

Language Independent Recommender Agent

Osman Yucel

The University of Tulsa 800 S Tucker Dr Tulsa, Oklahoma, USA
E-mail: osman-yucel@utulsa.edu

Sandip Sen

The University of Tulsa 800 S Tucker Dr Tulsa, Oklahoma, USA
E-mail: sandip-sen@utulsa.edu

Abstract

This paper presents a new “Language Independent Recommender Agent” (LIRA), using information distributed over any text-source pair on the web about candidate items. While existing review-based recommendation systems learn the features of candidate items and users’ preferences, they do not handle varying perspectives of users on those features. LIRA constructs agents for each user, which run regression algorithms on texts from different sources and builds trust relations. The key advantages of LIRA can be listed as : LIRA does not require reviews from target users, LIRA calculates trust values based on prediction accuracy instead of social connections or rating similarity, LIRA does not require the reviews to come from the same community or peer user group. Since ratings of the reviewers are not necessary for LIRA, we can collect and use reviews from different sources (web pages, professional critiques), as long as we know the corresponding item and source of that text. Since LIRA does not combine text from different sources, texts from different sources are not required to be in the same language. LIRA can utilize text from multiple languages, as long as sources are consistent with their language usage.

1 Introduction

The demand for accurate recommendation systems has dramatically increased with the growth of user engagement in varying e-Commerce portals and is facilitated by the rapid proliferation of personal and usage data on the Internet. An average user does not have enough time or resources to systematically explore the content on e-Commerce sites to identify items of interest. Hence there is a pressing need for automated and personalized assistance for exploring content at online stores and service outlets.

The recommendation systems must be able to learn about both the preferences of the users as well as the properties of the contents to provide accurate personalized recommendations to the users. Both of these requirements present complex challenges.

The first challenge is collecting information about content properties. The information about the items used by current recommendation systems is generally provided by the creator of the content or by a few expert individuals. This source bias makes limits the usefulness of the available information for the purpose of recommendation. This is because effective recommendation systems need to leverage subjective information from the perspective of individual users of the system. To illustrate this problem, consider the following recommendation request by a user on the *GoodReads* website:

“Id Like to read a page-turner book with a strong female lead ! i like YA , with humore and some supernatural aspects to it.” (*sic*)

The second critical challenge is learning the preferences of the users: humans are not perfectly rational beings and will not necessarily be aware of or be able to effectively articulate their own preferences.

There are a number of issues that prevent the recommendation systems from accurately modeling the users. Studies have, for example, exposed paradoxical transitive preferences from users (such as $A \succ B \succ C \succ A$) when asked to compare a set of items (Davis, 1958). In real life when people try to make decisions about candidate items we might be interested in, they usually base their decision about the candidates on what other people say about those items. They do not only consider what has been said about the items, but the source of that information is also important. For example, a movie critique(C) saying 'good' for a movie might not mean any thing for target user(T) while a blogger(B) using the same word for a movie might be very important for decision making process of T. Another example about real life is, in some cases it is not important if the source of information has liked the item or not. For example if T likes every movie which were defined as a 'World War II(WW2)' movie by B, T will probably like the next movie which was defined as 'WW2' by B, regardless of B's liking of that movie.

Existing systems interpret the reviews of users by relating them to their ratings. This introduces two problems: (i) since the rating system can differ among different web pages (different ranges, thumbs up/down, etc.), using reviews from different web pages is challenging and (ii) those systems can miss the fact that a feature of an item which causes one user to dislike an item, might cause another user to like the item. To illustrate this problem, we analyze a review by $User_A$ accompanied with a low rating:

"Romantic chick-flick; total waste of time!"

While it is obvious that $User_A$ did not like the item, the same review can suggest that $User_B$, who likes romantic movies, might like this movie. Another problem with existing review based systems is that they extract the features of the items and subsequently discard the information about the reviewers. For example, suppose $User_A$ reviewed a hotel $Hotel_X$ as:

"The view from the hotel was excellent!"

Let another user, $User_B$, review a hotel $Hotel_Y$ using the exact same words. Existing review based recommendation systems generally interpret these reviews as " $User_A$ and $User_B$ both care about the views from the hotels" and " $Hotel_X$ and $Hotel_Y$ both has nice views." Therefore these systems will recommend $Hotel_X$ to $User_B$ and $Hotel_Y$ to $User_A$. But if $Hotel_X$ has a nice city view, which $User_A$ likes, and $Hotel_Y$ has a nice nature view, which $User_B$ likes, these recommendations are misleading. LIRA avoids this problem by learning that $User_B$'s reviews are not predictive of $User_A$'s ratings.

Another problem that existing review based recommendation systems face is that not every user is willing to write reviews for the items they evaluate. Rating/Review ratios in *IMDB* and *GoodReads* are only 1/400 and 1/20 respectively. Because of this sparsity, *systems which require reviews from the target users will fail to produce recommendations for a large majority of the users*. LIRA does not require target users to write reviews and can recommend items to users as long as they rate items.

Finally, most of the existing review based systems depend on topic modelling approaches or finding similarities between different users' reviews. This limits their applicability to situations where all the reviews are in the same language. They also fail to work when the writing style of the users change, such as using "gr8" instead of "great". LIRA can use reviews from different languages for recommendation, and even different writing styles, as long as users are consistent about their language usage or writing style.

In this paper we present an agent based recommender system, LIRA, which creates agents for target users (A_T) and reviewers to create accurate recommendations. LIRA consists of targets user's agent, which processes messages from agents of other users, which sends the words sequences (N-Grams) in those users' reviews to target user's agent. LIRA agents, A_T s, learn from those messages and then make predictions for unknown items based on the learnt knowledge.

We summarize the problems confounding other recommender systems but which LIRA can handle:

1. LIRA can make use of different perspectives of the users on different features of items, such as liking the different types of views in different hotels.
2. LIRA does not depend on other users' or sources' ratings, it can work on the ratings of a target user, given that reviews are available for items rated by the target user. This property also enables LIRA to utilize reviews from different sources, such as blogs, reviews from different websites, etc.

3. LIRA does not depend on social connections, explicit statement, or rating similarity for calculating trust for a review source. It can calculate these trust values by measuring the capability of those sources to predict target user's rating in training set.
4. LIRA only needs Ratings of target users and does not require them to write reviews to get recommendations.
5. LIRA is language independent: it can leverage reviews written in any language.

2 Background and Related Work

In this section we provide the literature review on the methods we use, and also about the state of recommendation systems research.

2.1 TF-IDF

TF-IDF (Term Frequency - Inverse Document Frequency) is a measure of how important a word is to a document (Leskovec et al., 2014). It is widely used in the text mining approaches. TF-IDF value is calculated by multiplying two measures: Term Frequency(TF) and Inverse Document Frequency(IDF).

TF value is calculated using the following equation:

$$TF_{ij} = \frac{f_{ij}}{\max_k f_{kj}} \quad (1)$$

In the equation above, TF_{ij} is the Term Frequency value of word i in document j , f_{ij} is the frequency of i in j (number of occurrences), $\max_k f_{kj}$ is the frequency of the most frequent word k in document j . Thus, the most frequent term in document j gets a TF of 1, and other terms get fractions as their term frequency for this document.

IDF value is calculated using the following equation:

$$IDF_i = \log_2 \frac{N}{n_i} \quad (2)$$

In the equation above IDF_i is the Inverse Document Frequency of word i , n_i is the number of documents that word i appears in, and N is the total number of documents that we have in our collection.

Then TF-IDF value becomes:

$$TF - IDF_{ij} = TF_{ij} \times IDF_i \quad (3)$$

The terms with the highest TF-IDF score are often the terms that best characterize the topic of the document. We use TF-IDF because it can eliminate too frequent words, as we don't want the system to be affected by words which appears too frequently.

