribute to items by relying on the observed preferences of the users who are linked to her/him by social relations, directly or through a short path of links, as in SocialMF (Jamali and Ester, 2010). Moreover, having observed that, in the physical world, people

are likely to seek advice from both local friends and highly reputable users, some systems also take global reputation into account to improve recommendation performance. For instance, LOCABAL (Tang et al., 2013) computes users' reputation as a function of their importance in the social network and exploits this data to weight the impact of ratings in Matrix Factorization.

Despite the good recommendation results achieved by trust-based recommender systems, recent studies show that they are hardly accepted by people, who are concerned about the storage of personal information and the access to social relations (Burbach et al., 2018). It is thus vital to define trust models that can use data which is not perceived as personal.

For this purpose, we propose a compositional trust model and recommender system which rely on complementary information sources to obtain a twofold objective: (i) collecting rich evidence about user trust to improve Top-N recommendation, and (ii) adapting to possible restrictions on the information that can be used in the application domain of interest.

Indeed, various signs of trust can be used to compute users' global reputation without relying on sensitive information. For instance, social networks such as Booking.com (2019), Expedia.com (2001) and Yelp (2019a) publish anonymous feedback about users (expressed as endorsements to their profiles) and about their contributions (e.g., helpfulness of reviews) that can be used to assess reputation by ignoring the identity of the people who provided it. It is thus interesting to define a model that supports the interpretation of these types of feedback as an overall trustworthiness measure.

In this article we present the Multi-faceted Trust Model (MTM) as a framework to fuse *local trust between users* (inferred from direct social rela-

tions) with the following sources of information:

- The *quality of individual reviews* derived from the explicit feedback they receive from the social network.
- *Multi-dimensional user reputation* derived from the analysis and integration of different types of endorsements that users can receive with the quality of the reviews they author.

MTM makes it possible to separately include or exclude the components of trust to assess their relative impact on recommendation. This supports the evaluation of performance, e.g., when different types of anonymous feedback are considered, and when social relations are ignored.

We integrate MTM into a novel trust-based recommendation algorithm, denoted as LOCABAL+, which combines local trust and multi-dimensional global reputation in preference estimation. LOCABAL+ extends the LOCABAL recommender system, from which we take inspiration, as follows:

- It tunes Matrix Factorization by exploiting multi-faceted trust, which takes multiple aspects of user behavior into account, instead of only relying on social links.
- It regularizes social relations by means of rating similarity and multi-dimensional global reputation to exploit both properties in the selection of the like-minded users for rating prediction.
- It can be configured to use a subset of the facets of trust.

Experimental results show that LOCABAL+ achieves the best accuracy, error minimization and ranking results when it uses both global trust feedback

and social relations. However, it also outperforms state-of-the-art recommender systems based on Matrix Factorization and on K-Nearest Neighbors when it ignores social relations. We thus conclude that multi-faceted trust enhances recommendation performance in trust-based recommenders and it makes them more flexible with respect to the types of information that can be used in a specific application domain.

This work extends the preliminary multi-faceted trust model presented in (Mauro et al., 2019) as follows: firstly, we integrate the facets of trust in Matrix Factorization, instead of using a K-Nearest Neighbors model. Secondly, we perform more detailed and extensive experiments to evaluate recommendation performance: (i) we analyze the impact of the facets of trust on recommendation by tuning the components of MTM in a finer-grained way; (ii) we integrate all the analyzed recommender systems in the same framework to uniformly evaluate the performance of algorithms; (iii) we evaluate recommendation performance on new data (which was not previously used for training/validating the models) to comply with the application domains where models cannot be frequently optimized.

In the remainder of this paper, Section 2 presents our research questions and outlines the experiments to answer them. Section 3 provides background concepts and positions our work in the related one. Sections 4 and 5 present MTM and LOCABAL+. Section 6 describes the datasets used for the experiments and the instantiation of MTM on the available types of information. Section 7 presents the validation methodology we used and the evaluation results. Section 8 discusses the evaluation results and outlines our future work. Section 9 concludes the paper.

2. Research questions and experimental plan

In the Multi-faceted Trust Model we integrate diverse facets of trust and, in a specific application domain, one or more of them might not be available or usable. Therefore, besides assessing their overall value in improving Top-N recommendation, we separately study their impact on recommendation performance. We thus formulate the following research questions:

RQ1: *Can multi-faceted trust be used to improve the performance of a trust-based recommender system with respect to only relying on social links and rating similarity among users?*

RQ2: *What is the impact of the multi-dimensional reputation of users, of the quality of their contributions, and of social links, on collaborative recommendation performance?*

In order to answer these questions, we carry out experiments to measure the performance of LOCABAL+ on a spectrum of MTM configurations that tune in different ways the influence of the facets of trust we consider. We compare the performance of the algorithm when using or ignoring social links and trust statements about users and their contributions. We also compare the algorithm with the following baselines: the LOCABAL (Tang et al., 2013) and SocialMF (Jamali and Ester, 2010) trust-based recommender systems based on Matrix Factorization; SVD++ (Koren, 2008) and User-to-User Collaborative Filtering (Desrosiers and Karypis, 2011) which only use rating similarity; a user-to-user Collaborative Filtering algorithm (henceforth denoted as U2USocial) that employs friend relations to estimate ratings in a K-Nearest Neighbors approach, instead of using rating similarity.

We carry out the experiments on two large subsets the Yelp dataset (Yelp, 2019b). The first one, Yelp-Hotel, concerns accommodation facilities; the second one, Yelp-Food, is focused on restaurants.

We evaluate Top-k recommendation performance of algorithms with k=10 by taking the rating scores of the dataset as ground truth. Following the recent trends in the evaluation of recommender systems described in (Jannach et al., 2016), we measure their accuracy, error minimization, ranking capability and user coverage @k; see Section 7.1 for details.

3. Background and related work

3.1. Basic concepts: trust and reputation

Trust is generally described as a positive expectation that an agent has about other agents' behavior, from a subjective perspective. Gambetta (1988) defines it as "a particular level of the subjective probability with which an agent or group of agents will perform a particular action, both before [we] can monitor such action (or independently of [our] capacity of ever to be able to monitor it) and in a context in which it affects [our] own action". Moreover, both Gambetta (1988) and Golbeck and Hendler (2004) specify that a user trusts another one in a social network if (s)he believes that any future transaction with her/him will be rewarding rather than detrimental. On a more general perspective, Mui et al. (2002) (and similarly Misztal (1996)) elect "subjective probability" to subjective expectation, or degree of belief, to highlight that, more than a statistical probability, trust represents a belief status that an agent A has about another agent B's future behavior, given B's past behavior and her/his *reciprocity* of action within a society.

Different from trust, *reputation* describes a general "expectation about an agent's behaviour based on information about or observations of its past behaviour" (Abdul-Rahman and Hailes, 2000). Reputation has a global perspective and Mui et al. (2002) describe it as the "perception that an agent creates through past actions about its intentions and norms". According to Misztal (1996), reputation "helps us to manage the complexity of social life by singling out trustworthy people - in whose interest it is to meet promises".

While the previously described works analyze trust and reputation from the global viewpoint of agent-to-agent interaction, a few ones contextualize it in online collaboration systems and social networks, which are the scope of our present work. Noticeably, Mcnally et al. (2014) generalize trust relations by analyzing the occurrence of collaboration events that involve users; this makes it possible to link users because they have downloaded or bookmarked contents provided by other users, and so forth. These authors explain that reputation can derive from direct user-to-user interaction (e.g., when users are rated) or from indirect one; e.g., when they interact by virtue of some item. Moreover, it can derive from explicit trust statements, such as ratings, or from implicit ones like follower relations.

Before describing the state of the art on trust-based recommender systems, it is worth briefly discussing the main issues affecting them. Specifically, it may be questioned whether relying on social relations and global feedback about users is a safe approach to evaluate trustworthiness. Some trust-based recommender systems focus on user-to-user relations and ignore the feedback on user actions because there is a general opinion that the latter could be biased. While we obviously agree that this may be true, we point

out that *any* type of action that brings evidence about trust, including the establishment of friend relations, ratings, etc., could be performed with the aim of manipulating the reputation of some user. Therefore, data reliability assessment is a general pre-requisite for the development of recommender systems. Indeed, the weaknesses of some models adopted in e-commerce and collaboration sites have been analyzed to suggest how to improve the robustness of Reputation Management Systems; e.g., see (Resnick and Zeckhauser, 2002). However, Jøsang et al. point out that these systems are challenged by strategic manipulation and by various types of attacks which cannot always be detected by statistical analyses (Jøsang et al., 2007; Jøsang and Goldbeck, 2009). Therefore, Jøsang (2012) ultimately highlights the importance of strengthening legislation as a barrier to discourage malicious behavior.

3.2. Trust-based recommender systems

The homophily (McPherson et al., 2001) and social influence (Marsden and Friedkin, 1993) theories associate social links to user similarity. On this basis, social and trust-based recommender systems (Richthammer et al., 2017) exploit social networks as additional sources of information to complement rating data. These systems estimate user preferences by relying on the known social links existing between people; e.g., friend, follower and/or trust relations according to different inference techniques; see Table 1:

- AVG: average rating of selected (e.g., trusted) social links (Golbeck and Hendler, 2004, 2006; Liu and Lee, 2010; Parvin et al., 2019).
- KNN: K-Nearest Neighbors on social links (O'Donovan and Smyth, 2005; Massa and Avesani, 2007; Groh and Ehmig, 2007; Moradi and

Table 1: Overview of the main social and trust-based recommender systems cited in the paper; ACO stands for Ant Colony Optimization.

Citation	Algorithm	Technology	User filtering	Social regularization	Reputation
(Golbeck and Hendler, 2006)	FilmTrust	AVG	by trust	-	-
(Kuter and Golbeck, 2007)	SUNNY	AVG	by trust	probabilistic	-
(De Meo et al., 2018)	LGTR	AVG	social proximity	-	by feedback
(Parvin et al., 2019)	TCFACO	AVG	by trust	ACO	-
(O'Donovan and Smyth, 2005)	-	KNN	by trust	-	-
(Liu and Lee, 2010)	-	KNN	-	-	-
(Moradi and Ahmadian, 2015)	RTCF	KNN	by trust and rating similarity	-	-
(Jamali and Ester, 2009)	TrustWalker	Random Walk	by trust	-	-
(Deng et al., 2014)	RelevantTrustWalker	Random Walk	by trust relevancy	-	-
(Jamali and Ester, 2010)	SocialMF	MF	-	rating-based	-
(Yang et al., 2012)	CircleCon	MF	by circle	-	-
(Tang et al., 2013)	LOCABAL	MF	-	rating based	PageRank
(Guo et al., 2015)	TrustSVD	MF	-	-	-
(Yang et al., 2017)	TrustMF	MF	-	rating based	-
(Qian et al., 2016)	SoRS	MF	-	rating based	by rating conformity
(Yuan et al., 2011)	MF.FM	Collective MF	by group	rating based	-
(Ma et al., 2011b)	SOREG	MF	-	rating based	-
(Ma et al., 2011a)	RSTE	PMF	-	rating based	-
(Liu and Aberer, 2013)	SoCo	PMF	by context	-	-
(Du et al., 2017)	SIACC	Co-Clustering	-	rating based	PageRank

Ahmadian, 2015; Ardissono et al., 2017).

- MF: Matrix Factorization (in some cases with Random Walk) on the matrices of ratings and social links (Jamali and Ester, 2009, 2010; Ma et al., 2011b; Yang et al., 2012; Tang et al., 2013; Deng et al., 2014; Guo et al., 2015; Yang et al., 2017).

- PMF: Probabilistic Matrix Factorization on the matrices of ratings and social links (Ma et al., 2011a,c; Jiang et al., 2012; Liu and Aberer, 2013; Chaney et al., 2015).

- Probabilistic approaches on trust networks (Kuter and Golbeck, 2007;

Li et al., 2014).

- Co-clustering of ratings and trust matrices (Du et al., 2017).

