

September 2003

A Model-Based Approach to Recommending Partners

Frank Färber

Johann Wolfgang Goethe-Universität Frankfurt, ffaerber@wiwi.uni-frankfurt.de

Tim Weitzel

Johann Wolfgang Goethe-Universität Frankfurt

Tobias Keim

Univentures GmbH

Oliver Wendt

Univentures GmbH

Follow this and additional works at: <http://aisel.aisnet.org/wi2003>

Recommended Citation

Färber, Frank; Weitzel, Tim; Keim, Tobias; and Wendt, Oliver, "A Model-Based Approach to Recommending Partners" (2003).
Wirtschaftsinformatik Proceedings 2003. 24.
<http://aisel.aisnet.org/wi2003/24>

This material is brought to you by the Wirtschaftsinformatik at AIS Electronic Library (AISeL). It has been accepted for inclusion in Wirtschaftsinformatik Proceedings 2003 by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

In: Uhr, Wolfgang, Esswein, Werner & Schoop, Eric (Hg.) 2003. *Wirtschaftsinformatik 2003: Medien - Märkte - Mobilität*, 2 Bde. Heidelberg: Physica-Verlag

ISBN: 3-7908-0111-9 (Band 1)

ISBN: 3-7908-0116-X (Band 2)

© Physica-Verlag Heidelberg 2003

A Model-Based Approach to Recommending Partners

Frank Färber, Tim Weitzel

Johann Wolfgang Goethe-Universität Frankfurt

Tobias Keim, Oliver Wendt

Univentures GmbH

Abstract: Searching for and selecting qualified partners is a core task in many business contexts. Empirical research among Germany's top 1,000 firms discloses that internet-based platforms are effectively used as a personnel marketing channel but cannot increase the matching quality between jobs and candidates. Using e-recruitment as an example, we show how the matching quality can be substantially improved by means of a probabilistic latent aspect model developed in this paper. The underlying method incorporates findings from collaborative filtering and hybrid approaches to automated recommendation and is based on a model of personal attributes derived from research on team building and work psychology.

Keywords: personnel selection, collaborative filtering, probabilistic modeling

1 Introduction

The flexible building of partnerships is getting increasingly important in many business situations such as team building, the formation of joint ventures or consortia for bidding processes, and of course the recruitment of new employees [MaLa98, p. 146]. Partially driven by the progress in information and communication technology, the configuration of new partnerships takes place more frequently and, hence, needs to be accomplished efficiently. This implies growing technological requirements and a multitude of research questions in information systems.

At present, especially cost and time aspects are fostering the emergence of a variety of internet-based platforms allowing to search for and get in contact with potential partners [GaWa01, pp. 29ff.; LaPi99]. Examples range from Competence Site (www.competence-site.de) offering several marketplaces for different kinds of competencies to international career networks like Monster.com.

As opposed to the reduction of cost and time, the existing solutions merely improve matching quality by leveraging the great amount of online information on

partners and past matching processes. Many platforms employ simple Boolean search to support the matching process. This, however, is not efficient. Even though different application domains in which partners are matched to each other might differ in matching criteria and types of partners to be matched (e.g., individuals, teams, companies), usually underlying aspects that cannot be represented by certain keywords are crucial for a good match.

Many companies, e.g., do not search for specific skills such as "Java programming" when hiring new employees. They rather look for aspects such as versatility or leadership skills that might be derived from the type of experiences and activities a potential candidate discloses in his resume. These, however, can rarely be found by searching for keywords. Among others, this drawback of current internet-based partner matching was validated by a survey among the 1,000 largest German companies using recruitment as an example (see Chapter 2).

In contrast to the majority of online partner matching mechanisms currently employed, a great amount of information systems research has been conducted on supporting searchers in finding the right information or products online (e.g., [WeWo00]). Especially automated recommendation systems have proven their usefulness in matching different types of items such as movies or research articles with user preferences [Bre⁺98; Pop⁺01] (see Chapter 3).