2.2 Regression Algorithms

In this section we are providing a list of regression algorithms which can be used with LIRA.

2.2.1 Gaussian Process

A Gaussian Process is a collection of random variables, any finite number of which have (consistent) joint Gaussian distributions Williams and Rasmussen (2006). Every Gaussian Process can be fully specified with its mean function $m(x)$ and co-variance function $k(x, x')$. Given N data points, X_n, t_n where the inputs x are vectors of some input dimension, and targets t are the output values, the Gaussian process tries to infer the underlying function, $f(x)$ from the given data MacKay (1998). When that function is inferred a Gaussian process can use that function to predict t_{n+1} for a new data point x_{n+1} .

2.2.2 Support Vector Regression

The support vector machine (SVM), originally proposed by Vapnik (2013), is a learning algorithm (Cristianini and Shawe-Taylor, 2000; Palaniswami and Shilton, 2002) with the ability to provide function estimation. By using a mapping, $\Phi : X \rightarrow F$, where X is the domain and F is usually a high-dimensional feature space, support vector regression (SVR) operates in feature space to approximate unknown functions in an output space Y , thereby using nonlinear functions to linearly estimate an unknown regression.

2.2.3 Neural Networks

Artificial neural network is a widely used regression method (Wang, 2003). It aims to simulate working method of human brain to learn, and make predictions about data. It is implemented by creating rather simple nodes, called perceptrons, and creates fully connected layers to form a network. Perceptrons are linear separators, which gets the weighted sum of inputs and checks if that some is over or under the threshold value for that perceptron.

Actual learning in neural networks is done by learning the correct weight values between perceptrons. At every iteration, weights between every pair of nodes are updated using

$$W_{i,j} \leftarrow W_{i,j} + \alpha \times a_i \times \Delta_j, \quad (4)$$

where Δ_j is the error associated with node j , α is the learning rate, and a_i is the output of node i . When a neural network is learning, the only correct value which is known is the output value, therefore the error of the system can only be measured for the output layer. So the error associated with output node, o , is

$$\Delta_o = Err_o \times g(input_o). \quad (5)$$

The error made by output nodes are back-propagated to the hidden nodes to find their contribution on the error. So error for a hidden layer node h becomes:

$$\Delta_h = g(input_h) \times \sum_i W_{h,i} \Delta_i \quad (6)$$

2.3 Recommendation Systems

Recommendation techniques can be broadly divided into two categories: *non-personalized* and *personalized* recommendation techniques. *Non-personalized Recommendation* techniques compute the average statistics of the recommended items and recommend the same item set to all customers, i.e., they do not consider individual user preferences. For instance, a bookstore might recommend a bestseller book to all of its customers. *Non-personalized* systems have the advantage of being fast. Moreover, they do not suffer from the *cold start* problem¹, since they do not need an initial data (ratings, etc.) from each customer. Nevertheless, the quality of the results is low because of the lack of personalization.

Personalized Recommendation techniques provide better results by considering individual user preferences. Two well-known approaches are *content-based* and *collaborative filtering*. Moreover, some hybrid approaches have emerged to overcome the problems of these approaches. Adomavicius and Tuzhilin (2005) provide a detailed survey about recommender systems.

Content-based filtering algorithms recommend items that are similar to the ones that the user has liked. *Content-based* systems usually depend on textual content, e.g., keywords. Information Retrieval (IR) techniques, such as *TF-IDF* (*term frequency / inverse document frequency*) (Manning et al., 2008), can be applied on the keywords. The similarity among items can be determined by some scoring heuristics, such as *cosine similarity*, after the *TF-IDF* vectors are calculated. Apart from *IR* techniques, *Bayesian classifiers* and different machine learning techniques like *clustering decision trees* and *artificial neural networks* can be used for the similarity calculation (Adomavicius and Tuzhilin, 2005). Content-based

¹The problem of not being able to provide recommendations because of the lack of preference information about the user.

recommendation systems have been used by the researchers on various topics (Lops et al., 2011), such as music (Soleymani et al., 2015), books (Mooney and Roy, 2000), movies (Debnath et al., 2008), etc. However these systems need the items to be marked with features, which is usually objective, and cannot make use of user similarities.

Collaborative filtering algorithms identify people with overlapping interests. These algorithms rely on the assumption that people who have exhibited similar interests in the past will continue to have similar interests in the future. Instead of considering correlation of items as in *content-based* algorithms, correlation between users' preferences is examined. This is achieved by analyzing user ratings of items: if two users have provided similar ratings for many items, then it is concluded that these users have similar preferences. Another way of finding users with similar interests is to match demographic characteristics of users, e.g., age, education, geographic locations, gender, etc. After determining like-minded users, the transitive relationship between users and items is utilized. For instance, if user c_1 's preferences are similar to that of user c_2 's and user c_2 likes item s_1 , s_1 is recommended to c_1 if c_1 has not viewed item s_1 yet.

As in the case of *content-based* approaches, *IR* techniques, like *cosine measure*, can be used to calculate similarities between users. However, in *collaborative filtering*, similarity between vectors of the actual user-specified ratings is measured (Adomavicius and Tuzhilin, 2005). Moreover, different machine learning techniques are also applicable for identifying similarities between users (Adomavicius and Tuzhilin, 2005; Alpaydin, 2014; Mitchell, 1997).

RegSVD, proposed by Paterek (2007), improves on SVD by learning factor vectors directly on known ratings through a suitable objective function which minimizes prediction error. The proposed objective function is regularized in order to avoid over-fitting. Gradient descent is applied to minimize the objective function.

Collaborative filtering techniques have their own limitations. They work reliably only when there are sufficient number of users in the system with overlapping characteristics. Another concern is that when a new item is added to the system, it cannot be recommended to others until a number of people have rated it. Furthermore, these techniques are computationally expensive, and the recommendation process becomes cumbersome for millions of users and items. Collaborative filtering has also been tried on various domains such as e-commerce (Linden et al., 2003), movies (Ungar and Foster, 1998), etc. Another shortcoming is that these systems cannot identify when two users rate an item with the same rating with different reasons.

2.3.1 Recommendation with Topic Models

A topic model is a statistical model that generates documents (strings of words) from some set of topic clusters. Parametric estimation techniques and unsupervised clustering algorithms such as the Expectation-Maximization algorithm (Dempster et al., 1977) can fit models to observed documents. With the model, it is possible to estimate a likely distribution over topics on a document level and a distribution over words on the topic level. Topic models are an effective tool for characterizing documents since they are not necessary dependent on a language's underlying grammar, and yet can identify significant topic trends using the rate of each word's appearance. Some of the more popular existing topic model frameworks are Latent Semantic Analysis (LSA) (Deerwester et al., 1990) and Latent Dirichlet Allocation (LDA) (Blei et al., 2003).

Topic models are often used in recommender systems in a variety of creative ways (Bao et al., 2014; Chen and Wang, 2013; Hariri et al., 2011; Jakob et al., 2009; Liu et al., 2013). Generally, such recommender systems use topic models to generate latent features in their documents which can then be supplied to some supervised learning algorithm (Chen and Wang, 2013; Jakob et al., 2009; Paterek, 2007).

The fLDA topic model from Agarwal and Chen (2010) extended the LDA model to include ratings as well. The paper approaches the problem where a system has some users and items, and the goal of the recommender system is to make recommendations based on the user's rating history and a textual description of each item. The fLDA model simultaneously learns: (a) the correlation between items' documents and features using LDA, and (b) the users' affinity towards each topic, to produce (c) a rating distribution for each user i and item j . The rating distribution inside the model is defined as a Normal (or binomial) distribution: $y_{ij} \sim \text{Normal}(\alpha_i + \beta_j + s_i^T \bar{z}_j, \sigma^2)$ where y_{ij} is the rating, α_i is the user's rating

bias, β_j is the item's rating bias, s_i is the user's affinity towards each topic as a vector, and \bar{z} is a latent variable measuring how each word in the description of item j correlates with a topic.