Differently, we generate personalized recommendations by relying on a compositional, multi-faceted trust model that includes complementary data about user behavior: i.e., not only local trust among users inferred from social relations, but also quality of user contributions (derived from the anonymous global feedback they receive) and multi-dimensional global reputation derived from diverse types of information, among which anonymous endorsements to user profiles. The integration of these facets of trust supports a rich computation of reputation based on complementary aspects of user behavior. Moreover, it makes it possible to compensate trust evidence in application domains in which some types of information are not available or cannot be used. In particular, LOCABAL+ works with or without using social links.

Some research about recommender systems studies the differences between trust and friends networks. Guo et al. (2015), Ma et al. (2011b) and Li et al. (2018) find out that, different from explicit trust relations (such as those among Epinions users (Epinions, 2019)), friendship does not strictly imply preference similarity: user preferences are strongly correlated among trusted neighbors but they are only slightly positively correlated among "trust-alike" neighbors such as friends in social networks (Guo et al., 2015). Several authors recognize the importance of limiting the social context to the user's local proximity; for instance, Massa and Avesani (2007) and Yuan et al. (2011) prove that recommendation accuracy decreases when indirect social connections (i.e., paths of social links) are used to estimate user preferences. Moreover, Yang et al. (2012) point out that users may trust different sub-

sets of friends regarding different domains. In order to deal with this issue, authors propose various methods to filter the neighboorhood used for rating prediction; e.g., Yang et al. (2017) use category-specific circles and Yuan et al. (2011) use thematic groups to steer Matrix Factorization. Moreover, KNN and AVG systems select neighbors by ranking the users directly linked to the current user in the social/trust network on the basis of their rating similarity (Massa and Avesani, 2007; Liu and Lee, 2010; Li et al., 2014; Moradi and Ahmadian, 2015; Ardissono et al., 2017; Parvin et al., 2019). Analogously, social regularization is used to increase the impact of like-minded users in Matrix Factorization: e.g., TrustMF (Yang et al., 2017) applies social regularization to users' direct social links and Yuan et al. (2011) applies it to the members of thematic groups; RSTE (Ma et al., 2011a) and SOREG (Ma et al., 2011b) integrate trust and rating similarity in Probabilistic Matrix Factorization, and Ma et al. (2011c) use tag-based similarity to build a larger social context for regularization. Finally, SocialMF (Jamali and Ester, 2010) employs rating similarity to regularize the impact of users who are reachable through a short path of social links in Random Walk. Other systems achieve similar filtering results by combining trust-based and item-based recommendation, as in TrustWalker (Jamali and Ester, 2009), or by filtering the users of the trust networks according to rating similarity, as in TCFACO (Parvin et al., 2019) and RelevantTrustWalker (Deng et al., 2014). Finally, Du et al. (2017) apply co-clustering to the matrices of ratings and social relations in order to identify like-minded users within social connections.

In comparison, we extend social regularization by tuning the impact of users on the Matrix Factorization process on the basis of both rating similar-

ity and global multi-dimensional reputation. In other words, we select neighbors for rating prediction by privileging users who are trustworthy *and* like-minded. This approach improves prediction accuracy because it enhances the quality of the rating information used to estimate preferences.

Building on social influence theories, Guo et al. (2015) propose TrustSVD that extends SVD++ (Koren, 2008) to jointly factorize the rating and trust matrices: they learn a truster model that describes how people are influenced in item evaluation by their parties's opinions. TrustMF (Yang et al., 2017) learns both the truster and trustee models to consider the fact that, in a social network, people mutually influence each other. Jiang et al. (2012) investigate the relation between social influence and personal preferences. Finally, some researchers leverage the local and global perspectives of social influence building on the observation that in the physical world humans ask for opinions from both local friends and highly reputable people; e.g., Qian et al. (2016); Tang et al. (2013); Hu et al. (2018). Specifically, in LOCABAL, Tang et al. (2013) combine rating similarity and social links with users' global reputation, which is based on the PageRank (Page et al., 1999) score as a measure of importance in the social network.

We use both local trust and global reputation to steer personalized recommendation. However, we propose a multi-faceted trust model to determine users' trustworthiness on the basis of complementary types of information about their behavior, including social relations and global feedback from other users. As previously noticed, this makes it possible to tune Matrix Factorization on the basis of the contributions of trustworthy, like-minded people, using public, anonymous information. To the best of our knowledge,

the only other work that employs global feedback about users is LGTR by De Meo et al. (2018), which defines global reputation on the basis of the feedback collected by user actions and of a local context depending on social relations. However, in neighbor identification, LGTR discards rating similarity, which is very useful to select like-minded users for preference prediction. Moreover, LGTR does not take review quality and endorsements to user profiles into accout. Our model is thus more general than this one.

Some trust-based recommender systems assume the existence of both positive and negative evidence about trust as, e.g., in the social networks where users can rate other users positively or negatively (Li et al., 2011; Victor et al., 2011; Rafailidis and Crestani, 2017). In our work, we start from the consideration that the trust models provided by several social networks are only based on the expression of "likes". Therefore, we propose a model that can also work on positive-only feedback to comply with them.

Some works associate users' reputation to rating conformity; e.g., O'Donovan and Smyth (2005) and Li et al. (2013) base reputation on the percentage of ratings provided by a user that agree with those of the other people. Su et al. (2017) cluster users on the basis of rating similarity and consider the largest cluster as the "honest" group. In the SoRS recommender system Qian et al. (2016) derive reputation by iteratively calculating the correlation of the historical ratings provided by a user and the quality of items emerging from the rating scores they receive. However, review conformity does not fully characterize quality; e.g., Victor et al. (2011) point out that controversial reviews must be considered and matched to individual preferences. We thus leave this aspect for our future work.

4. Multi-faceted Trust Model (MTM)

MTM is aimed at computing users' trustworthiness in the context of recommender systems. It integrates local trust between users (inferred from social relations, in line with trust-based recommender systems research) with the public, anonymous feedback received by users and by their contributions in a social network. MTM is compositional and supports the inclusion or exclusion of facets of trust to comply with the requirements of the application domain of interest. We identified the classes of evidence about trust of MTM by analyzing the information publicly provided by social networks and e-commerce sites such as Yelp (2019a), Booking.com (2019), Expedia.com (2001), LibraryThing (2019), Amazon.com (2006), Ciao (Danetsoft, 2019) and Epinions (2019). However, we generalized those indicators to enhance the applicability of our model to heterogeneous domains. In the following we describe each class in detail.

4.1. Quality of individual contributions on an item

By individual contribution on an item we mean a piece of information that a user provides about it. A contribution is usually a review associated with a rating score but we describe our model at a more general level because in some online services users can post different types of content; e.g., the Yelp social network allows to write both reviews and tips about items.

Let $\mathcal{U}$ be the set of users and $\mathcal{I}$ the set of items of a service. Given $v \in \mathcal{U}$ and $i \in \mathcal{I}$ we denote an individual contribution provided by v on i as $contr_{vi}$. Then, we define the quality of $contr_{vi}$ ($fContr_{vi}$, in [0, 1]) on the basis of the amount of feedback that $contr_{vi}$ has received. We measure quality in

a relative way with respect to the most popular contributions on the same item in order to be robust with respect to item-specific biases:

- In the social networks that provide both positive and negative feedback about contributions we take inspiration from the definition of gold standard helpfulness of reviews defined in (Kim et al., 2006; Raghavan et al., 2012; O'Mahony and Smyth, 2018). In those works, helpfulness is defined as the ratio between the number of positive evaluations ($positiveVotes_{contr_{vi}}$) and the total number of evaluations ($positiveVotes_{contr_{vi}}$+$negativeVotes_{contr_{vi}}$) that a contribution receives:

$$helpfulness_{contr_{vi}} = \frac{positiveVotes_{contr_{vi}}}{positiveVotes_{contr_{vi}} + negativeVotes_{contr_{vi}}} \quad (1)$$

  We define the quality of $contr_{vi}$ in a relative way with respect to the best contribution on the same item as follows:

$$fContr_{vi} = \frac{helpfulness_{contr_{vi}}}{\max_{a \in \mathcal{U}} helpfulness_{contr_{ai}}} \quad (2)$$

- In the social networks that only support positive feedback we compute quality as the ratio between the overall number of appreciations obtained by $contr_{vi}$ ($appreciations_{contr_{vi}}$) and the maximum number of appreciations received by the other contributions on the same item:

$$fContr_{vi} = \frac{appreciations_{contr_{vi}}}{\max_{a \in \mathcal{U}} appreciations_{contr_{ai}}} \quad (3)$$

4.2. Multi-dimensional global reputation

We define *multi-dimensional global reputation* building on heterogeneous types of information about users to capture different aspects of their behavior. Let $\mathcal{U}$ be the set of users, $\mathcal{I}$ the set of items, $v \in \mathcal{U}$ and $i \in \mathcal{I}$. We define the following sub-classes of trust evidence:

(P) *Importance of the user in the social network* (imp_v in [0, 1]), based on her/his social connections. Similar to LOCABAL, we use PageRank (Page et al., 1999) to model this type of indicator. PageRank estimates the relative importance of nodes in a graph by counting the number and quality of the links that enter them, under the assumption that being referenced by others is a quality sign. We compute imp_v as:

$$imp_v = \frac{1}{1 + log(rank_v)} \tag{4}$$

where $rank_v \in [1, |U|]$ is the PageRank score of v and the most important user is ranked with 1; see (Tang et al., 2013).

(U) *Global feedback about the user's profile:*

- *User profile endorsements and public recognition* ($fEndors_v$, in [0, 1]). This class represents the global types of feedback that user profiles receive from the social network. It may have different instances representing individual trust indicators. We consider the appreciations that a user v receives from the other members of $\mathcal{U}$ (e.g., "likes"), public assessments of reputation which some social networks grant to their best contributors, and the number of friends, fans, or followers in the social networks that disclose the

number but not the identity of users. Similar to the evaluation of the feedback on user contributions, we compute the value of each trust indicator as the ratio between the number of appreciations received by v, denoted as $appreciations_v$, and the maximum number of appreciations received by a user $a \in \mathcal{U}$:

$$fEndors_v = \frac{appreciations_v}{\max_{a \in \mathcal{U}} appreciations_a} \tag{5}$$

In this way we are able to assign a value that indicates the importance of each user profile with respect to the profiles of the other users of the community, on the basis of public or anonymous data.

- *Visibility* (vis_v in [0, 1]). This class is aimed at estimating how popular v becomes, thanks to her/his contributions. Intuitively, the visibility describes the impact of the user's contributions in the social network as observed from the feedback they receive. We compute vis_v as the ratio between the number of appreciations received by v and the total number of contributions provided by her/him, normalized by the maximum number of appreciations acquired by the other members of $\mathcal{U}$:

$$vis_v = \frac{appreciations_v}{\max_{a \in \mathcal{U}} appreciations_a * |Contributions_v|} \tag{6}$$

where $Contributions_v$ is the set of contributions authored by v.

(Q) *Quality of the user as a contributor* (q_v, in [0, 1]) with respect to the other members of $\mathcal{U}$.

This class is aimed at providing an overall evaluation of the user by considering the feedback received by all her/his contributions. As in the previous cases, we compute quality in a relative way with respect to the best contributor of the social network. Specifically:

- In the social networks that only provide positive feedback about contributions we define q_v as follows:

$$q_v = \frac{\sum_{c_1 \in Contributions_v} appreciations_{c_1}}{\max_{a \in \mathcal{U}} \sum_{c_2 \in Contributions_a} appreciations_{c_2}} \tag{7}$$

  where $appreciations_{c_1}$ is the number of appreciations received by contribution c_1 and c_1 is authored by user v (analogously for c_2).

- In the social networks that provide positive and negative feedback about reviews, we apply Equation 7 by replacing $appreciations_{c_n}$ with $fContr_{c_n}$ computed according to Equation 2; i.e., we compute the relative quality of the user's contributions by taking both the positive and negative votes they receive into account.