We try to apply these techniques to the partner matching problem considering the specifics of partner attributes as opposed to information content or product features (see Chapter 4). Our model-based approach is capable of dealing with underlying aspects and, thereby, addresses a major disadvantage of existing online matching mechanisms.

2 An example of partner matching: personnel recruitment

Matching or selecting partners is the objective of the recruitment process. In this chapter, we detail the above mentioned drawbacks of current matching platforms using online recruitment as an example. We, therefore, first describe the different phases of searching and hiring new employees (see section 2.1). We then draw conclusions on the impact of the internet on this process from a survey on modern recruitment practices that we conducted among the largest 1,000 companies in Germany (see section 2.2). A more detailed description of the results of this survey can be found in König et al. [Kön⁺03].

2.1 The recruitment process

Partnership building is a very important part of the human resources management functions. Especially, the evaluation of partners (i.e., employees, candidates) is a core element in recruitment as well as personnel development. In the following, we will focus on the process of recruiting new employees as depicted in Figure 1 [Albe98, pp. 56; Schn95, pp. 25].



Figure 1: Recruitment process

The process steps *employer branding* and *personnel attraction* involve activities undertaken by a company to establish a positive employer image and to attract candidates for specific job positions. While employer branding is more focused on the long-term strategic positioning of a company in the labor market using different marketing measures, personnel attraction aims at generating a sufficient number of applicants qualified for a specific job. The latter is divided into passive marketing activities such as job postings and active search activities such as the use of head hunters.

With the applications being received by the recruitment department, a workflow is initiated and has to be managed. This process, which is called *applicant management*, is actually a support function for the following *pre-selection* and *selection* processes. The objective of pre-selection and selection is to evaluate the quality of the candidate compared to the requirements defined by the job profile. The first of these two steps includes the screening of resumes and other application documents in order to decide if the candidate enters the selection stage. Finally, invited candidates are assessed using interviews and other assessment methods.

2.2 Results from a survey on modern recruitment practices in Germany

The survey on modern recruitment practices concentrated on the use of information and communication systems along the recruitment process as described above. The main system domains analyzed were corporate homepages, internet job portals and applicant management systems. We achieved a response rate of 19.6% (N=196).

The survey showed that both, the corporate homepage and internet job portals, are frequently used among the 1,000 largest German companies. 80% of the respondents stated that they frequently make use of the corporate homepage to attract candidates. Traditional print media (internet job portals) follows with 54% (48%)

of frequent or very frequent usage. However, when looking at the effectiveness, online channels fall behind traditional print media. While 62% of the respondents consider search activities in print media as effective, only 54% (38%) consider the corporate homepage (internet job portals) as effective.

A closer look at the benefits generated by internet-based recruitment so far shows that e-recruitment instruments have mainly been used as a cost-efficient alternative for personnel marketing measures to increase the total number of applicants (see Figure 2). Consequently, online job ads are considered the main advantage of internet job portals.