Similarly, the TopicMF model by Bao et al. (2014) incorporates a topic model based on matrix factorization; this model considers both ratings and reviews of items in its recommendation. Their model combines this topic model using matrix factorization with another recommender system using matrix factorization to produce a combined model. Given I users, J items, $D = I \times J$ reviews available, W total words and an observed user-to-item rating matrix \mathbf{r} and a review-to-word frequency matrix \mathbf{F} , the algorithm attempts to find the best factorization for K hidden topics: $\mathbf{r} = \mathbf{u}_{I \times K} \mathbf{v}_{J \times K}^T$, $\mathbf{F} = \mathbf{\Theta}_{D \times K} \mathbf{\Phi}_{W \times K}^T$. The hidden matrices are learned by minimizing the least squares equation: $\min_{\mathbf{u}, \mathbf{v}, \mathbf{\Theta}, \mathbf{\Phi}} \|\mathbf{r} - \mathbf{u}\mathbf{v}^T\| + \|\mathbf{F} - \mathbf{\Theta}\mathbf{\Phi}^T\| + \gamma$ where γ is an additional term that includes bias values and measures to prevent learned parameters from growing too large.

McAuley and Leskovec (2013) have proposed the HFT algorithm, which creates topics to identify the hidden topics in the reviews to describe the items of interest, and also the users' interest in those items. They are also using their algorithm to identify the categories of those items automatically.

These systems work on reviews provided by the users, but in the process of combining them into topics they lose the information about the reviewers. For example if two people both mention that they like the 'view' of two different hotels, those items and users will be considered similar by these systems. However the types of the views can be different and the users opinion of a 'good view' might also differ between users.

2.3.2 Hybrid Recommendation Systems

Sparsity of available data makes the personalization of rating prediction difficult, as it makes it hard to link users and products to each other. To overcome this problem some researchers focused on generating hybrid recommendation systems to combine the strengths of different approaches (Kazienko and Musiał, 2006). Mostly the base of this hybrids are combination of content-based (Kazienko and Musiał, 2006; Schein et al., 2002; Levi et al., 2012) and collaborative recommendation systems (Sarwar et al., 2001). Some researchers tried to incorporate the information from social networks to this hybrid recommendation systems either as social tags (Middleton et al., 2004), or social trust (Sugiyama et al., 2004). Some of the more recent studies focused on creating hybrid recommendation systems which also make use of sentiment analysis approaches (Homoceanu et al., 2011; Yang et al., 2013). Leung et al. (2008) proposed a method to extract association rules from data and use it as a supplement to enhance the recommendation system.

2.3.3 Recommendation Using Trust

The area of trust-based recommender systems has been the object of extensive study for the past years. Indeed, trust has been shown to provide significant improvements to classical collaborative filtering techniques. The main difference between most of these trust-enhanced methods is the acquisition of trust values between actor pairs.

Golbeck (2005) introduced *TidalTrust* algorithm, which uses a modified breadth-first search, to estimate the trust by using transitive rules. In this study the trust of user u for user v is calculated using u 's trusted users' trust on v . Massa and Avesani (2007); Massa and Bhattacharjee (2004) proposed a similar method where instead of only looking at one level trust transitivity, they remove cycles from the trust network. Using a *trust propagation horizon* phase, they allow the trust to diffuse along the network.

While the studies above require an explicit statement about direct trust between users, it is very difficult to get that information from user explicitly. Therefore some other studies focused on defining their own implicit trust values based on implicit information. O'Donovan and Smyth (2005) proposed a trust value based on the similarity of two user's ratings, while Wang et al. (2011) generated a trust metric based on the similarity among users' tastes.

3 Methodology

In this section we define how LIRA works. We first describe the preprocessing steps, i.e., how the data is prepared, and then explain how LIRA uses that data. While the dataset includes detailed meta-data about users, items, and reviews, our recommender system will only examine the textual content and numeric rating associated with the records. We do not need the rating to accompany the review so we create two sets of triples: $(Reviewer, Item, Review)$ and $(User, Item, Rating)$. This property shows the reviews and ratings are only related through the items. Note that sets of $Users$ and $Reviewers$ are not necessarily disjoint. If a person rates items but does not write reviews, they will only appear in the set of $Users$. Likewise, if a person writes reviews but do not provide ratings, such as bloggers or critiques, they only appear in the set of $Reviewers$. However if a person both writes reviews and provides ratings, they appear both in $Reviewers$ and $Users$ sets.

3.1 Preparing Data

For this study we have used “Amazon Reviews” dataset in 5 categories McAuley et al. (2015a,b). In addition to these, we have also used “GoodReads” dataset which consists of book ratings, and reviews from many different languages.

In all datasets we used, we run the following pre-processing steps:

1. Remove products which do not have any reviews, along with ratings for those products.
2. Remove users who have less than 10 ratings, along with ratings by those users.
3. Remove reviews for products which do not have any ratings.

Information about resulting dataset is shown in Table 1.

Table 1 Properties of Datasets

	Raters	Reviewers	Items	Ratings	Reviews
Apps	20104	87271	13105	334927	752325
Automotive	360	2928	1579	4971	18822
Baby	4382	19445	6883	67060	159780
Beauty	5123	22359	11654	91824	195763
Music	1763	5541	3539	41152	64546
GoodReads	5094	3011	106213	518855	163549

Instead of following the traditional method of randomly splitting the data into 2 sets (training and testing), we use a slightly different approach. Because our algorithm works on every user separately, we split the data at the user level. We chose a random split of 80%-20% of ratings for each individual user and use the former as the training and trust calculation set, and the latter for testing our algorithm. All of our datasets are large enough to validate with hold-out method.

3.2 LIRA details

LIRA uses a framework which utilizes separate learners for every available source (See Figure 1). It learns the writing behavior and predictiveness of reviews from each source separately. For this example, the rating scale used is [1, 5].

We start the initial processing by creating data points which consists of the reviews by every source separately and marking them with target user’s rating for the corresponding item. That data point is then converted to a feature vector where the features are the words in the review and their values are their respective TF-IDF (Leskovec et al., 2014) values. Note that, LIRA is a general framework, within which any text to feature vector conversion approach can be used at this step.

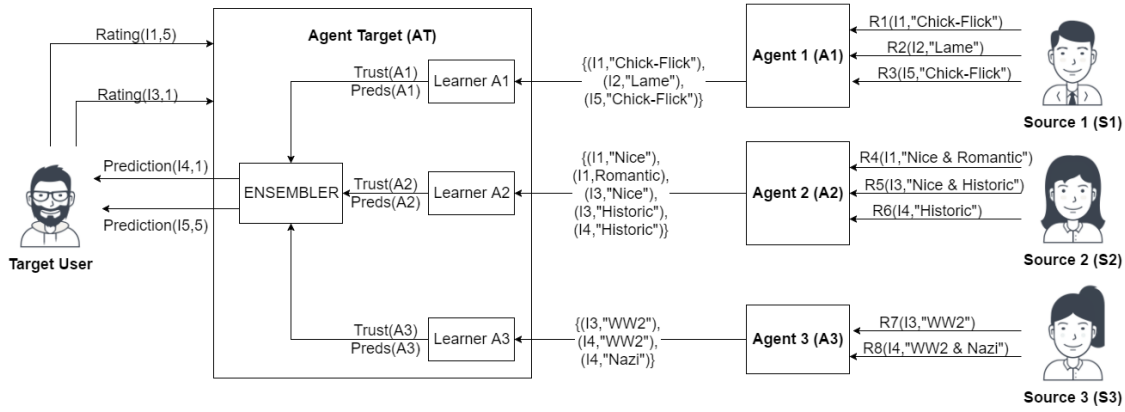


Figure 1: Overview of LIRA

3.2.1 Text Regression

We choose a machine-learning regression algorithm, such as Neural Network (Wang, 2003), Gaussian Process (Williams and Rasmussen, 2006), etc. to learn to predict the target user’s ratings for unseen items. Any regression algorithm can be used at this step of LIRA. In Figure 1 *Learner A1* will only find the common item *I1* among target user’s ratings and source *A1*’s reviews. *Learner A1* will learn that the word sequence ‘**Chick-Flick**’ is likely to result in the target user rating that item 5. Similarly *Learner A2* will find 2 common items, *I1* and *I3*, getting the reviews from *A2* for those items. *Learner A2* will learn that the word ‘**Romantic**’ by *A2* means high ratings by target user and ‘**Historic**’ by *A2* means low ratings by target user. *Learner A2* will also realize that, word ‘**Nice**’ by *A2* is not useful for predicting target user’s ratings.