The previously described classes of trust evidence are generic and most of them could be mapped to multiple indicators. For instance, Yelp supports different types of endorsements to user profiles, such as "thanks" and "Elite" recognition. In other cases, the social relation between users might be mapped to friends, follower and trust links. In order to obtain a single value representing a user v's multi-dimensional reputation, we fuse these indicators by computing their average, assuming that they additively contribute to increasing v's trustworthiness. Let's consider a set $\mathcal{F}$ of indicators that are

instances of the P, U and Q classes. We define the *multi-dimensional global reputation* of v, denoted as mgr_v (in [0, 1]) as:

$$mgr_v = \frac{\sum_{l=1}^{|\mathcal{F}|} C_l * indicator_l}{\sum_{l=1}^{|\mathcal{F}|} C_l} \tag{8}$$

where C_l can be set to 1 to take the trust indicator (e.g., social links) into account, 0 otherwise. We assume that each indicator is computed according to the method defined for the class to which it belongs.

It is worth noting that in Equation 8 all the indicators have the same weight because, for simplicity, we assume that they equally contribute to v's reputation. In our future work, we plan to carry out a deeper analysis to understand the impact of different weighting schemes on recommendation performance. For this purpose, we will carry out experiments with LOCABAL+ by setting these weights to different values in [0, 1].

4.2.1. Multi-faceted trust

The multi-faceted trust, mft_{vi} (in [0, 1]), describes the overall trust in the rating provided by a user v on an item i, given v's multi-dimensional reputation and the quality of her/his contribution about i. We use multi-faceted trust values in LOCABAL+ to tune the influence of rating scores in the Matrix Factorization process used to learn the latent user and item vectors; see Section 5.3. We define mft_{vi} as follows:

$$mft_{vi} = \beta * mgr_v + C(1 - \beta) fContr_{vi} \tag{9}$$

where

- mgr_v is v's multi-dimensional reputation; see Equation 8.
- $fContr_{vi}$ is the quality of the contribution provided by v on item i and is computed according to Equations 2 or 3, depending on the type of feedback (positive/negative) that contributions can receive.
- β takes values in the [0, 1] interval and balances the relative weight of mgr_v and $fContr_{vi}$ in the computation of mft_{vi}. The higher β, the stronger is the impact of multi-dimensional reputation on trust.
- C can be either 0 or 1 and is used to ignore or use the feedback on contributions in the evaluation of mft_{vi}; by default, $C = 1$.

As discussed below in Section 7, the best configuration of the β parameter in Equation 9 depends on the dataset to which the trust model has to be applied and it can be empirically found by using the MTM model within a recommender system (LOCABAL+ in our case) and checking its performance. In the datasets we have considered, the best values are somehow low (e.g., 0.1 or 0.3), which means that the global feedback on user contributions, represented by $fContr_{vi}$, is very useful to steer recommendation.

5. Recommendation model

We describe LOCABAL+ incrementally, starting from the main concepts that characterize the LOCABAL trust-based recommender system.

5.1. Basic Collaborative Filtering with Matrix Factorization

Basic Collaborative Filtering builds on the assumption that, if people rated items similarly in the past, they will do it again in the future. Thus,

it uses rating similarity in preference estimation. The algorithms based on Matrix Factorization assume that a few latent patterns influence rating behavior and they perform a low-rank matrix factorization on the users-items rating matrix; e.g., see SVD++ (Koren, 2008). Given the following notation:

- $\mathcal{U} = \{u_1, \ldots, u_n\}$ is the set of users and $\mathcal{I} = \{i_1, \ldots, i_m\}$ is the set of items.
- $\mathbf{R} \in \mathbb{R}^{n \times m}$ is the users-items rating matrix.
- $\mathbf{R}_{xy}$ is the rating score given by user $u_x \in \mathcal{U}$ to item $i_y \in \mathcal{I}$, if any:
  - $\mathcal{O} = \{< u_x, i_y > \mid \mathbf{R}_{xy} \neq 0\}$ is the set of known ratings (ground truth)
  - $\mathcal{T} = \{< u_x, i_y > \mid \mathbf{R}_{xy} = 0\}$ is the set of unknown ratings

Assuming K latent factors, $\mathbf{u}_x \in \mathbb{R}^K$ denotes the user preference vector of u_x and $\mathbf{i}_y \in \mathbb{R}^K$ denotes the item characteristic vector of i_y.

In order to learn these vectors, the recommender system solves the following optimization problem:

$$\min_{\mathbf{U},\mathbf{I}} \sum_{<u_x, i_y> \in \mathcal{O}} (\mathbf{R}_{xy} - \mathbf{u}_x^T \mathbf{i}_y)^2 + \lambda(||\mathbf{U}||_F^2 + ||\mathbf{I}||_F^2) \tag{10}$$

where

- $\mathbf{U} = [\mathbf{u}_1, \ldots, \mathbf{u}_n] \in \mathbb{R}^{K \times n}$ and $\mathbf{I} = [\mathbf{i}_1, \ldots, \mathbf{i}_m] \in \mathbb{R}^{K \times m}$.
- $||.||_F$ denotes the Frobenius Norm and $||\mathbf{U}||_F^2 + ||\mathbf{I}||_F^2$ are the regularization terms to avoid over-fitting.
- $\lambda > 0$ controls the impact of $\mathbf{U}$ and $\mathbf{I}$ on regularization.

5.2. LOCABAL

LOCABAL (Tang et al., 2013) extends Collaborative Filtering based on Matrix Factorization in two ways:

1. It exploits a user's local social context to learn her/his preference vector by considering both rating similarity and social relations, regularized on the basis of the former. In this way, rating estimation can benefit from the contribution of users who are socially related to the current user but, at the same time, have similar preferences as her/him.
2. It relies on the user's global social context, represented by her/his reputation, to weight the contribution of rating similarity in Matrix Factorization. Global reputation is computed using PageRank as described in Section 4.2, Equation 4, page 17.

In detail:

- Let $\mathbf{T} \in \mathbb{R}^{n \times n}$ be the users-users social relation matrix. $\mathbf{T}_{uz} \neq 0$ denotes the existence of a direct social link between $u_x \in \mathcal{U}$ and $u_z \in \mathcal{U}$. Zero values mean that users are not socially related.

- Let $\mathcal{N}_x = \{u_z \mid \mathbf{T}_{xz} = 1\}$ be the set of u_x's direct social links.

- Let $\mathbf{S} \in \mathbb{R}^{n \times n}$ be a users-users trust matrix whose cells represent the strength of the social relations between users, depending on their rating similarity. For $u_z \in \mathcal{N}_x$, $\mathbf{S}_{xz} = \sigma(u_x, u_z)$, where $\sigma(u_x, u_z)$ is the Cosine similarity of u_x and u_z's rating vectors.

LOCABAL solves the following optimization problem:

$$\min_{\mathbf{U},\mathbf{I},\mathbf{H}} \sum_{<u_x,i_y>\in\mathcal{O}} w_x(\mathbf{R}_{xy} - \mathbf{u}_x^T\mathbf{i}_y)^2 + \alpha \sum_{x=1}^{n} \sum_{u_z\in N_x} (\mathbf{S}_{xz} - \mathbf{u}_x^T\mathbf{H}\mathbf{u}_z)^2 + \lambda(||\mathbf{U}||_F^2 + ||\mathbf{I}||_F^2 + ||\mathbf{H}||_F^2) \tag{11}$$

where

- w_x in [0, 1] is u_x's global reputation computed by applying Equation 4. This weight tunes the contribution given by rating similarity so that highly reputable users influence the Matrix Factorization process more strongly than the other ones.
- $\alpha >= 0$ tunes the contribution given by u_x's local social context.
- $\mathbf{H} \in \mathbb{R}^{K\times K}$ captures user preference correlation: if u_x and u_z are strongly connected in $\mathbf{S}_{xz}$, then their preferences should be tightly correlated via $\mathbf{H}$. We remind that K is the number of latent factors.
- $\lambda >= 0$ controls the impact of $\mathbf{U}$, $\mathbf{I}$ and $\mathbf{H}$ on regularization.

As discussed in Section 3, this algorithm estimates trust by relying on social links. If this information is not available, LOCABAL cannot be applied or it reduces to SVD (Koren et al., 2009), by setting $w_x = 1$ and ignoring the trust matrix that cannot be computed. Our MTM model is aimed at providing a more general solution, which can be applied to complementary types of evidence about trust as can be found in social networks.

5.3. LOCABAL+

LOCABAL+ extends LOCABAL in two ways:

1. It models global context by taking multi-faceted trust into account.
2. It tunes social regularization on the basis of both rating similarity and multi-dimensional global reputation.

We consider the following optimization problem to be solved in order to learn the user preference and item characteristic vectors:

$$\min_{\mathbf{U},\mathbf{I},\mathbf{H}} \sum_{<u_x,i_y>\in\mathcal{O}} mft_{xy}(\mathbf{R}_{xy} - \mathbf{u}_x^T\mathbf{i}_y)^2 + \alpha \sum_{x=1}^{n} \sum_{u_z\in N_x} mgr_z(\mathbf{S}_{xz} - \mathbf{u}_x^T\mathbf{H}\mathbf{u}_z)^2 \\ +\lambda(||\mathbf{U}||_F^2 + ||\mathbf{I}||_F^2 + ||\mathbf{H}||_F^2) \tag{12}$$

where:

- mft_{xy} represents the multi-faceted trust towards user u_x in the context of item i_y; see Equation 9 in page 20. This weight tunes the estimation of ratings in the Matrix Factorization process by taking users' global reputation and quality of contributions into account; i.e., by looking at users from a broad perspective on their behavior. We assume that we can estimate missing ratings more precisely by giving more importance to the ratings authored by users whose multi-faceted trust is high.

- $\mathbf{S} \in \mathbb{R}^{n\times n}$ is a users-users trust matrix such that, for $u_z \in \mathcal{N}_x$, $\mathbf{S}_{xz}$ is set to the Pearson Correlation similarity (PC) of u_x and u_z's rating vectors, limited to the set of items rated by both users:

$$PC(u_x, u_z) = \frac{\sum_{i_y\in\mathcal{I}_{xz}} (r_{xy} - \bar{r}_x)(r_{zy} - \bar{r}_z)}{\sqrt{\sum_{i_y\in\mathcal{I}_{xz}} (r_{xy} - \bar{r}_x)^2 \sum_{i_y\in\mathcal{I}_{xz}} (r_{zy} - \bar{r}_z)^2}} \tag{13}$$

  where $\mathcal{I}_{xz}$ is the set of items rated by both u_x and u_z, r_{xy} is the rating

given by u_x to i_y (and analogously for r_{zy}), $\bar{r}_x$ ($\bar{r}_z$) is the mean value of u_x's (u_z's) ratings.

As suggested in (Ricci et al., 2011), we use Pearson Correlation similarity, instead of the Cosine similarity used in LOCABAL, because the latter does not consider the differences in the mean and variance of the ratings made by u_x and u_z. Pearson similarity removes the effects of mean and variance.

- mgr_z is the multi-dimensional global reputation of u_z and tunes preference correlation in the **H** matrix (which depends on rating similarity, given **S**) on the basis of u_z's multi-dimensional global reputation. By adding the mgr_z factor we impose that, the more reputable are u_x's friends, the higher impact they have in the estimation of their own similarity with u_x. Therefore, highly trustworthy users influence social regularization more than the others.

As mft_{xy} and mgr_z are based on a compositional model, they can be computed by using a subset of the trust facets considered so far; e.g., by ignoring social links, feedback on user profiles or feedback about user contributions. In those cases, LOCABAL+ runs with a lower amount of information about users but it can still work as a trust-based recommender system. Specifically, in the experiments we carried out, the algorithm reaches satisfactory performance results also with partial evidence about trust; see Section 7.

Obviously, the flexibility of MTM comes with a cost, i.e., the effort needed to map the facets of trust to the types of information available in the application domain in which the recommender system is used. This effort consists

of understanding the semantics of the evidence about trust (e.g., types of feedback that are provided by users) and choosing the corresponding classes of trust in MTM. However, as previously discussed, we defined these classes by analyzing several social networks to abstract from the particular types of information they offer and to model trust in a general way. The next section describes the datasets we used for our experiments and the mappings we defined to apply LOCABAL+ to these datasets. Moreover, it sketches the work that should be done to map MTM to a different type of social network in order to give the reader a broader idea of the work to be done.

6. Datasets

For our experiments we use two subsets of the (Yelp, 2019b) dataset to analyze data about user behavior in different domains: accommodations versus restaurants.