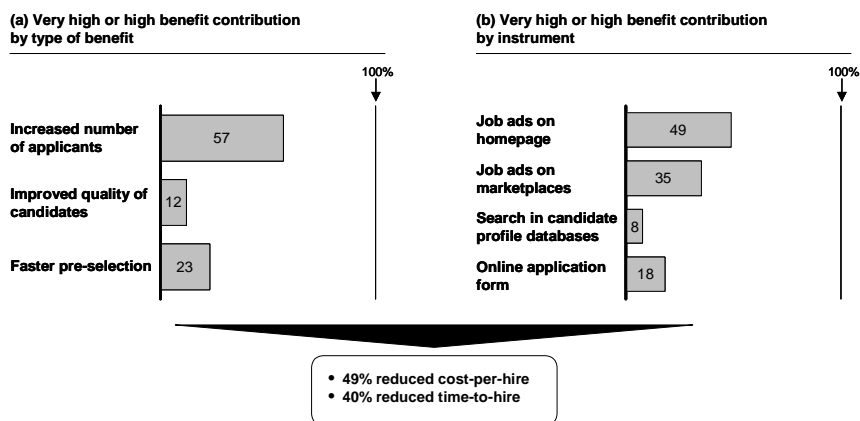


Figure 2: Benefits from e-recruitment

Generally, benefits of information technology in the recruitment domain can be clustered into three categories: reduction of costs, reduction of time-to-hire and increase in matching quality. According to our survey results, cost-per-hire and time-to-hire have already been reduced by a significant number of companies (49% resp. 40%). The improvement of matching quality, however, plays a less important role so far as only 12% of the respondents stated that they identified better candidates through the internet. Also, only 8% consider actively searching candidates in the resume databases of career portals as beneficial.

We identified two main reasons why the improvement of matching quality has not been a major benefit driver so far. First, applicant data is often not captured or stored in a structured digital format throughout the entire application workflow. As a consequence, filtering and rating mechanisms cannot be applied to these applications. Only 12% of applications are received through web-based forms as opposed to unstructured e-mail applications and paper-based applications. 21% enter candidate profiles that go beyond pure contact information into an applicant tracking system. The second main reason is the lack of adequate matching functionality provided by internet career portals. 44% of the respondents actively searched for qualified candidates in the resume databases of job portals, but only 8% were able

to generate significant benefits. This illustrates that simple keyword-based search functionality makes it difficult for users to find the right candidates even though they might be in the system.

3 An overview on automated recommendation

Recommender systems have evolved with the interactive environment of the internet. While users of this vast communication network have access to large amounts of information items and product descriptions, they have difficulties to find the right information or preferred products. In every day life, the process of finding and choosing the right things is usually supported by recommendations from other people that we trust or assume to have similar tastes, or we rely on reviews by trusted sources such as renowned newspapers.

However, with the large amounts of information about preferences and interests being captured on the internet (e.g., in the form of site visits and transactions) this data can be used to automatically infer recommendations to individual users [ReVa97, p. 56; Sar⁺00, pp. 158].

Techniques being applied in automated recommendation systems have their roots in information filtering and related fields such as information retrieval and categorization [FoDu92; BeCr92]. Hence, recommender systems are usually distinguished into those using *content-based filtering* (see section 3.1) and those using *collaborative filtering* (see section 3.2). While content-based methods recommend objects similar to those a user has preferred in the past, methods based on collaborative filtering identify other users with tastes similar to the current user and recommend objects those users have preferred [BaSh97, p. 66; Bre⁺98, p. 43].

3.1 Content-based filtering

Techniques for retrieving or filtering different kinds of objects based on their content or descriptions are called content-based filtering. They are mostly used for retrieving documents with unstructured textual information [Res⁺94; BaSh97]. The main steps of content-based filtering are the extraction of features representing the objects to be recommended, the capturing of either implicit or explicit user preferences towards these objects, and the prediction of the likelihood whether the user likes an object he has not seen yet based on the features of this object and the preference profile that has been created for this user (see Figure 3).




Figure 3: General process of content-based recommendation

Several methods have been proposed on how to extract and represent features from unstructured information [BaSh97, p. 67]. Text documents are often represented by a word-by-document matrix with the frequency of a word in a document as elements of the matrix [FoDu92, p. 53]. Thus, documents can be seen as vectors with dimensions equal to the number of words in the vocabulary. In order to match documents with user preferences, the latter are also represented as vectors of keywords.

A very simple measure to match preference profiles with the content of text documents is the standard cosine similarity function applied on the vector space representation of both the documents and the preference profiles:

$$s(d, p) = \frac{\sum_w n(d, w) \cdot n(p, w)}{\sqrt{\sum_w n(d, w)^2} \sqrt{\sum_w n(p, w)^2}} \quad (1)$$

with $n(d, w)$ being the term frequencies, i.e., the number of times a word w occurs in document d , and $n(p, w)$ representing the weight of word w in preference profile p [Hofm99, p. 294].