3.2.2 Training

In the training phase, LIRA creates learners for every target user-source pair, to learn how the source’s reviews are predictive of the target user’s ratings. Hence, for every target user, $User_T$, and source, S , pair, LIRA creates $Learner_{TS}$ and feeds $(Review_{SI}, Rating_{TI})$ pairs to this learner. Here, $Review_{SI}$ is source S ’s review of $Item_I$ and $Rating_{TI}$ is $User_T$ ’s rating of $Item_I$.

Since the reviews are in free-form text, they are first converted to feature vectors using TF-IDF. It should be noted that text-to-feature conversion is done separately for every source in LIRA. Therefore, as long as a given source is consistent with its individual writing style and language usage, stop-word elimination will eliminate the words that are not informative only for that particular source’s reviews. For example if the source starts every review with a **Merhaba**², this word will be identified as a stop-word, and will be eliminated.

These feature vectors are used to train, $Learner_{TS}$, using a plug-and-play regression algorithm. The algorithm for training LIRA is presented in Algorithm 1. Note that if a source has not reviewed any items rated by the target user, it cannot be used as a predictor for this user’s rating, and is eliminated from consideration (see line 4).

3.2.3 Trust Calculation


We have mentioned in Section 2.3.3 that existing recommendation systems assign a trust value between users either by using rating similarities between users or by analyzing social connections between them. Unlike those systems, we decided to assign trust values to represent their capability of making predictions for the target user.

²Hello in Turkish.

Algorithm 1: LIRA Training For a Target User U_{ser_T} .

Input: $Ratings_T$, set of ratings provided by T
Input: $Sources$, set of sources who provides reviews
Output: $LearnerSet_T$, set of trained regression models for every source

- 1 $LearnerSet_T = \{\}$;
- 2 **foreach** $S \in Sources$ **do**
- 3 $Reviews_{TS} \leftarrow getReviewsBySourceForItemsRatedByT(Ratings_T, S)$;
- 4 **if** $Reviews_{TS} = \emptyset$ **then**
- 5 skip source;
- 6 $EliminateStopWords(Reviews_{TS}, Ratings_T)$;
- 7 $FeatureVectors \leftarrow TextToFeatureConversion(Reviews_{TS})$;
- 8 $Learner_{TS}.train(FeatureVectors, Ratings_{TS})$;
- 9 $LearnerSet_T.add(Learner_{TS})$

Figure 2: Behavior of adjusted sigmoid function with different β values.

As mentioned in Section 3.1, the available data set is split into training and trust calculation sets. First, all the learners are trained using the training set. Then we make all the learners predict every item in the trust calculation set. This approach is similar to the use of validation set in common supervised machine learning approaches. Root Mean Square Error (RMSE) for learner L is calculated as:

$$RMSE(L) = \sqrt{\frac{\sum_{I \in P_L} (Prediction(L, I) - r_I)^2}{|P_L|}}. \quad (7)$$

In Equation 7, P_L is the set of items which L can make predictions on, r_I is the actual rating, and $Prediction(L, I)$ is the rating prediction of L for item I .

This error measure is useful but not sufficient for a complete evaluation of our system. If we only consider the RMSE, a $Learner_A$ which made only one, albeit perfectly accurate, prediction will have the same trust value as a $Learner_B$ who was able to perfectly predict a large number of items. To address this issue, we include the number of predictions in our trust calculation. We use a sigmoid function to ensure that the magnitude of this count does not suppress the error value. As the number of predictions is non-negative, we choose the sigmoid function given in Equation 8 such that $sigmoid(0) = 0$ and $sigmoid(\infty) = 1$. We introduce the β parameter to adjust how quickly our sigmoid function reaches ~ 1 (see Figure 2).

$$sigmoid(x) = \frac{2}{1 + e^{-\beta x}} - 1. \quad (8)$$

We need to combine the error value and number of common items for every learner, to get the final trust value. Trust should be positively correlated with the number of common items. So we choose the sigmoid function, with the number of common items as argument, as the numerator of the trust expression. The




Figure 3: Behavior of trust function with different α and β values.

error measure calculated must be negatively correlated with the trust value. As the best case has 0 errors, to ensure the trust value lies between $[0, 1]$ we use $(1 + \text{errors})$ as the denominator term. For tuning the relative importance of the number of common items and error, we use an α parameter as the exponent of the denominator. The resultant trust value is then calculated as

$$\text{trust}(L) = \frac{\text{sigmoid}(|P_L|)}{(1 + \text{RMSE}(L))^\alpha}. \quad (9)$$

Figure 3 shows how the importance of RMSE and the number of predictions change according to the values assigned to α and β parameters.

In some cases a prediction may be made by a new learner, where they do not have any previous predictions to calculate a trust value. To avoid losing the information from those learners, we assign them a default trust value. Since we do not know about the capabilities of this learner yet, we conservatively assign them the minimum possible trust value, which is calculated by assuming they have made a single prediction with maximal prediction error possible. Since the maximum value of error in a prediction can be the difference between maximum and minimum values in the rating scale we end up with the following expression (Equation 10) for the trust value of the new learner (l_{new}):

$$\text{trust}(l_{new}) = \frac{\text{sigmoid}(1)}{(1 + \text{Rating}_{max} - \text{Rating}_{min})^\alpha}. \quad (10)$$

Algorithm 2: LIRA Trust Calculation For a Target User U_{ser_t}

Input: $Sources$, set of sources who provides reviews
Input: $Predictors_t$, set of regression models previously trained for t separately for every source
Input: $Ratings_t$, set of ratings provided by t
Output: $TrustVector_t$, a vector which holds t 's trust values for the sources

```

1 foreach  $s \in Sources$  do
2    $Learner_s \leftarrow Predictors_t.get(s)$ ;
3    $Count_s = 0$ ;
4   foreach  $i \in Ratings_t$  do
5     if  $s$  has a review for  $i$  then
6        $Count_s \leftarrow Count_s + 1$ ;
7        $Text \leftarrow getReview(s, i)$ ;
8        $FeatureVector \leftarrow TextToFeatureConversion(Text)$ ;
9        $Prediction \leftarrow Learner_s.predict(FeatureVector)$ ;
10       $Error_s.add(Rating_{ti}, Prediction)$ ;
11  $Trust_{ts} \leftarrow CalculateTrust(Error_s, Count_s)$ ;

```

3.2.4 Predicting target user's ratings for candidate items

In the prediction step of LIRA, we predict the rating for the unseen items by the target user. To achieve this, we will use the regression model trained in Section 3.2.2 and trust values calculated in Section 3.2.3.