The Yelp dataset contains information about the users of the social network and about a large set of businesses including food, accommodation, transportation, health, education and so forth. Yelp members can establish bidirectional friend relations to share posts; moreover, they can establish stricter unidirectional fan relations to get access to the contributions provided by other users. The dataset is structured as follows:

- Each item (business) is associated with a list of tags representing the categories to which it belongs; e.g., a restaurant might be associated with the "Indian" tag to specify the type of cuisine it offers. The full list of Yelp categories is available at
  `https://www.yelp.com/developers/documentation/v3/category_list`.

- Each item is associated with the rating scores and with textual reviews and tips provided by the members of Yelp. Every user can post a contribution (including review+rating, and possibly tip) on the same item. Item ratings take values in a [1,5] Likert scale where 1 is the worst value and 5 is the best one.

- User contributions are associated with the appreciations they receive from Yelp members; i.e., "useful", "funny" and "cool" for reviews, "like" for tips.

- The dataset publishes friend relations but it only provides the number of fans of each user. Therefore, only the former data can be used to infer direct *trust-alike* relations among users; see Section 3.2.

- The dataset publishes various types of endorsement that user profiles can receive: e.g., every year Yelp rewards its most valuable contributors by attributing them the status of Elite users. Moreover, each user profile can receive *compliments* by other Yelp users; e.g., "write more", "thanks" and "great writer".

Notice that both compliments and appreciations represent positive feedback about users and contributions; moreover, the dataset reports the number of compliments and appreciations but not the identities of the people who provided them. This type of feedback thus represents an important anonymous source of trust information that can be used by a recommender system.

Table 2: Statistics about the Yelp-Hotel dataset.

Measure				Value
#Users				654
#Items (businesses)				1081
#Ratings				10081
#Friend relations				11554
Sparsity of users-items rating matrix				0.9857
Sparsity of users-users friends matrix				0.9729
Measure	Min	Max	Mean	Median
#Elite years of individual users	0	13	4.4052	4
#Compliments received by individual users	0	45018	724.1177	110
#Fans of individual users	0	1803	91.9740	41
#Appreciations on reviews provided by ind. users	0	5194	112.7064	53
#Appreciations on tips provided by ind. users	0	152	2.5107	0
#Appreciations received by individual reviews	0	559	3.9897	3
#Friends of individual users	0	224	17.6667	7

6.1. Yelp-Hotel

Yelp-Hotel is obtained by filtering the complete Yelp dataset on users who provided at least 10 ratings and on businesses tagged with at least one category associated with accommodation facilities. The tags used to filter the dataset are: Hotels, Mountain Huts, Residences, Rest Stops, Bed & Breakfast, Hostels, Resorts.

Table 2 provides information about this dataset. It can be noticed that user profiles receive various types of feedback; e.g., the median number of Elite years is 4 and the median number of compliments to user profiles is 110. Also anonymous fans contribute to global reputation (median = 41). Moreover, the dataset contains a relatively high amount of feedback about user contributions: the median number of appreciations is 53 for reviews and 0 for tips. The number of compliments, fans, appreciations, etc. reaches very high values in some cases: for each type of feedback, the distribution

Table 3: Statistics about the Yelp-Food dataset.

Measure				Value
#Users				8432
#Items (businesses)				8157
#Ratings				198759
#Friend relations				160891
Sparsity of users-items rating matrix				0.9971
Sparsity of users-users friends matrix				0.9977
Measure	Min	Max	Mean	Median
#Elite years of individual users	0	13	1.4154	0
#Compliments received by individual users	0	24635	55.0733	4
#Fans of individual users	0	1803	10.1950	2
#Appreciations on reviews provided by ind. users	0	9023	81.4163	27
#Appreciations on tips provided by ind. users	0	154	0.3876	0
#Appreciations received by individual reviews	0	642	3.4539	1
#Friends of individual users	0	1231	19.0810	4

of individuals (users or contributions) has a long tail. Both the users-items and the users-users friends matrices are sparse.

6.2. Yelp-Food

Yelp-Food, about 10 times larger than Yelp-Hotel, is obtained by filtering the complete Yelp dataset on users who provided at least 10 ratings and on businesses located in the cities of Phoenix, Toronto, Pittsburgh which are tagged with at least one category describing a type of restaurant (e.g., "Indian" and "Italian") for a total of 85 categories.[2]

[2]The selection of businesses to define the Yelp-Food is based on the following tags: American, Argentine, Asian Fusion, Australian, Austrian, Bangladeshi, Belgian, Brasseries, Brazilian, British, Cambodian, Cantonese, Catalan, Chinese, Conveyor Belt Sushi, Cuban, Czech, Delis, Empanadas, Falafel, Filipino, Fish & Chips, French, German, Greek, Hawaiian, Himalayan/Nepalese, Hot Pot, Hungarian, Iberian, Indian, Indonesian, Irish, Italian, Japanese, Japanese Curry, Korean, Latin American, Lebanese, Malaysian, Mediterranean, Mexican, Middle Eastern, Modern European, Mongolian, New Mexican Cuisine, Noodles, Pakistani, Pan Asian, Persian/Iranian, Peruvian, Piadina, Pizza, Poke, Polish, Polynesian, Portuguese, Ramen, Russian, Salad, Scandinavian, Scottish, Seafood,

As shown in Table 3, the median number of compliments, fans and appreciations is very low and it reaches higher values only for the number of appreciations on reviews provided by individual users. Moreover, for each type of feedback, the distributions of users and reviews have long tails.

6.3. Trust indicators for both datasets

Let $\mathcal{U}$ and $\mathcal{I}$ be the sets of users and items of the dataset; let $u, v \in \mathcal{U}$ and $i \in \mathcal{I}$. We define the following trust indicators:

- *Quality of individual contributions on an item* ($fContr_{vi}$, in [0, 1]).

  For this class of trust evidence we apply Equation 3, which is suitable for positive-only feedback. We map $appreciations_{contr_{vi}}$ to the possibly different types of feedback that a contribution $contr_{vi}$ can receive:

  - $appreciations_{contr_{vi}} = useful_{contr_{vi}} + fun_{contr_{vi}} + cool_{contr_{vi}}$ for item reviews
  - $appreciations_{contr_{vi}} = like_{contr_{vi}}$ for tips

  where $useful_{contr_{vi}}$ is the number of "useful" appreciations received by $contr_{vi}$ (a review), fun is a shortener for "funny" and so forth.

- *Multi-dimensional global reputation* (mgr_v in [0, 1]).

  Following the approach described in Section 4.2, we compute the multi-dimensional global reputation of a user v by fusing in Equation 8 the

Shanghainese, Sicilian, Singaporean, Soup, Southern, Spanish, Sri Lankan, Steakhouses, Sushi Bars, Syrian, Tacos, Tapas Bars, Tapas/Small Plates, Teppanyaki, Tex-Mex, Thai, Turkish, Ukrainian, Vegan, Vegetarian, Vietnamese, Wraps.

indicators described in the remainder of this section:

$$mgr_v = \frac{C_1 imp_v + C_2 elite_v + C_3 lup_v + C_4 opLeader_v + C_5 vis_v + C_6 q_v}{C_1 + C_2 + C_3 + C_4 + C_5 + C_6} \tag{14}$$

(P) *Importance of the user in the social network* (imp_v in [0, 1]).

In order to compute the PageRank score of users, we transform each bidirectional friend relation into two unidirectional social links. In this way, we can apply the approach described in Section 4.2 and Equation 4 to compute reputation on the basis of the connections among the users of the social network.

(U) *Global feedback on the user's profile.*

We consider the following trust indicators:

- $elite_v$ (in [0, 1]). We map the number of years in which v has the Elite status to $appreciations_v$ in Equation 5:

$$elite_v = \frac{\#EliteYears_v}{\max_{a \in \mathcal{U}} \#EliteYears_a} \tag{15}$$

- lup_v (degree of liking of user profile, in [0, 1]). We map the number of compliments ("more", "thanks" - *thks*, "great writer" - *gw*) received by v to $appreciations_v$ in Equation 5:

$$lup_v = \frac{more_v + thks_v + gw_v}{\max_{a \in \mathcal{U}}(more_a + thks_a + gw_a)} \tag{16}$$

where $more_v$ is the number of "more" compliments received by v, and similar for the other variables.

– $opLeader_v$ (opinion leader degree, in [0, 1]). The number of anonymous fans of a user v, $fans_v$, can be interpreted as a global recognition of her/his profile. We thus map this number to $appreciations_v$ in Equation 5:

$$opLeader_v = \frac{fans_v}{\max_{a \in \mathcal{U}} fans_a} \tag{17}$$

– vis_v (visibility, in [0, 1]). We map the number of compliments received by v to $appreciations_v$, and the reviews and tips ($Revs_v \cup Tips_v$) authored by v to $Contributions_v$ in Eq. 6:

$$vis_v = \frac{more_v + thks_v + gw_v}{\max_{a \in \mathcal{U}}(more_a + thks_a + gw_a)|Revs_v \cup Tips_v|} \tag{18}$$

(Q) *Quality of the user as a contributor* (q_v in [0, 1]).

We assume that the quality of a contributor depends on both the reviews and tips authored by her/him. Therefore, for this indicator, we map $Contributions_v$ to the reviews and the tips provided by v. Moreover, we map $appreciations_c$ to the amount of feedback obtained by these contributions:

$$q_v = \frac{\sum_{c_1 \in Revs_v \cup Tips_v} useful_{c_1} + fun_{c_1} + cool_{c_1} + like_{c_1}}{\max_{a \in \mathcal{U}} \sum_{c_2 \in Revs_a \cup Tips_a} useful_{c_2} + fun_{c_2} + cool_{c_2} + like_{c_2}} \tag{19}$$

6.4. Instantiation of MTM in a different application domain

Le'ts consider, as a further example of instantiation of MTM, the LibraryThing (2019) social network that publishes information about books.

LibraryThings enables its members to create their own virtual libraries and to tag and review books. Users can establish friend relations to watch and take inspiration from the libraries created by other people; moreover, they can visualize the reviews published in the social network and they can express positive-only feedback about the helpfulness of each review. Users are not enabled to endorse other users' profiles. LibraryThing discloses the social relations among users and the number of helpfulness votes received by each review. Trust indicators can be mapped to MTM trust classes as follows:

- *Quality of individual contributions on an item* ($fContr_{vi}$). For this class of trust evidence we apply Equation 3 of page 16, which is suitable for positive-only feedback. For each review r published in the social network, we thus map $appreciations_r$ to the number of helpful votes that r has received.

- *Multi-dimensional global reputation* (mgr_v). We compute mgr_v by fusing in Equation 8 the imp_v and q_v indicators respectively describing v's importance in the social network and her/his quality as a contributor:

$$mgr_v = \frac{C_1 imp_v + C_6 q_v}{C_1 + C_6} \tag{20}$$

  We can compute imp_v as v's PageRank score by transforming bidirectional friend relations to pairs of unidirectional social links. Moreover, q_v can be defined as the ratio between the total number of helpful votes received by v's reviews and the maximum number of helpful votes received by the other members of the social network.

7. Validation of LOCABAL+

7.1. Evaluation metrics

As mentioned in Section 2, we evaluate recommendation algorithms on the basis of *accuracy* and *error minimization* (i.e., the ability to provide correct results), *ranking capability* (i.e., the ability to correctly sort items depending on their ground truth relevance to the user) and *user coverage* (i.e., the percentage of users for whom the recommender is able to find items that are likely to be relevant). This is in line with the recent trends in the evaluation of recommender systems, which do not exclusively focus on accuracy to provide a broader view on performance; e.g., see (Jannach et al., 2016). Before describing the evaluation metrics in detail we introduce the notation we use:

- U is the set of users and I the set of items; R is the set of ground truth ratings and $\hat{R}$ the set of estimated ones.
- r_{ui} is the rating score that $u \in U$ has given to $i \in I$ and $\hat{r}_{ui}$ is the rating score estimated by the recommender system.
- $Relevant_u$ is the set of items that u has positively rated; in a [1, 5] Likert scale we define $Relevant_u = \{i \in I \mid r_{ui} > 3\}$.
- $Recommended_u$ is the set of items that the system suggests to u: $Recommended_u = \{i \in I \mid \hat{r}_{ui} > 3\}$.