An important assumption of the cosine similarity function is that words are considered orthogonal or independent from each other [FoDu92, p. 53]. However, this is not very realistic in many cases as the use of words in documents is interdependent. As an alternative approach, latent semantic indexing or analysis (LSI, LSA) [FoDu92, p. 53; Hofm99] assumes an underlying or latent structure of the usage of words. Instead of representing high-dimensional vectors of words, a lower dimensional latent semantic space is used which provides information beyond the lexical level by revealing semantic relations between words. Mathematically, LSA can be based on singular-value decomposition or on a probabilistic latent semantic aspect model [Hofm99; HoPu99].

3.2 Collaborative filtering

Collaborative filtering aims at identifying items that the current user has not chosen yet based on the subjective opinions of other users with similar tastes [Res⁺94, p. 175; BaSh97, p. 66]. The content or descriptions of items are not considered, i.e., all that is known about an object is a unique identifier. The advantage of collaborative filtering compared to pure content-based systems is obvious. They can deal with any type of object independent of the complexity of their features [BaSh97, p. 67].

Users with preferences similar to the current user are those that have agreed with the active user's opinion in the past. They are called *neighbors* [Sar⁺00, p. 160]. The process of neighborhood formation and the associated algorithms are the core of the overall collaborative filtering-based recommendation process as it is depicted in Figure 4. While in content-based filtering the object features as well as the content-related user preferences have to be represented, in collaborative filtering only the actual behavior or opinion data has to be stored. Hence, no preprocessing comparable to the extraction of features in content-based filtering is necessary.

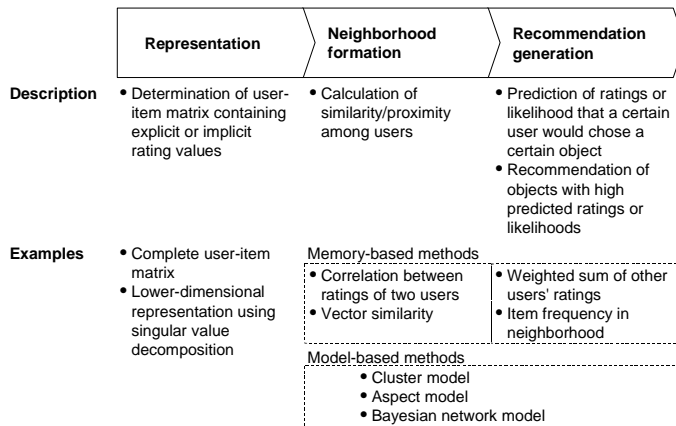


Figure 4: General process of collaborative filtering-based recommendation

The behavior or opinion data is expressed as numerical ratings or votes which are represented in the form of a user-item rating matrix R containing the rating values $r_{i,j}$ which stand for the rating of user i for item j . The matrix R is also referred to as the original representation of user ratings [Bre⁺98, p. 44; Sar⁺00, p. 161]. With the set of users being $U=\{u_1, \dots, u_m\}$ and the set of items (or objects) being $I=\{i_1, \dots, i_n\}$ the user-item matrix has the dimension $m \times n$.

Two types of collaborative filtering methods are distinguished [Bre⁺98, p. 44; Sar⁺01, p. 287]: *memory-based methods* and *model-based methods*. The main dif-

ference between those two types is that memory-based methods always use the entire rating data to generate recommendations while model-based systems determine parameters of a prediction model in advance which is then used to predict user-item ratings.