For a candidate item, i , we first retrieve every source which has written a review for i and also has a learner trained for target user U_{ser_t} . Every source's review is separately fed into their respective learner for t , which calculates predictions by individual sources. These predictions are then combined using a weighted average approach, where the weights correspond to trust values for each source. To avoid losing any predictive value in reviews from new sources, which did not have any reviews that can be used in the trust calculation step, we set their trust level to a default value explained in Equation 10.

$$Prediction(t, i) = \frac{\sum_{s \in (Sources_i \cap TrainedSources_t)} Trust_{ts} \times Prediction_s(t, i)}{\sum_{s \in (Sources_i \cap TrainedSources_t)} Trust_{ts}} \quad (11)$$

The prediction LIRA makes for a $TargetUser_t - Item_i$ pair is calculated using the formula in Equation 11. In this equation $Trust_{ts}$ is the trust value calculated for $Source_s$ for $TargetUser_t$'s ratings, as described in Section 3.2.3. $Prediction_s(t, i)$ is the prediction made by $Source_s$'s review for $TargetUser_t - Item_i$. $Source_i$ is the set of sources who have provided reviews for $Item_i$ and $TrainedSources_t$ is the set of sources with a trained learner for $TargetUser_t$. Note that, for some $TargetUser_t - Item_i$ pairs, $Sources_i \cap TrainedSources_t = \emptyset$. Therefore LIRA will not be able to make predictions for those pairs.

Learners which do not have a review written for the item we are trying to predict the rating for cannot make prediction and we exclude those sources. For example, *Learner A2* will not be able to make a prediction for *I5*.

The final prediction for an item I_X , r_{I_X} is calculated as:

$$r_{I_X} = \frac{\sum_{u \in Rev(I_X)} Trust(L_u) \times Prediction(L_u, I_X)}{\sum_{u \in Rev(I_X)} Trust(L_u)}, \quad (12)$$

where $Rev(I_X)$ is the set of reviewers who have provided review for item I_X , $Trust(u)$ is the trust value for that reviewer, and $Prediction(L_u, I_X)$ is the prediction by *Learner* u , L_u , for item I_X .

3.3 Scalability

Given the description of the approach, it may appear that this approach is computationally expensive because of the need for separately training a learner for every user-source pair, where every reviewer is

Algorithm 3: LIRA Rating Prediction For a Target User, $User_t$ and Candidate Item, $Item_i$

Input: $Sources_i$, set of sources who have provided reviews for i

Input: $Predictors_i$, set of regression models previously trained for t separately for every source

Input: $TrustVector_t$, a vector which holds t 's trust values for the sources

Output: \hat{r}_{ti} , predicted rating of t for i

```

1  $TotalTrust \leftarrow 0$ ;
2  $RawPrediction \leftarrow 0$ ;
3 foreach  $s \in Sources_i$  do
4   if  $s \in Predictors_t$  then
5     if  $s \in TrustVector_t$  then
6        $SourceTrust \leftarrow TrustVector_t.get(s)$ ;
7     else
8        $SourceTrust \leftarrow defaultTrust$ ;
9      $Learner_s \leftarrow Predictors_T.get(s)$ ;
10     $Text \leftarrow getReview(s, i)$ ;
11     $FeatureVector \leftarrow TextToFeatureConversion(Text)$ ;
12     $Prediction \leftarrow Learner_s.predict(FeatureVector)$ ;
13     $TotalTrust \leftarrow TotalTrust + SourceTrust$ ;
14     $RawPrediction \leftarrow RawPrediction + SourcePrediction \times SourceTrust$ ;
15  $\hat{r}_{ti} = RawPrediction/TotalTrust$ ;

```

used as a source for every other user. For a system with N users, where each user is also a reviewer, the number of learned models will be $N \times (N - 1)$. However, some properties of the algorithm can be used to redress this problem as follows:

1. It is immediately obvious that creating the classifiers for separate users are independent processes. Therefore this approach is easily parallelizable.
2. More importantly, LIRA is designed to be an agent-based system. So, every target user can run their own agent, and therefore their own learners, on their local computer, allowing them to get recommendations from the system without the burden of running LIRA for any other user³. This way they will be training only one classifier which contains N sub-learners, where N is the number of sources, which will not be computationally prohibitive.

Another approach, to make the system even more efficient and part of our planned future work, is to train only a fixed number of most trustworthy sources for each target user. Assuming the top K sources per user is identified, the number of total models in the system will be reduced to $N \times K$. Periodically, the pool of trusted sources can be updated by a process of exploration. This problem can be represented as an instance of Multi-Arm Bandit problem (Gittins et al., 2011) where reviewers correspond to the available options and their accuracy on predicting the target user's ratings is the pay-off for the corresponding option.

4 Experimental Setup

In this section we explain the experiments we run, and the motivation behind them.

4.1 Selection of Methods to Use

We evaluate our algorithm on three datasets to demonstrate the general applicability and domain independence of our approach. We selected 'Automotive', 'Baby', and 'Digital Music' as the experimental

³This means that LIRA implementations should be distributed and not susceptible to the bottlenecks and single-point-of-failure issues plaguing centralized approaches.

domains because of their distinct characteristics. We are assuming that the opinions are mostly objective and based on the quality of the item for the ‘Automotive’ domain, that the opinions are mostly subjective and base on the preference of users for the ‘Music’ domain, and that the opinions will be a combination of objective quality and subjective personal preference in the ‘Baby’ domain.

For selecting appropriate method for regression and the tokenizer used in text to feature conversion, we run LIRA using the candidate methods on 100 randomly selected users and test which method yields the best results.

4.1.1 Regression Method Selection

To choose which regression method to use, we have tested LIRA using following regression learning algorithms: Neural Network, Gaussian Process, and SMOReg. For these algorithms, we used the default settings from WEKA. For the number of hidden nodes is $(attributes + classes)/2$, learning rate is 0.3, and sigmoid function has been used as the activation function. Kernel function used for both Gaussian Process and SMOReg is *Polynomial Kernel Function* (Fan et al., 1995).

4.1.2 Tokenizer Selection

To choose which tokenizer method to use, we have tested LIRA using following tokenizers: Word Tokenizer which only extracts words from a given text, and N-Gram tokenizer, which tries to extract N-Grams, including unigrams (words)(Mayfield and McNamee, 2003). We limited the N-Grams to unigrams, bigrams, and trigrams, to reduce to complexity and prevent the data from getting too sparse.

4.2 All Ratings Experiment

For this set of experiments we have used all the ratings and reviews in the dataset. LIRA works separately for every user, so for every user LIRA goes over every other users’ review to predict the rating of the target user.

As LIRA requires a separate trust calculation set, it splits the given training set 75%-25%⁴, and then uses the larger subset for training and the smaller one for trust calculation. Similar to the split explained in Section 3.1, this split is also done at the user level. This process is also used for the next set of experiments.

We have chosen to compare our algorithm against state of the art algorithms in both rating based and review based categories. We have chosen RegSVD (Paterek, 2007) as the rating based competitor, and HFT (McAuley and Leskovec, 2013) as the review based competitor.

4.3 100 User Experiment

This set of experiments are done to test LIRA’s capability of making use of the reviews from other sources. We randomly choose 100 users from the dataset as the group of target users. In our case. other sources then correspond to reviews from other users not in the chosen set. In general, however, these others sources can be taken from completely different platforms, such as websites, blogs, critique reviews, etc.

For the 100 users chosen, we only use their own ratings and not their own reviews, but use the reviews from all the users in the dataset, to predict their ratings. Hence, we consider all the users, except the selected 100, as outside sources for reviews. Note that the reviews from users, except the target user, are also used in the prediction, since those 99 users are considered on-site reviewers.

Since the competing algorithms are only capable of working with on-site data, we test them using ratings and reviews of the 100 selected users.

This approach is motivated by the following real life scenario. Consider a newly built movie recommendation web site which has few (in our example 100) users. As this system will only have the ratings and reviews of these few users, the only rating data that can be used for the recommendation systems is that. However LIRA can make use of outside sources for the recommendation, such as critique reviews, fan blogs, etc.