We evaluate recommendation accuracy and error minimization by means of the following metrics:

- Precision: $\text{P@k} = \frac{1}{|U|} \sum_{u \in U} \text{P}_u\text{@k}$, where $\text{P}_u\text{@k} = \frac{|Recommended_u \cap Relevant_u|}{|Recommended_u|}$
- Recall: $\text{R@k} = \frac{1}{|U|} \sum_{u \in U} \text{R}_u\text{@k}$, where $\text{R}_u\text{@k} = \frac{|Recommended_u \cap Relevant_u|}{|Relevant_u|}$
- Accuracy: $\text{F1@k} = 2 * \frac{\text{P@k * R@k}}{\text{P@k + R@k}}$
- Root Mean Squared Error: $\text{RMSE@k} = \sqrt{\frac{1}{|\hat{R}\text{@k}|} \sum_{\hat{r}_{ui} \in \hat{R}\text{@k}} (r_{ui} - \hat{r}_{ui})^2}$
- Mean Absolute Error: $\text{MAE@k} = \frac{1}{|\hat{R}\text{@k}|} \sum_{\hat{r}_{ui} \in \hat{R}\text{@k}} |r_{ui} - \hat{r}_{ui}|$

As far as ranking capability is concerned we use the following metrics:

- Mean Reciprocal Rank, which measures the placement of the first relevant items in recommendation lists:

  $\text{MRR@k} = \frac{1}{|U|} \sum_{u \in U} \frac{1}{rank_u}$, where $rank_u$ is the position of the first relevant item in the list generated for user u.

- Mean Average Precision, which measures the average correct positioning of items in the recommendation lists:

  $\text{MAP@k} = \frac{1}{|U|} \sum_{u \in U} \frac{1}{|Relevant_u|} \sum_{x=1}^{\text{k}} \text{P}_u\text{@x} * Rel_u(x)$

  where $Rel_u(x) = 1$ if the item in position x of the list for u is relevant to her/him, 0 otherwise.

Finally, we measure User Coverage (shortened to UCov in the tables showing the evaluation results) as the percentage of users of the dataset for whom the algorithm finds at least one item $i \in I$ such that $\hat{r}_{ui} > 3$, i.e., an item that the system evaluates as relevant to the user.

Table 4: Configurations of LOCABAL+ used in the experiments.

Configuration	Social relations	Feedback on user profiles	Feedback on user contributions
LOCABAL+	Yes	Yes	Yes
LOC+noF	Yes	Yes	-
LOC+noE	Yes	-	Yes
LOC+noS	-	Yes	Yes

7.2. Methodology applied in the experiments

We consider various configurations of MTM to evaluate the performance of LOCABAL+ when using all the facets of trust available in the YELP-Hotel/Yelp-Food datasets, or a subset of them. We are interested in understanding whether the algorithm can provide good recommendation results when we omit different sources of evidence in order to assess its applicability to social networks that disclose different types of information about users. Specifically, we consider the following cases, summarized in Table 4:

- LOCABAL+. This is the algorithm applied to the complete information available in the dataset (social relations, feedback about users and feedback about contributions). It computes multi-dimensional global reputation with $C_1 = \cdots = C_6 = 1$ in Equation 14, and multi-faceted trust for item i with $C = 1$ in Equation 9.

- LOC+noF. This configuration ignores the feedback about reviews and tips; thus, it only relies on multi-dimensional global reputation, which is computed by taking social links and global feedback on user profiles into account. In detail, LOC+noF is obtained by switching off the quality of the user as a contributor (q_v) in the computation of multi-dimensional global reputation ($C_6 = 0$ in Equation 14) and the feedback

received by the specific contribution ($fContr_{vi}$) in the computation of multi-faceted trust for item i ($C = 0$ in Equation 9). LOC+noF is useful to understand whether, by only using social information and anonymous feedback on user profiles, the recommender system is able to generate useful suggestions.

- LOC+noE. This configuration ignores the global feedback on user profiles, i.e., the trust indicators of class U in Section 4.2 (user profile endorsements and public recognition, visibility) in the computation of multi-dimensional global reputation (i.e., $C_2 = C_3 = C_4 = C_5 = 0$ in Equation 14). LOC+noE is particularly interesting because not all of the social networks manage profile endorsements; e.g., we mentioned in Section 6 that LibraryThing does not support this type of feedback. Therefore we are interested in understanding whether the recommender system can achieve good performance by only relying on social links and feedback on user contributions.

- LOC+noS. This configuration ignores social relations; it is obtained by switching off the importance of users (imp_v) in the computation of multi-dimensional global reputation ($C_1 = 0$ in Equation 14) and the social regularization component of LOCABAL+ ($\alpha = 0$ in Equation 12). LOC+noS helps understand whether, thanks to the exploitation of public, anonymous feedback about users and user contributions, LOCABAL+ can generate good recommendations in the application domains where the information about social links is unavailable. As previously discussed, this is an important aspect for the applicability

of trust-based recommender systems, given the growing sensibility of users towards privacy protection.

In the experiments we use the Surprise (Hug, 2018) implementation of U2UCF and SVD++ and the RecQ implementation of LOCABAL (Coder-Yu, 2019). LOCABAL+ and U2USocial are developed by extending the implementations of LOCABAL and U2UCF respectively. All the algorithms are integrated in Surprise to uniformly evaluate their performance.

On each dataset we organize the evaluation as follows: we first validate the algorithms on 90% of the dataset by running Grid Search to find the best configuration of parameters with respect to MAP, using 5 cross-fold validation. All the executions are performed having set 50 latent factors. Then we additionally test the best configuration obtained from Grid Search on the remaining 10% of the dataset to measure the performance of the algorithms on new data in order to check their impact in a dynamic environment where new ratings are continuously provided.

7.3. Evaluation results

7.3.1. Yelp-Hotel

Table 5 compares the performance achieved by each configuration of LOCABAL+ to that of the baselines (U2UCF, SocialMF, LOCABAL, SVD++ and U2USocial) on Yelp-Hotel. In this table, as well as in the following ones, the best values are in boldface. The significance level of results, reported in the second row, is obtained by separately comparing each configuration with all the baselines. The rows describing performance are sorted by MAP, from the best one to the worst one, in order to highlight the ranking capabilities of the algorithms.

Table 5: Performance@10 on Yelp-Hotel dataset (the best results are in boldface). The "$\diamond$" symbol means that results are significant at $p < 0.05$ with respect to all the baselines except for SocialMF; the "†" means that results are significant at $p < 0.01$ with respect to all the baselines except for U2UCF.

	α	β	P	R	F1	MAP	RMSE	MAE	MRR	UCov
Significance	-	-	0.01	-	0.01	†	0.02	0.02	-	$\diamond$
LOCABAL+	0.9	0	0.7919	0.7389	**0.7645**	**0.5303**	0.8922	0.671	**0.6125**	0.7543
LOC+noF	0.1	0.3	0.7923	0.7381	0.7642	0.5288	0.8927	0.6708	0.6087	0.7523
LOC+noE	0.3	0.3	0.7923	0.7377	0.764	0.5285	0.8938	0.6716	0.6086	0.7517
LOC+noS	-	0.1	**0.7931**	0.7368	0.7639	0.5274	**0.8916**	**0.6702**	0.6082	0.7513
U2UCF	-	-	0.76	**0.7399**	0.7498	0.5215	0.9582	0.7264	0.5982	0.6127
SocialMF	-	-	0.7757	0.7261	0.7501	0.5116	0.9238	0.6954	0.6055	0.7655
LOCABAL	0.1	-	0.7732	0.7259	0.7488	0.5112	0.9281	0.6994	0.6078	0.7698
SVD++	-	-	0.7595	0.717	0.7376	0.4994	0.976	0.7383	0.5993	**0.7755**
U2USocial	-	-	0.7503	0.7233	0.7366	0.4798	1.0085	0.773	0.5336	0.2589

As shown in the central portion of the table, in this dataset the LOCABAL+ configurations outperform the baselines in all measures except for (i) Recall that is dominated by U2UCF, and (ii) User Coverage, where SVD++ is the best algorithm, followed by LOCABAL and SocialMF. The loss in user coverage is however compensated by higher accuracy and ranking capability as LOCABAL+ is the best algorithm in terms of F1, MAP and MRR. Noticeably, LOC+noS is the most precise algorithm, excelling in P, RMSE and MAE. In the following we analyze the LOCABAL+ configurations.

LOCABAL+ obtains its best performance with $\alpha = 0.9$ and $\beta = 0$. The value of α shows that the algorithm strongly relies on users' multi-dimensional reputation to steer social regularization: instead of minimizing the impact of the local social context of users ($\mathbf{S}_{xz}$ in Equation 12), as done in LOCABAL (where $\alpha = 0.1$ uniformly flattens the impact of the local social context across users), LOCABAL+ tunes social regularization on the basis of the preferences of the most similar and reputable friends. Differently, $\beta = 0$

means that the multi-faceted trust mtf_{xy} that tunes the impact of ratings in Matrix Factorization is computed by ignoring users' reputation; therefore, for this purpose, the algorithm only relies on the feedback ($fContr_{vi}$) received by the reviews and tips associated with the ratings. This is different from LOCABAL, which tunes the impact of ratings on the basis of users' PageRank score. In summary, in Yelp-Hotel, LOCABAL+ steers social regularization by multi-dimensional reputation and weights the impact of ratings on the basis of the publicly recognized value of user contributions, which emerges as a good source of information to identify reliable ratings.

It should however be noticed that, in LOCABAL+, multi-dimensional global reputation is computed by taking multiple types of trust evidence into account, i.e., PageRank score, user profile endorsements and quality of the user as a contributor that, in turn, derives from the feedback on contributions. Therefore, it is difficult to say which type of evidence brings the most useful information. In order to clarify the situation we analyze the other configurations of LOCABAL+.

LOC+noF ignores the feedback on user contributions and is optimized with $\alpha = 0.1$ and $\beta = 0.3$. It has lower performance than LOCABAL+ but it outperforms all the baselines in Precision, F1, MAP, RMSE, MAE and MRR. The value of α dramatically weakens the role of social regularization in the Matrix Factorization process with respect to LOCABAL+ (it is flattened to 10% as in LOCABAL, but it is much weaker due to the presence of the mgr_z term within the nested summation of Equation 12). Moreover, $\beta = 0.3$ means that the values of multi-faceted trust computed by the algorithm are reduced to 30% of the multi-dimensional reputation. However, given the

weak role of social regularization, reputation is central to learning the user preference and item characteristic vectors. These findings are coherent with the hypothesis that, in this dataset, the feedback on user contributions is a very useful type of information to learn user preferences but show that, even by only employing social links and the feedback on user profiles, the algorithm can achieve satisfactory results.

LOC+noE ignores the feedback on user profiles. It obtains its best performance with $\alpha = \beta = 0.3$: with respect to LOCABAL+, the algorithm weights social regularization much less but it partially takes multi-dimensional reputation into account (30%) in the computation of multi-faceted trust. With respect to LOC+noF, LOC+noE increases a little bit the role of social regularization in the Matrix Factorization process. The algorithm outperforms the baselines and is generally worse than LOCABAL+. Moreover, it performs slighly worse than LOC+noF: it has the same precision and very similar F1, MAP and MRR but it has lower recall, RMSE, MAE and User Coverage. We explain these findings with the fact that, as LOC+noE ignores the trust feedback received by user profiles, it generally misses useful information for preference prediction. In the dataset, user endorsements have high median values (e.g., the median number of compliments received by individual users is 110 and the median number of fans is 41). Therefore, when the algorithm ignores them, it has fewer chances to recognize highly reputable users. The value of β also shows that the feedback on user contributions determines the value of multi-faceted trust by 70%. Once more, it looks like the feedback on user contributions has an important role in defining trust; however, the feedback on user profiles is useful as well.

LOC+noS ignores social relations among users: it only employs anonymous trust statements and anonymous social information (number of fans) to learn user preferences. This means that it is not possible to compute the importance of users in the social network (imp_v) and that social regularization does not make sense. We obtain LOC+noS by forcing parameter $C_1 = 0$ in Equation 14 and $\alpha = 0$ in Equation 12. This algorithm outperforms the baselines in all measures except for User Coverage and Recall; moreover, it has the best RMSE and MAE of all the algorithms and configurations of LOCABAL+. This supports the hypothesis that, in this dataset, anonymous trust feedback is a precious source of information to be used in a recommender system, and that correct ratings can be predicted without using personal data about social relations.

Figure 1 shows the variation of MAP for all the configurations of LOCABAL+, depending on α and β. By setting α to a constant value (Figure 1b), MAP decreases when β grows. This means that, having blocked the influence of the social component of LOCABAL+, the best results are achieved when the feedback on contributions makes ratings more influent on the Matrix Factorization process. Moreover, by setting β to a constant value (Figure 1a), MAP improves when α increases and the best results are achieved with α=0.9, i.e., when the social component, in combination with the other facets of trust, strongly influences preference estimation.

Table 6 summarizes the evaluation results on new data (10% of Yelp-Hotel). It can be seen that the results are fairly consistent with those of Table 5: LOCABAL+ outperforms the baselines in all performance measures except for User Coverage, as previously, dominated by SVD++ and

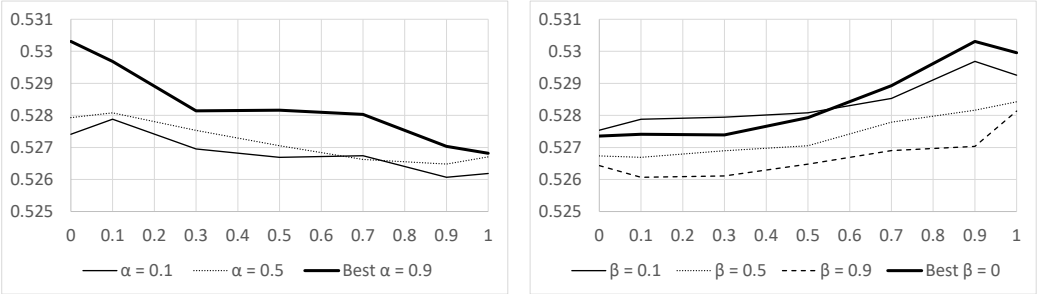

(a) MAP variation with respect to β **(b)** MAP variation with respect to α

Figure 1: MAP variation on Yelp-Hotel. The Y axis represents MAP; the X axis represents β in Figure 1a and α in the Figure 1b.

Table 6: Performance@10 on the new data of the Yelp-Hotel dataset.

	α	β	P	R	F1	MAP	RMSE	MAE	MRR	UCov
LOCABAL+	0.9	0	**0.8253**	**0.7703**	**0.7968**	**0.4818**	**0.8801**	**0.6601**	0.5341	0.6748
SocialMF	-	-	0.8044	0.7687	0.7862	0.4808	0.918	0.701	**0.5399**	0.7055
LOCABAL	0.1	-	0.8041	0.761	0.782	0.4736	0.9369	0.7158	0.5378	0.7014
U2UCF	-	-	0.7906	0.7684	0.7794	0.4723	0.9627	0.7251	0.5239	0.5767
SVD++	-	-	0.7867	0.7607	0.7735	0.4709	0.9832	0.7457	0.5324	**0.7076**
U2USocial	-	-	0.7759	0.7561	0.7658	0.4518	1.0597	0.7975	0.4952	0.2802

SocialMF. The main difference in this case is that SocialMF also dominates MRR. We can conclude that LOCABAL+ can be used in dynamic application domains without frequently optimizing the model.

7.3.2. Yelp-Food

Table 7 shows the evaluation results on the Yelp-Food dataset. The LOCABAL+ configurations outperform the baselines with statistically significant results in all measures except for User Coverage. Moreover, SVD++ has the highest coverage and LOCABAL is the second best, followed by SocialMF, but all of them are less accurate and have lower ranking capability than the LOCABAL+ configurations. In this case, LOCABAL+ and

Table 7: Performance@10 on Yelp-Food dataset.

	α	β	P	R	F1	MAP	RMSE	MAE	MRR	UCov
Significance			0.01	0.01	0.01	0.01	0.01	0.01	0.04	0.01
LOCABAL+	0.7	0.3	0.7769	**0.7528**	0.7647	**0.5993**	0.9791	0.7406	**0.7166**	0.8499
LOC+noE	0.9	0.3	0.7769	0.7527	0.7646	0.5993	0.9793	0.7407	0.7165	0.8497
LOC+noS	-	0.3	0.777	0.7527	0.7647	0.5992	0.9789	0.7403	0.7164	0.8497
LOC+noF	0.5	0.7	**0.7783**	0.752	**0.7649**	0.5986	**0.9768**	**0.7389**	0.7163	0.8485
SocialMF	-	-	0.76	0.7312	0.7453	0.5741	1.0358	0.7828	0.7092	0.8642
LOCABAL	0.1	-	0.7586	0.7294	0.7437	0.5724	1.041	0.7866	0.7093	0.8656
U2UCF	-	-	0.7488	0.7337	0.7412	0.5642	1.0796	0.8072	0.6795	0.7025
SVD++	-	-	0.7483	0.7116	0.7295	0.5539	1.0761	0.812	0.7017	**0.8698**
U2USocial	-	-	0.7729	0.7357	0.7539	0.5474	1.0741	0.8178	0.6061	0.2295

LOC+noF are the best performing algorithms but, as LOCABAL+ has the best Recall, MAP and MRR (and LOC+noF the worst ones among LOCABAL+ configurations), we consider LOCABAL+ as the preferable one.

In this dataset, all the LOCABAL+ configurations take multi-dimensional global reputation into account in the computation of multi-faceted trust. Specifically, $\beta = 0.3$ in LOCABAL+, LOC+noE and LOC+noS; moreover, $\beta = 0.7$ in LOC+noF that overlooks the feedback on user contributions. Moreover, social regularization has a medium to high role in the Matrix Factorization process, with the strongest influence in LOCABAL+ ($\alpha = 0.7$) and LOC+noE ($\alpha = 0.9$). In summary, the configurations exploit both multi-faceted trust and social regularization to obtain their best performance but the exclusion of feedback on user contributions raises the importance of the other facets of trust. This can be explained by the fact that users and user contributions receive a low amount of feedback for the evaluation of trustworthiness; e.g., the median number of appreciations is 1 for reviews and 0 for tips and the median number of endorsements to user profiles, includ-

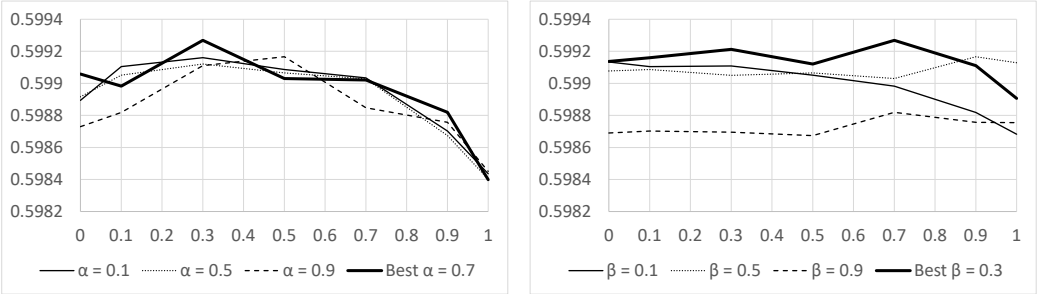

(a) MAP variation with respect to β.

(b) MAP variation with respect to α.

Figure 2: MAP variation on Yelp-Food. The Y axis represents MAP; the X axis represents β in Figure 2a and α in Figure 2b.

ing Elite years, compliments and fans, is 6. Thus, the algorithms rely on the joint contribution of all the sources of evidence about trust. Another relevant observation is that the performance of LOC+noE and LOC+noS is similar to that of LOCABAL+. This can be explained by assuming that, in this dataset, the feedback on user profiles and social relations play complementary roles and can replace each other without a major loss of performance.

Figure 2 shows the variation of MAP for all the LOCABAL+ configurations depending on α and β. It can be noticed that by setting α to a constant value (Figure 2b), MAP first slightly increases but, when $\beta > 0.3$, it quickly decreases. This means that, regardless of the influence of the social component of LOCABAL+, the algorithm benefits from a moderate support by multi-dimensional global reputation. Notice also that, by setting β to a constant value (Figure 2a), results are not particularly affected by the value of α; therefore, social regularization has a generally constant contribution to recommendation performance.

Table 8 shows the evaluation of algorithms on new data. It can be seen that LOCABAL+ outperforms the baselines in all performance metrics, in-

Table 8: Performance@10 on the new data of the Yelp-Food dataset.

	α	β	P	R	F1	MAP	RMSE	MAE	MRR	UCov
LOCABAL+	0.7	0.3	**0.7935**	**0.787**	**0.7903**	**0.5835**	**0.9701**	**0.7376**	**0.6584**	**0.8053**
U2UCF	-	-	0.769	0.7699	0.7694	0.5606	1.0423	0.7807	0.6378	0.7193
U2USocial	-	-	0.7798	0.7708	0.7753	0.5472	1.0568	0.8055	0.5913	0.2595
SocialMF	-	-	0.7818	0.7489	0.765	0.5461	1.0403	0.7829	0.6387	0.8011
LOCABAL	0.1	-	0.7814	0.7478	0.7642	0.5451	1.0449	0.7864	0.6384	0.8025
SVD++	-	-	0.7742	0.7389	0.7561	0.5339	1.0712	0.8056	0.6314	0.8033

cluding MRR. Thus, we can conclude that also in Yelp-Food LOCABAL+ can be employed in dynamic environments without requiring a frequent optimization of parameters.

8. Discussion and future work

The experimental results show that, in the Yelp-Hotel and Yelp-Food datasets, LOCABAL+ outperforms all the baseline recommender systems in accuracy, error minimization and ranking capability with a minor loss of User Coverage with respect to SVD++, LOCABAL and SocialMF. Thus, we can say that LOCABAL+ generates better suggestions to marginally fewer people, with a clear positive gain in performance.

Noticeably, LOCABAL+ outperforms the baselines even though MTM is configured to ignore different types of evidence about trust; this finding shows the flexibility of our approach towards the lack of user information. This is generally important because some classes of trust information defined in MTM might not be available in specific recommendation domains. Moreover, it is a key achievement in relation to the management of personal data because, different from the other trust-based recommender systems, LOCABAL+ can work by ignoring data about social links, and by only relying

on public anonymous information. This aspect is more and more important given the increasing sensibility of users towards disclosing personal data.

The evaluation results on new data confirm that, on both datasets, LOCABAL+ achieves the best performance in all the measures except for User Coverage (actually, in Yelp-Food LOCABAL+ outperforms the other algorithms in User Coverage as well, but this is not true in Yelp-Hotel). Thus, we conclude that LOCABAL+ can be applied, without losing recommendation capability, to dynamic environments in which it can not be frequently optimized.

The evaluation results are useful to answer our research questions:

RQ1: *Can multi-faceted trust be used to improve the performance of a trust-based recommender system with respect to only relying on social links and rating similarity among users?* Our experiments enable us to positively answer RQ1 because they provide consistent results on two datasets having different characteristics (e.g., size, distributions of global feedback on users and contributions, etc.); see Section 6. Specifically, the superior results of LOCABAL+ with respect to the baselines and, in particular, to LOCABAL, show that by exploiting multiple facets of trust (including global anonymous feedback on users and user contributions), the performance of the recommender system significantly improves. It is worth noticing that, thanks to MTM, trust-based methods can be successfully applied without using personal information about social links. This fact represents a key aspect of our approach compared to existing work on trust-based recommender systems.