⁴This 75%-25% split is made on the training set. Since the whole dataset was split as 80% training set and 20% test set, the ratios on the whole split becomes 60% model building set, 20% weight calculation set, and 20% test set.

Recommendation services with only such small number of users (100 users in this case), corresponds to very sparse datasets. LIRA is not adversely affected by this sparsity, as it is able to make use of reviews from other outside sources.

We made sure that the same set of 100 randomly selected users and exactly the same training and test sets are used to test competing algorithms and LIRA.

4.4 Error Measurement

To calculate the error made by our algorithms and competing algorithms, we used the two most commonly used error measurement approaches in recommendation systems: Mean Absolute Error(MAE) (Willmott and Matsuura, 2005) and Root Mean Squared Error(RMSE) (Levinson, 1946).

4.4.1 Mean Absolute Error(MAE)

Mean Absolute Error(MAE) is a very intuitive error measurement approach. As the name suggests MAE is an average of the absolute errors made while predicting each instance (see Equation 13). MAE is primarily used when the values to be predicted are in the same scale and every error made by the predictor is equally important.

$$MAE = \frac{\sum_{(u,i)|R_{u,i}} |r_{u,i} - \hat{r}_{u,i}|}{|\{(u,i)|R_{u,i}\}|} \quad (13)$$

In this equation, $R_{u,i}$ is a boolean value signaling whether user u has rated item i in our data set, as we can only measure the error if we know the actual rating. Variables $r_{u,i}$ and $\hat{r}_{u,i}$ denote the actual and predicted rating respectively. Finally the sum of the absolute differences between actual and predicted ratings are divided by total number of predictions to obtain the average.

4.4.2 Root Mean Squared Error(RMSE)

Root Mean Squared Error(RMSE) calculates the standard deviation of the errors on the instances by the predictor. RMSE is primarily used when the greater errors are to be penalized harsher. Since error values are squared, RMSE of uniformly distributed errors will be lower than randomly distributed errors, given that MAE remains constant.

The formulation for MAE is given in Equation 14.

$$RMSE = \sqrt{\frac{\sum_{(u,i)|R_{u,i}} (r_{u,i} - \hat{r}_{u,i})^2}{|\{(u,i)|R_{u,i}\}|}} \quad (14)$$

In this equation, $R_{u,i}$ is a boolean showing whether user u has a rating on item i in our data set, as we can only measure the error if we know the actual rating. Variables $r_{u,i}$ and $\hat{r}_{u,i}$ denote the actual and predicted rating respectively. Finally the sum of the squared differences between actual and predicted ratings are divided by total number of predictions and square-rooted to get RMSE value.

5 Results

In this section we will present and analyze the results from the above-mentioned set of experiments.

5.1 Selection of Methods to Use

5.1.1 Regression Method Selection

We compare error per user by means of MAE and RMSE for three candidate regression algorithms: Neural Network, Gaussian Process, and SMOReg.

Box-plots in Figure 4 shows that Gaussian Process algorithm outperforms other two algorithms, and hence we choose Gaussian Process algorithm to conduct our further experiments.

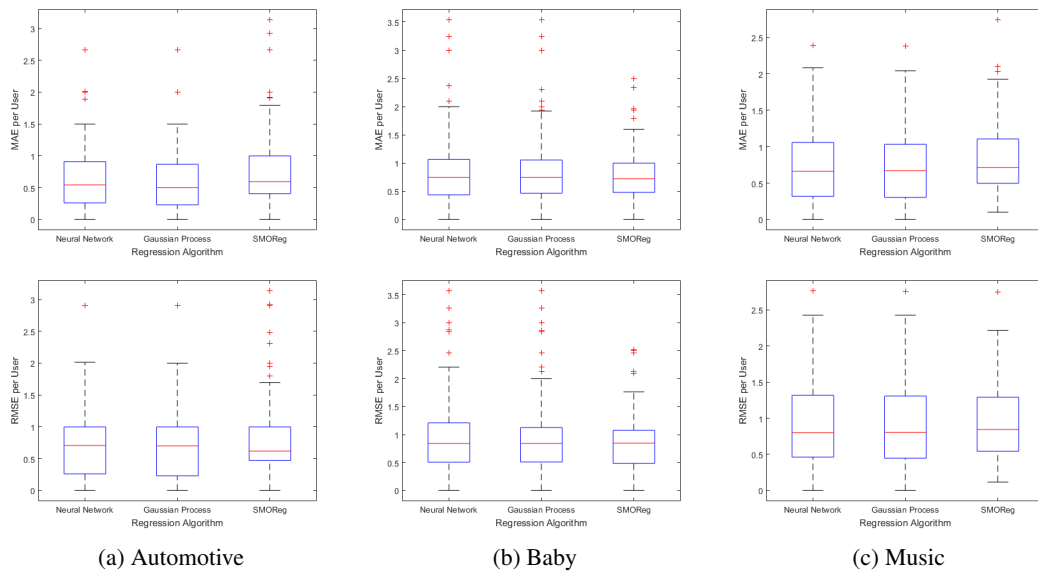


Figure 4: Box-Plot of errors per user made by LIRA using each candidate regression algorithm

5.1.2 Tokenizer Selection

We compare error per user by means of MAE and RMSE for two tokenizer, Word Tokenizer which only extracts words, and N-Gram tokenizer which also extracts N-Grams.

Box-plots in Figure 5, and data in Table 2 shows that the average errors with or without N-Grams are approximately same. Given that including N-Grams do not increase the accuracy of the system, we follow the *Occam's Razor* principle and conduct further experiments with only the Word Tokenizer.

Table 2 Average error of LIRA using two tokenizers.

	Word Tokenizer		N-Gram Tokenizer	
	MAE	RMSE	MAE	RMSE
Automotive	0.7945	0.8804	0.8181	0.9011
Baby	0.7965	0.8831	0.7913	0.8752
Music	0.8293	0.9683	0.8289	0.9676

5.2 All Ratings Experiment

The results of this experiment compare LIRA against the competing algorithms, in the condition where all the data available comes from the same environment, therefore all 3 algorithms are able to make use of all the data. The resulting MAE and RMSE values are presented in Table 3, and bold values mark the minimum MAE and RMSE in every domain. These results show that LIRA outperforms both competing algorithms on all six domains. Since LIRA is language independent, it is able to generate better predictions than competing algorithms, especially in the *GoodReads* domain. It can also be seen that while both competing algorithms suffer from limited data in small sized domains, such as *Automotive*, LIRA is still able to make accurate predictions.

As explained in Section 3.2.4, for $TargetUser_t - Item_i$ pairs where t does not have learners trained for any of the sources who reviewed i , LIRA will fail to make any prediction. Therefore, the coverage of LIRA is not 100%.

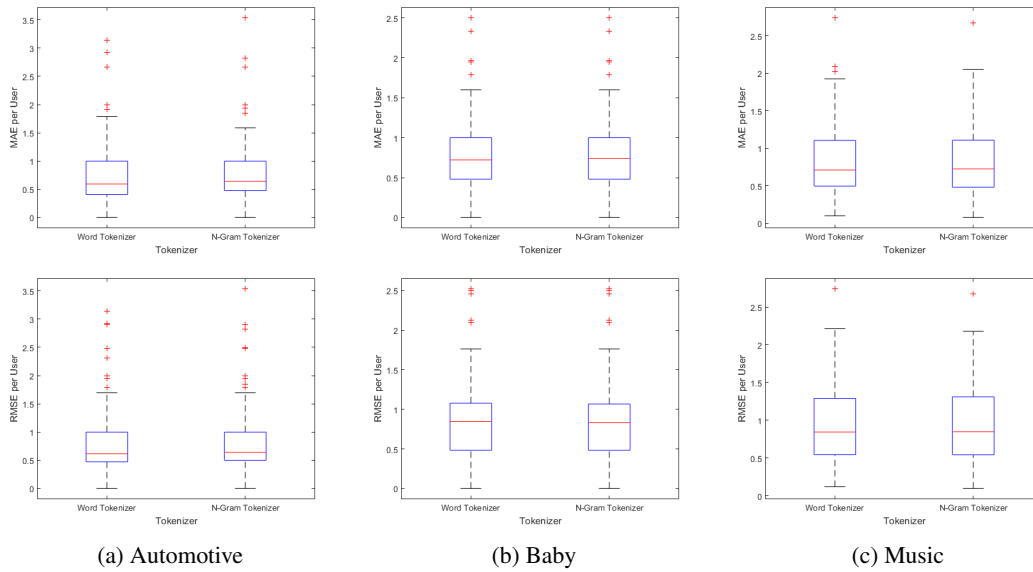


Figure 5: Box-Plot of errors per user made by LIRA using each tokenizer.

Table 3 Results of LIRA against competing algorithms in “All Ratings Experiment”

	RegSVD		HFT		LIRA		
	MAE	RMSE	MAE	RMSE	MAE	RMSE	Coverage
Apps	1.032	1.351	1.312	1.813	0.919	1.152	(0.74)
Automotive	1.702	2.081	1.262	1.804	0.632	0.944	(0.78)
Baby	1.037	1.340	1.243	1.710	0.867	1.111	(0.82)
Beauty	1.079	1.420	1.091	1.520	0.843	1.120	(0.81)
Music	0.757	1.022	0.864	1.231	0.655	0.958	(0.98)
Goodreads	1.015	1.396	0.818	1.093	0.748	0.963	(0.86)

We applied paired t-tests to verify if the improvements LIRA provides are significant. The results show that all the improvements are significant at $p < 0.001$ level (see Table 4).

Table 4 t values for t-test comparison of LIRA against competing algorithms.

	LIRA vs. RegSVD	LIRA vs. HFT
Apps	29.808	74.779
Automotive	19.546	10.613
Baby	14.915	24.037
Beauty	23.527	21.869
Music	7.681	16.005
Goodreads	67.618	20.865

5.3 100 User Experiment

The results of this experiment compare LIRA against the competing algorithms, in the condition where data from external sources are available. Since only LIRA is able to make use of external sources, competing algorithms will only be able to make use of data from the selected 100 users.

Table 5 Results of LIRA against competing algorithms in “100 User Experiment”

	RegSVD		HFT		LIRA		
	MAE	RMSE	MAE	RMSE	MAE	RMSE	Coverage
Apps	2.223	2.604	1.594	2.940	1.034	1.384	(0.76)
Automotive	2.323	2.682	1.559	2.070	0.677	0.998	(0.80)
Baby	2.842	3.100	1.190	1.550	0.814	1.177	(0.87)
Beauty	2.604	2.972	1.481	1.820	0.939	1.172	(0.83)
Music	2.086	2.472	1.136	1.670	0.601	0.886	(0.98)
Goodreads	2.128	2.472	1.207	1.687	0.790	1.024	(0.84)

The resulting MAE and RMSE values are presented in Table 5, and bold values mark the minimum MAE and RMSE in every domain. These results show that LIRA outperforms both competing algorithms on all six domains. These results show that LIRA continues to produce low error rates and is not significantly affected by small size of the dataset. On the other hand, error rate of the competing rating-based algorithm, RegSVD, drastically increases when the number of ratings is low. Error rate of competing review-based system, HFT, also increases, as it is only capable of making use of reviews which are accompanied with ratings⁵.

6 Conclusions

While most of the work on recommendation systems is based on ratings and rating similarities, limited representation power of ratings highlight the need for new recommendation systems which can make use of the other information available to improve recommendation accuracy. With the introduction of Web 2.0 and burgeoning of online platforms, which rely on user-generated content, a rapidly increasing volume of on online and off-site sources make available item reviews that can provide valuable insights about user preferences and can be used to improve the coverage and accuracy of existing recommendation systems. Reviews written by users are more representative and detailed representation of a user’s preference for an item. Descriptive reviews also express users’ opinions about the features of candidate items, and help us develop a deeper understanding of their preferences.

In this paper we have proposed an agent-based approach, using the trust aspect of agent societies, that uses regression algorithms to predict the ratings of a target user for a candidate item. Our approach creates agents for both target users and reviewers to create accurate recommendations. LIRA consists of target user’s agent, which uses messages, in the form of words or sequence of words in their reviews, from the other agents. The agent of the target user learns predictive knowledge from those messages which are subsequently used for producing recommendations. LIRA assigns trust to each source’s agent using their accuracy on a set of unknown items. Final prediction of LIRA is calculated using a trust-weighted average of sources’ individual prediction. Among other benefits, this approach allows the recommendation system to differentiate two users, who use the same word for different reasons, by inferring from their past reviews where identical words have different predictive values for the target user when used by different sources.

Another key feature of LIRA is that it does not use the ratings of the sources and uses only target user’s ratings. This results in the LIRA approach having two key advantages. Firstly, it can use reviews from different sources, such as different websites or fan blogs, without needing those texts to be accompanied with a rating value. This makes LIRA useful for addressing the “cold start” problem of new recommender systems which do not have many users to write reviews for them. Secondly, by not taking the sources ratings into account, LIRA treats reviews as descriptive texts. For example, if one or more sources complain about a movie, because they found it “too romantic”, LIRA can infer that the movie might be recommended to a target user who does enjoy romantic movies.

⁵It should be noted that *100 User Experiment* is conducted to show that accuracy of LIRA does not get affected by fewer number of users. Since available data is a subset of data used in *All Ratings Experiment*, which shows the difference is significant, significance values are not necessary for *100 User Experiment*.

Since LIRA does not deal with the meaning of the words, and only use them as signals between agents, the language in which the review was written in becomes irrelevant. As long as sources are consistent with their choice of language or writing style, LIRA will be able to effectively leverage their reviews. The language-independent property of LIRA, whereby reviews written in different languages can be gainfully utilized, is a distinctive advantage compared to other review-based recommender systems which combines reviews from different sources for purposes such as topic modelling (McAuley and Leskovec, 2013; Agarwal and Chen, 2010).

Last, but not the least, a key advantage of LIRA is that it does not require the target users to write reviews. It is well-known that people write reviews much more infrequently than they rate items online. Hence LIRA can be used in many scenarios where review based recommender systems that require target users to write reviews can be used.

We presented experimental results from two sets of experiments. In the first set there was a rating accompanying every review. This experiment showed that analyzing reviews from different sources separately, allows LIRA to outperform state-of-art algorithms. Domains where LIRA was more successful than its competitors are likely those where people are more capable and likely to write more descriptive reviews. The second set of experiments show that when there is very little user and rating data available, only LIRA can make use of other sources of information to make accurate recommendations and thus outperform its competitors.

One of the properties of LIRA, which might look like a disadvantage, is that it needs to train regression algorithms for every user-source pair. However it should be mentioned that the agents work independently from each other which makes this approach highly parallelizable. We are currently also working on an approach which provides the same advantages that LIRA provides, without having to create separate learners for every source.