RQ2: *What is the impact of the multi-dimensional reputation of users, of the*

quality of their contributions, and of social links, on collaborative recommendation performance? Our experiments show that LOCABAL+ can work without using data about social relations, with a minor loss of performance, if it can use other types of trust feedback. Regarding the other facets of trust, the experiments show that their relative importance depends on the amount and quality of the feedback about users and user contributions available to the recommender system. The two datasets we selected are interesting because the diverse distributions of feedback they provide determine slightly different behavior of the LOCABAL+ configurations. In Yelp-Hotel, which provides a large amount of feedback about user contributions (median = 53 against 27 of Yelp-Food), the algorithm obtains the best performance results by privileging this type of information over multi-dimensional global reputation; moreover, performance results are clearly affected by the omission of the feedback on user contributions. Differently, in the Yelp-Food dataset, which stores scarser feedback about contributions, the algorithm obtains the best results by balancing the influence of all types of global feedback, including the endorsements to user profiles, and performance increases if feedback on user contributions is ignored.

We thus conclude that global feedback about user trust is a useful type of information in Top-N recommendation. However, in order to decide which types of evidence should be used, it is important to analyze the characteristics of the data to which the recommender system is applied and the features of the various types of feedback that can be used.

Before closing this section it is worth noting that our current work focuses

on numerical aspects of the evidence about trust, by measuring the amount of feedback that users and user contributions receive. However, the content of the global feedback provided by users is itself a further source of information about trust that can be analyzed to acquire information about users' behavior. In this perspective, we plan to extend our model to the analysis of the content of reviews (and/or microblogs), which has been largely studied to evaluate their quality (Huang et al., 2015; Mudambi and Schuff, 2010; Chua and Banerjee, 2015; Qazi et al., 2016; Korfiatis et al., 2012; Krishnamoorthy, 2015; Kim et al., 2006), to steer personalized recommendation (Raghavan et al., 2012; Alahmadi and Zeng, 2015; Shen et al., 2019; Hernández-Rubio et al., 2019) and to guide the explanation of recommendations (O'Mahony and Smyth, 2018; Musto et al., 2019).

9. Conclusions

This paper has described the Multi-faceted Trust Model (MTM) and the LOCABAL+ recommender system, which combine social links with global anonymous feedback about users and user contributions to enhance collaborative Top-N recommendation. LOCABAL+ extends the LOCABAL recommender system with multi-faceted trust and with an enhanced regularization of social relations based on both rating similarity and users' multi-dimensional global reputation.

LOCABAL+ has various advantages with respect to the state-of-the art trust-based recommender systems. In particular, being based on the MTM compositional model, it can be configured to work with different types of evidence about trust, such as anonymous public feedback on user profiles and

user contributions, as well as information about social relations. Interestingly, LOCABAL+ can work by ignoring the information about social links, which has been recently found to be problematic in the user acceptance of trust-based recommender systems because it is considered as personal information. Another advantage of LOCABAL+ is that the extension to social regularization makes it possible to select users to steer Matrix Factorization in a more selective way with respect to only considering rating similarity.

Experiments carried out on two public datasets of item reviews show that, with a minor loss of user coverage, LOCABAL+ outperforms state-of-the art trust-based recommender systems and Collaborative Filtering in accuracy, error minimization and ranking capability both when it uses complete information from the datasets and when it ignores social relations (or other types of feedback on users and contributions). It thus represents a flexible approach to trust-based recommendation, suitable to comply with specific data management requirements.

10. Acknowledgments

This work was supported by the University of Torino which funded the conduct of the research and preparation of the article.

References

Abdul-Rahman, A., Hailes, S., 2000. Supporting trust in virtual communities, in: Proceedings of the 33rd Hawaii International Conference on System Sciences-Volume 6 - Volume 6, IEEE Computer Society, Washing-

ton, DC, USA. pp. 6007–. URL: http://dl.acm.org/citation.cfm?id=820262.820322.

Alahmadi, D.H., Zeng, X.J., 2015. ISTS: Implicit social trust and sentiment based approach to recommender systems. Expert Systems with Applications 42, 8840 – 8849. URL: http://www.sciencedirect.com/science/article/pii/S0957417415004960, doi:https://doi.org/10.1016/j.eswa.2015.07.036.

Amazon.com, 2006. Amazon.com: online shopping for electronics, apparel, etc., http://www.amazon.com.

Ardissono, L., Ferrero, M., Petrone, G., Segnan, M., 2017. Enhancing collaborative filtering with friendship information, in: Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization, ACM, New York, NY, USA. pp. 353–354. URL: http://doi.acm.org/10.1145/3079628.3079629, doi:10.1145/3079628.3079629.

Booking.com, 2019. Booking.com, https://www.booking.com.

Burbach, L., Nakayama, J., Plettenberg, N., Ziefle, M., Valdez, A.C., 2018. User preferences in recommendation algorithms: The influence of user diversity, trust, and product category on privacy perceptions in recommender algorithms, in: Proceedings of the 12th ACM Conference on Recommender Systems, ACM, New York, NY, USA. pp. 306–310. URL: http://doi.acm.org/10.1145/3240323.3240393, doi:10.1145/3240323.3240393.

Chaney, A.J., Blei, D.M., Eliassi-Rad, T., 2015. A probabilistic model for using social networks in personalized item recommendation, in: Proceedings

of the 9th ACM Conference on Recommender Systems, ACM, New York, NY, USA. pp. 43–50. URL: http://doi.acm.org/10.1145/2792838.2800193, doi:10.1145/2792838.2800193.

Chua, A., Banerjee, S., 2015. Understanding review helpfulness as a function of reviewer reputation, review rating, and review depth. Journal of the Association for Information Science and Technology 66, 354–362. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/asi.23180, doi:10.1002/asi.23180.

Coder-Yu, 2019. Locabal.py.

Danetsoft, 2019. dvd.ciao.co.uk, http://data.danetsoft.com/dvd.ciao.co.uk.

De Meo, P., Fotia, L., Messina, F., Rosaci, D., Sarn, G.M., 2018. Providing recommendations in social networks by integrating local and global reputation. Information Systems 78, 58 – 67. URL: http://www.sciencedirect.com/science/article/pii/S0306437916306548, doi:https://doi.org/10.1016/j.is.2018.07.002.

Deng, S., Huang, L., Xu, G., 2014. Social network-based service recommendation with trust enhancement. Expert Systems with Applications 41, 8075 – 8084. URL: http://www.sciencedirect.com/science/article/pii/S0957417414004102, doi:https://doi.org/10.1016/j.eswa.2014.07.012.

Desrosiers, C., Karypis, G., 2011. A comprehensive survey of neighborhood-based recommendation methods, in: Ricci, F., Rokach, L., Shapira, B.,

Kantor, P.B. (Eds.), Recommender Systems Handbook. Springer US, Boston, MA, pp. 107–144. URL: https://doi.org/10.1007/978-0-387-85820-3_4, doi:10.1007/978-0-387-85820-3_4.

Du, X., Liu, H., Jing, L., 2017. Additive co-clustering with social influence for recommendation, in: Proceedings of the Eleventh ACM Conference on Recommender Systems, ACM, New York, NY, USA. pp. 193–200. URL: http://doi.acm.org/10.1145/3109859.3109883, doi:10.1145/3109859.3109883.

Epinions, 2019. Epinions datasets, http://www.trustlet.org/epinions.html.

Expedia.com, 2001. Expedia, http://www.expedia.com.

Gambetta, D., 1988. Can we trust trust?, in: Gambetta, D. (Ed.), Trust: Making and Breaking Cooperative Relations. Blackwell, pp. 213–237.

Golbeck, J., Hendler, J., 2004. Accuracy of metrics for inferring trust and reputation in semantic web-based social networks, in: Motta, E., Shadbolt, N.R., Stutt, A., Gibbins, N. (Eds.), Engineering Knowledge in the Age of the Semantic Web, Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 116–131.

Golbeck, J., Hendler, J., 2006. Filmtrust: movie recommendations using trust in web-based social networks, in: CCNC 2006. 2006 3rd IEEE Consumer Communications and Networking Conference, 2006., pp. 282–286. doi:10.1109/CCNC.2006.1593032.

Groh, G., Ehmig, C., 2007. Recommendations in taste related domains: Collaborative filtering vs. social filtering, in: Proceedings of the 2007 International ACM Conference on Supporting Group Work, ACM, New York, NY, USA. pp. 127–136. URL: http://doi.acm.org/10.1145/1316624.1316643, doi:10.1145/1316624.1316643.

Guo, G., Zhang, J., Yorke-Smith, N., 2015. TrustSVD: Collaborative filtering with both the explicit and implicit influence of user trust and of item ratings, in: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI Press. pp. 123–129. URL: http://dl.acm.org/citation.cfm?id=2887007.2887025.

Hernández-Rubio, M., Cantador, I., Bellogín, A., 2019. A comparative analysis of recommender systems based on item aspect opinions extracted from user reviews. User Modeling and User-Adapted Interaction 29, 381–441. URL: https://doi.org/10.1007/s11257-018-9214-9, doi:10.1007/s11257-018-9214-9.

Hu, G.N., Dai, X.Y., Qiu, F.Y., Xia, R., Li, T., Huang, S.J., Chen, J.J., 2018. Collaborative filtering with topic and social latent factors incorporating implicit feedback. ACM Trans. Knowl. Discov. Data 12, 23:1–23:30. URL: http://doi.acm.org/10.1145/3127873, doi:10.1145/3127873.

Huang, A.H., Chen, K., Yen, D.C., Tran, T.P., 2015. A study of factors that contribute to online review helpfulness. Comput. Hum. Behav. 48, 17–27. URL: http://dx.doi.org/10.1016/j.chb.2015.01.010, doi:10.1016/j.chb.2015.01.010.

Hug, N., 2018. Surprise, a Python library for recommender systems.

Jamali, M., Ester, M., 2009. TrustWalker: A random walk model for combining trust-based and item-based recommendation, in: Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, New York, NY, USA. pp. 397–406. URL: http://doi.acm.org/10.1145/1557019.1557067, doi:10.1145/1557019.1557067.

Jamali, M., Ester, M., 2010. A matrix factorization technique with trust propagation for recommendation in social networks, in: Proceedings of the Fourth ACM Conference on Recommender Systems, ACM, New York, NY, USA. pp. 135–142. URL: http://doi.acm.org/10.1145/1864708.1864736, doi:10.1145/1864708.1864736.

Jannach, D., Resnick, P., Tuzhilin, A., Zanker, M., 2016. Recommender systems — beyond matrix completion. Commun. ACM 59, 94–102. URL: http://doi.acm.org/10.1145/2891406, doi:10.1145/2891406.

Jiang, M., Cui, P., Liu, R., Yang, Q., Wang, F., Zhu, W., Yang, S., 2012. Social contextual recommendation, in: Proceedings of the 21st ACM International Conference on Information and Knowledge Management, ACM, New York, NY, USA. pp. 45–54. URL: http://doi.acm.org/10.1145/2396761.2396771, doi:10.1145/2396761.2396771.

Jøsang, A., 2012. Robustness of trust and reputation systems: Does it matter?, in: Dimitrakos, T., Moona, R., Patel, D., McKnight, D.H. (Eds.), Trust Management VI, Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 253–262.

Jøsang, A., Goldbeck, J., 2009. Challenges for robust trust and reputation systems, in: 5th Int. workshop on Security and trust management (STM 2009), Saint Malo, France.

Jøsang, A., Ismail, R., Boyd, C., 2007. A survey of trust and reputation systems for online service provision. Decis. Support Syst. 43, 618–644. URL: http://dx.doi.org/10.1016/j.dss.2005.05.019, doi:10.1016/j.dss.2005.05.019.

Kim, S.M., Pantel, P., Chklovski, T., Pennacchiotti, M., 2006. Automatically assessing review helpfulness, in: Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Stroudsburg, PA, USA. pp. 423–430. URL: http://dl.acm.org/citation.cfm?id=1610075.1610135.

Koren, Y., 2008. Factorization meets the neighborhood: A multifaceted collaborative filtering model, in: Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, New York, NY, USA. pp. 426–434. URL: http://doi.acm.org/10.1145/1401890.1401944, doi:10.1145/1401890.1401944.

Koren, Y., Bell, R., Volinsky, C., 2009. Matrix factorization techniques for recommender systems. Computer 42, 30–37. URL: http://dx.doi.org/10.1109/MC.2009.263, doi:10.1109/MC.2009.263.