References

- Adomavicius, G. and Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE transactions on knowledge and data engineering*, 17(6):734–749.
- Agarwal, D. and Chen, B.-C. (2010). flda: matrix factorization through latent dirichlet allocation. In *Proceedings of the third ACM international conference on Web search and data mining*, pages 91–100. ACM.
- Alpaydin, E. (2014). *Introduction to machine learning*. MIT press.
- Bao, Y., Fang, H., and Zhang, J. (2014). Topicmf: Simultaneously exploiting ratings and reviews for recommendation. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 2–8.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022.
- Chen, L. and Wang, F. (2013). Preference-based clustering reviews for augmenting e-commerce recommendation. *Knowledge-Based Systems*, 50:44–59.
- Cristianini, N. and Shawe-Taylor, J. (2000). *An introduction to support vector machines*. Cambridge University Press Cambridge.
- Davis, J. M. (1958). The transitivity of preferences. *Behavioral science*, 3(1):26–33.
- Debnath, S., Ganguly, N., and Mitra, P. (2008). Feature weighting in content based recommendation system using social network analysis. In *Proceedings of the 17th international conference on World Wide Web*, pages 1041–1042. ACM.
- Deerwester, S. C., Dumais, S. T., Landauer, T. K., Furnas, G. W., and Harshman, R. A. (1990). Indexing by latent semantic analysis. *JASIS*, 41(6):391–407.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)*, pages 1–38.

- Fan, J., Heckman, N. E., and Wand, M. P. (1995). Local polynomial kernel regression for generalized linear models and quasi-likelihood functions. *Journal of the American Statistical Association*, 90(429):141–150.
- Gittins, J., Glazebrook, K., and Weber, R. (2011). *Multi-armed bandit allocation indices*. John Wiley & Sons.
- Golbeck, J. A. (2005). *Computing and Applying Trust in Web-based Social Networks*. PhD thesis, University of Maryland at College Park, College Park, MD, USA. AAI3178583.
- Hariri, N., Zheng, Y., Mobasher, B., and Burke, R. (2011). Context-aware recommendation based on review mining. *General Co-Chairs*, page 27.
- Homoceanu, S., Loster, M., Lofi, C., and Balke, W.-T. (2011). Will i like it? providing product overviews based on opinion excerpts. In *Commerce and Enterprise Computing (CEC), 2011 IEEE 13th Conference on*, pages 26–33. IEEE.
- Jakob, N., Weber, S. H., Müller, M. C., and Gurevych, I. (2009). Beyond the stars: exploiting free-text user reviews to improve the accuracy of movie recommendations. In *Proceedings of the 1st international CIKM workshop on Topic-sentiment analysis for mass opinion*, pages 57–64. ACM.
- Kazienko, P. and Musiał, K. (2006). *Recommendation framework for online social networks*. Springer.
- Leskovec, J., Rajaraman, A., and Ullman, J. D. (2014). *Mining of massive datasets*. Cambridge University Press.
- Leung, C. W.-k., Chan, S. C.-f., and Chung, F.-l. (2008). An empirical study of a cross-level association rule mining approach to cold-start recommendations. *Knowledge-Based Systems*, 21(7):515–529.
- Levi, A., Mokryn, O., Diot, C., and Taft, N. (2012). Finding a needle in a haystack of reviews: cold start context-based hotel recommender system. In *Proceedings of the sixth ACM conference on Recommender systems*, pages 115–122. ACM.
- Levinson, N. (1946). The wiener (root mean square) error criterion in filter design and prediction. *Journal of Mathematics and Physics*, 25(1):261–278.
- Linden, G., Smith, B., and York, J. (2003). Amazon.com recommendations: Item-to-item collaborative filtering. *Internet Computing, IEEE*, 7(1):76–80.
- Liu, H., He, J., Wang, T., Song, W., and Du, X. (2013). Combining user preferences and user opinions for accurate recommendation. *Electronic Commerce Research and Applications*, 12(1):14–23.
- Lops, P., De Gemmis, M., and Semeraro, G. (2011). Content-based recommender systems: State of the art and trends. In *Recommender systems handbook*, pages 73–105. Springer.
- MacKay, D. J. (1998). Introduction to gaussian processes. *NATO ASI Series F Computer and Systems Sciences*, 168:133–166.
- Manning, C. D., Raghavan, P., Schütze, H., et al. (2008). *Introduction to information retrieval*, volume 1. Cambridge university press Cambridge.
- Massa, P. and Avesani, P. (2007). Trust metrics on controversial users: balancing between tyranny of the majority and echo chambers. *International Journal on Semantic Web and Information Systems*, 3(1):39–64.
- Massa, P. and Bhattacharjee, B. (2004). Using trust in recommender systems: an experimental analysis. In *Trust Management*, pages 221–235. Springer.
- Mayfield, J. and McNamee, P. (2003). Single n-gram stemming. In *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*, pages 415–416. ACM.
- McAuley, J. and Leskovec, J. (2013). Hidden factors and hidden topics: understanding rating dimensions with review text. In *Proceedings of the 7th ACM conference on Recommender systems*, pages 165–172. ACM.
- McAuley, J., Pandey, R., and Leskovec, J. (2015a). Inferring networks of substitutable and complementary products. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794. ACM.

- McAuley, J., Targett, C., Shi, Q., and van den Hengel, A. (2015b). Image-based recommendations on styles and substitutes. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 43–52. ACM.
- Middleton, S. E., Shadbolt, N. R., and De Roure, D. C. (2004). Ontological user profiling in recommender systems. *ACM Transactions on Information Systems (TOIS)*, 22(1):54–88.
- Mitchell, T. M. (1997). *Machine Learning*. McGraw-Hill.
- Mooney, R. J. and Roy, L. (2000). Content-based book recommending using learning for text categorization. In *Proceedings of the fifth ACM conference on Digital libraries*, pages 195–204. ACM.
- O’Donovan, J. and Smyth, B. (2005). Trust in recommender systems. In *Proceedings of the 10th international conference on Intelligent user interfaces*, pages 167–174. ACM.
- Palaniswami, M. and Shilton, A. (2002). Adaptive support vector machines for regression. In *Neural Information Processing, 2002. ICONIP’02. Proceedings of the 9th International Conference on*, volume 2, pages 1043–1049. IEEE.
- Paterek, A. (2007). Improving regularized singular value decomposition for collaborative filtering. In *Proceedings of KDD cup and workshop*, volume 2007, pages 5–8.
- Sarwar, B., Karypis, G., Konstan, J., and Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*, pages 285–295. ACM.
- Schein, A. I., Popescul, A., Ungar, L. H., and Pennock, D. M. (2002). Methods and metrics for cold-start recommendations. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 253–260. ACM.
- Soleymani, M., Aljanaki, A., Wiering, F., and Veltkamp, R. C. (2015). Content-based music recommendation using underlying music preference structure. In *Multimedia and Expo (ICME), 2015 IEEE International Conference on*, pages 1–6. IEEE.
- Sugiyama, K., Hatano, K., and Yoshikawa, M. (2004). Adaptive web search based on user profile constructed without any effort from users. In *Proceedings of the 13th international conference on World Wide Web*, pages 675–684. ACM.
- Ungar, L. H. and Foster, D. P. (1998). Clustering methods for collaborative filtering. In *AAAI workshop on recommendation systems*, volume 1, pages 114–129.
- Vapnik, V. (2013). *The nature of statistical learning theory*. Springer science & business media.
- Wang, J., Yin, J., Liu, Y., and Huang, C. (2011). Trust-based collaborative filtering. In *Fuzzy Systems and Knowledge Discovery (FSKD), 2011 Eighth International Conference on*, volume 4, pages 2650–2654. IEEE.
- Wang, S.-C. (2003). Artificial neural network. In *Interdisciplinary computing in java programming*, pages 81–100. Springer.
- Williams, C. K. and Rasmussen, C. E. (2006). Gaussian processes for machine learning. *the MIT Press*, 2(3):4.
- Willmott, C. J. and Matsuura, K. (2005). Advantages of the mean absolute error (mae) over the root mean square error (rmse) in assessing average model performance. *Climate research*, 30(1):79–82.
- Yang, D., Zhang, D., Yu, Z., and Wang, Z. (2013). A sentiment-enhanced personalized location recommendation system. In *Proceedings of the 24th ACM Conference on Hypertext and Social Media*, pages 119–128. ACM.