Korfiatis, N., Garcia-Baríocanal, E., Sánchez-Alonso, S., 2012. Evaluating content quality and helpfulness of online product reviews: The interplay of review helpfulness vs. review content. Electronic Commerce Research

and Applications 11, 205 – 217. URL: http://www.sciencedirect.com/science/article/pii/S1567422311000639, doi:https://doi.org/10.1016/j.elerap.2011.10.003.

Krishnamoorthy, S., 2015. Linguistic features for review helpfulness prediction. Expert Systems with Applications 42, 3751 – 3759. URL: http://www.sciencedirect.com/science/article/pii/S0957417414008239, doi:https://doi.org/10.1016/j.eswa.2014.12.044.

Kuter, U., Golbeck, J., 2007. Sunny: A new algorithm for trust inference in social networks using probabilistic confidence models, in: Proceedings of the 22Nd National Conference on Artificial Intelligence - Volume 2, AAAI Press. pp. 1377–1382. URL: http://dl.acm.org/citation.cfm?id=1619797.1619866.

Li, H., Bhowmick, S.S., Sun, A., 2011. Casino: Towards conformity-aware social influence analysis in online social networks, in: Proceedings of the 20th ACM International Conference on Information and Knowledge Management, ACM, New York, NY, USA. pp. 1007–1012. URL: http://doi.acm.org/10.1145/2063576.2063721, doi:10.1145/2063576.2063721.

Li, H., Wu, D., Mamoulis, N., 2014. A revisit to social network-based recommender systems, in: Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval, ACM, New York, NY, USA. pp. 1239–1242. URL: http://doi.acm.org/10.1145/2600428.2609554, doi:10.1145/2600428.2609554.

Li, X., Fang, H., Yang, Q., Zhang, J., 2018. Who is your best friend?: Ranking social network friends according to trust relationship, in: Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization, ACM, New York, NY, USA. pp. 301–309. URL: http://doi.acm.org/10.1145/3209219.3209243, doi:10.1145/3209219.3209243.

Li, Y.M., Wu, C.T., Lai, C.Y., 2013. A social recommender mechanism for e-commerce: Combining similarity, trust, and relationship. Decision Support Systems 55, 740 – 752. URL: http://www.sciencedirect.com/science/article/pii/S0167923613000705, doi:https://doi.org/10.1016/j.dss.2013.02.009.

LibraryThing, 2019. LibraryThing - a home for your books, https://www.librarything.com.

Liu, F., Lee, H.J., 2010. Use of social network information to enhance collaborative filtering performance. Expert Syst. Appl. 37, 4772–4778. URL: http://dx.doi.org/10.1016/j.eswa.2009.12.061, doi:10.1016/j.eswa.2009.12.061.

Liu, X., Aberer, K., 2013. SoCo: A social network aided context-aware recommender system, in: Proceedings of the 22Nd International Conference on World Wide Web, ACM, New York, NY, USA. pp. 781–802. URL: http://doi.acm.org/10.1145/2488388.2488457, doi:10.1145/2488388.2488457.

Ma, H., King, I., Lyu, M.R., 2011a. Learning to recommend with explicit and implicit social relations. ACM Trans. Intell. Syst. Technol. 2, 29:1–

29:19. URL: http://doi.acm.org/10.1145/1961189.1961201, doi:10.1145/1961189.1961201.

Ma, H., Zhou, D., Liu, C., Lyu, M.R., King, I., 2011b. Recommender systems with social regularization, in: Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, ACM, New York, NY, USA. pp. 287–296. URL: http://doi.acm.org/10.1145/1935826.1935877, doi:10.1145/1935826.1935877.

Ma, H., Zhou, T.C., Lyu, M.R., King, I., 2011c. Improving recommender systems by incorporating social contextual information. ACM Trans. Inf. Syst. 29, 9:1–9:23. URL: http://doi.acm.org/10.1145/1961209.1961212, doi:10.1145/1961209.1961212.

Marsden, P., Friedkin, N., 1993. Network studies of social influence. Sociological Methods & Research 22, 127–151. URL: https://doi.org/10.1177/0049124193022001006, doi:10.1177/0049124193022001006, arXiv:https://doi.org/10.1177/0049124193022001006.

Massa, P., Avesani, P., 2007. Trust-aware recommender systems, in: Proceedings of the 2007 ACM Conference on Recommender Systems, ACM, New York, NY, USA. pp. 17–24. URL: http://doi.acm.org/10.1145/1297231.1297235, doi:10.1145/1297231.1297235.

Mauro, N., Ardissono, L., Hu, Z.F., 2019. Multi-faceted trust-based collaborative filtering, in: Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization, ACM, New York, NY, USA.

pp. 216–224. URL: http://doi.acm.org/10.1145/3320435.3320441, doi:10.1145/3320435.3320441.

Mcnally, K., O'Mahony, M.P., Smyth, B., 2014. A comparative study of collaboration-based reputation models for social recommender systems. User Modeling and User-Adapted Interaction 24, 219–260. URL: http://dx.doi.org/10.1007/s11257-013-9143-6, doi:10.1007/s11257-013-9143-6.

McPherson, M., Smith-Lovin, L., Cook, J.M., 2001. Birds of a feather: Homophily in social networks. Annual Review of Sociology 27, 415–444. URL: https://doi.org/10.1146/annurev.soc.27.1.415, doi:10.1146/annurev.soc.27.1.415.

Misztal, B., 1996. Trust in modern societies. Polity Press, Cambridge, MA.

Moradi, P., Ahmadian, S., 2015. A reliability-based recommendation method to improve trust-aware recommender systems. Expert Systems with Applications 42, 7386 – 7398. URL: http://www.sciencedirect.com/science/article/pii/S0957417415003553, doi:https://doi.org/10.1016/j.eswa.2015.05.027.

Mudambi, S., Schuff, D., 2010. What makes a helpful online review? a study of customer reviews on Amazon.com. MIS Quarterly 34, 185–200.

Mui, L., Mohtashemi, M., Halberstadt, A., 2002. A computational model of trust and reputation, in: Proceedings of the 35th Annual Hawaii International Conference on System Sciences, pp. 2431–2439. doi:10.1109/HICSS.2002.994181.

Musto, C., Lops, P., de Gemmis, M., Semeraro, G., 2019. Justifying recommendations through aspect-based sentiment analysis of users reviews, in: Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization, ACM, New York, NY, USA. pp. 4–12. URL: http://doi.acm.org/10.1145/3320435.3320457, doi:10.1145/3320435.3320457.

O'Donovan, J., Smyth, B., 2005. Trust in recommender systems, in: Proceedings of the 10th International Conference on Intelligent User Interfaces, ACM, New York, NY, USA. pp. 167–174. URL: http://doi.acm.org/10.1145/1040830.1040870, doi:10.1145/1040830.1040870.

O'Mahony, M.P., Smyth, B., 2018. From opinions to recommendations, in: Brusilovsky, P., He, D. (Eds.), Social Information Access: Systems and Technologies. Springer International Publishing, Cham, pp. 480–509. URL: https://doi.org/10.1007/978-3-319-90092-6_13, doi:10.1007/978-3-319-90092-6_13.

Page, L., Brin, S., Motwani, R., Winograd, T., 1999. The PageRank Citation Ranking: Bringing Order to the Web. Technical Report 1999-66. Stanford InfoLab. URL: http://ilpubs.stanford.edu:8090/422/. previous number = SIDL-WP-1999-0120.

Parvin, H., Moradi, P., Esmaeili, S., 2019. TCFACO: Trust-aware collaborative filtering method based on ant colony optimization. Expert Systems with Applications 118, 152 – 168. URL: http://www.sciencedirect.com/science/article/pii/S0957417418306225, doi:https://doi.org/10.1016/j.eswa.2018.09.045.

Qazi, A., Syed, K.B.S., Raj, R.G., Cambria, E., Tahir, M., Alghazzawi, D., 2016. A concept-level approach to the analysis of online review helpfulness. Computers in Human Behavior 58, 75 – 81. URL: http://www.sciencedirect.com/science/article/pii/S0747563215302995, doi:https://doi.org/10.1016/j.chb.2015.12.028.

Qian, F., Zhao, S., Tang, J., Zhang, Y., 2016. Sors: Social recommendation using global rating reputation and local rating similarity. Physica A: Statistical Mechanics and its Applications 461, 61 – 72. URL: http://www.sciencedirect.com/science/article/pii/S0378437116302102, doi:https://doi.org/10.1016/j.physa.2016.05.025.

Rafailidis, D., Crestani, F., 2017. Learning to rank with trust and distrust in recommender systems, in: Proceedings of the Eleventh ACM Conference on Recommender Systems, ACM, New York, NY, USA. pp. 5–13. URL: http://doi.acm.org/10.1145/3109859.3109879, doi:10.1145/3109859.3109879.

Raghavan, S., Gunasekar, S., Ghosh, J., 2012. Review quality aware collaborative filtering, in: Proceedings of the Sixth ACM Conference on Recommender Systems, ACM, New York, NY, USA. pp. 123–130. URL: http://doi.acm.org/10.1145/2365952.2365978, doi:10.1145/2365952.2365978.

Resnick, P., Zeckhauser, R., 2002. Trust among strangers in Internet transactions: empirical analysis of eBay's reputation system. Emerald Group

Publishing Limited. volume 11. pp. 127–157. doi:https://doi.org/10.1016/S0278-0984(02)11030-3.

Ricci, F., Rokach, L., Shapira, B., 2011. Introduction to Recommender Systems Handbook. Springer US, Boston, MA. pp. 1–35. URL: https://doi.org/10.1007/978-0-387-85820-3_1, doi:10.1007/978-0-387-85820-3_1.

Richthammer, C., Weber, M., Pernul, G., 2017. Reputation-enhanced recommender systems, in: Steghöfer, J.P., Esfandiari, B. (Eds.), Trust Management XI, Springer International Publishing, Cham. pp. 163–179.

Shen, R.P., Zhang, H.R., Yu, H., Min, F., 2019. Sentiment based matrix factorization with reliability for recommendation. Expert Systems with Applications URL: http://www.sciencedirect.com/science/article/pii/S0957417419303951, doi:https://doi.org/10.1016/j.eswa.2019.06.001.

Su, K., Xiao, B., Liu, B., Zhang, H., Zhang, Z., 2017. TAP: A personalized trust-aware QoS prediction approach for web service recommendation. Knowledge-Based Systems 115, 55 – 65. URL: http://www.sciencedirect.com/science/article/pii/S0950705116303331, doi:https://doi.org/10.1016/j.knosys.2016.09.033.

Tang, J., Hu, X., Gao, H., Liu, H., 2013. Exploiting local and global social context for recommendation, in: Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, AAAI Press. pp. 2712–2718. URL: http://dl.acm.org/citation.cfm?id=2540128.2540519.

Victor, P., Cornelis, C., Cock, M.D., Teredesai, A.M., 2011. Trust- and distrust-based recommendations for controversial reviews. IEEE Intelligent Systems 26, 48–55. doi:10.1109/MIS.2011.22.

Yang, B., Lei, Y., Liu, J., Li, W., 2017. Social collaborative filtering by trust. IEEE Transactions on Pattern Analysis and Machine Intelligence 39, 1633–1647. doi:10.1109/TPAMI.2016.2605085.

Yang, X., Steck, H., Liu, Y., 2012. Circle-based recommendation in online social networks, in: Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, New York, NY, USA. pp. 1267–1275. URL: http://doi.acm.org/10.1145/2339530.2339728, doi:10.1145/2339530.2339728.

Yelp, 2019a. Yelp, https://www.yelp.com.

Yelp, 2019b. Yelp dataset challenge, https://www.yelp.com/dataset_challenge.

Yuan, Q., Chen, L., Zhao, S., 2011. Factorization vs. regularization: Fusing heterogeneous social relationships in top-n recommendation, in: Proceedings of the Fifth ACM Conference on Recommender Systems, ACM, New York, NY, USA. pp. 245–252. URL: http://doi.acm.org/10.1145/2043932.2043975, doi:10.1145/2043932.2043975.

11. Author contributions

Liliana Ardissono has the following contributions in the present paper: Conceptualization; Funding acquisition; Methodology; Resources; Supervision; Roles/Writing - original draft.

Noemi Mauro has the following contributions in the present paper: Conceptualization; Data curation; Investigation; Methodology; Software; Validation; Roles/Writing - original draft.