The determination of similarity between users' tastes, called neighborhood formation, is the first step in collaborative filtering. The earliest similarity measure that was used in this context was the Pearson correlation coefficient [Res⁺94; Bre⁺98, p. 44]. Hence, the similarity of the preference profiles of two users $a, i \in U$ is defined as:

$$\text{sim}(a, i) = \frac{\sum_{j \in I_a \cap I_i} (r_{a,j} - \bar{r}_a)(r_{i,j} - \bar{r}_i)}{\sqrt{\sum_{j \in I_a \cap I_i} (r_{a,j} - \bar{r}_a)^2 \sum_{j \in I_a \cap I_i} (r_{i,j} - \bar{r}_i)^2}} \quad (2)$$

where user a is the active user, user i is any other user, I_a and I_i are the sets of items that users a resp. i have recorded ratings for, and \bar{r}_a resp. \bar{r}_i are the mean ratings for those users defined as:

$$\bar{r}_x = \frac{1}{|I_x|} \sum_{j \in I_x} r_{x,j} \quad (3)$$

The similarities can then be used as weights to calculate predicted ratings $p_{a,j}$ of the active user for unseen items. Breese et al., e.g., use a weighted sum of the ratings of other users [Bre⁺98, p. 44].

Determining similarities between user preferences constitutes an explicit step in memory-based algorithms before ratings are predicted. In contrast to this, model-based approaches estimate or learn parameters of a probabilistic model based on past ratings used as training data. The ratings are then predicted with the estimated model. Examples for such probabilistic models include clustering models [Bre⁺98, p. 46; UnFo98], Bayesian network models [Bre⁺98, p. 46], and latent aspect models [HoPu99]. The latter will be reviewed in more detail in Chapter 4.

3.3 Hybrid approaches

The quality of pure collaborative filtering-based recommendation systems often suffers from the sparsity of the original rating matrix, i.e., the potential overlap of rated objects might be too small to infer good recommendations even for users with similar tastes. Pure content-based systems, on the other hand, do not produce good results if object features are too divers.

The main idea behind hybrid approaches is to combine both methods. Similarity measures for items as described in section 1.2.1 can be used to infer additional ratings for objects the user has not seen yet but are similar to objects already rated. Alternatively, ratings for an object as a whole can be treated as ratings of its features such as the director of a movie. In this case, the rating matrix does not consist of votes for objects but of votes for all possible objects features.

Melville et al. propose a hybrid algorithm using the term content-boosted collaborative filtering [Mel⁺02]. They supplement original user ratings with pseudo user ratings leading to a pseudo user-item matrix V consisting of the following elements:

$$v_{u,i} = \begin{cases} r_{u,i} & \text{if user } u \text{ rated item } i \\ c_{u,i} & \text{otherwise} \end{cases} \quad (4)$$

The elements $r_{u,i}$ denote the actual ratings and $c_{u,i}$ are ratings predicted by a content-based system. Collaborative filtering is then performed using this comparably dense matrix. Melville et al. apply weighting schemes to consider different confidences in the pure content-based predictions and in the correlation-based user similarity [Mel⁺02]. Balabanović and Shoham, Good et al., and Sarwar et al. provide other examples for hybrid approaches [BaSh97; Good⁺99; Sar⁺01].

4 A partner recommending model

The findings of the survey on current internet-based recruitment practices showed that searching for candidates is often only supported by methods that are not adequate to achieve a good matching quality between jobs and candidates. In the following sections we propose a method that exploits the advantages of collaborative filtering and hybrid methods to automated recommendations. We use a probabilistic latent aspect approach which belongs to the class of model-based collaborative filtering. The following section 4.1 introduces a model of personal attributes. Its structure and classification is based on research on team building and work psychology. Section 4.2 then introduces the composition of the model and its parameters. The final section of this chapter specifies the determination of model parameters and shows how existing methods can be adapted to lead to better matching results.

4.1 Partner attributes

Different types and aspects of human attributes can be distinguished. The main categorization dimensions that are considered in our model are the following (see also Figure 5):

- The differentiation *readily detectable* vs. *underlying* attributes refers to whether the attributes can be easily observed (e.g., hair color) or require a more or less extensive assessment (e.g., leadership skills). Readily detectable attributes also include those attributes that can be derived from known facts (e.g., mathematical skills derived from a diploma in mathematics). However, the quality of a match often builds on underlying attributes, such as a person's attitudes or values, that are difficult to appraise [Jack96, p. 57]. Also, underlying attributes often influence social factors as opposed to task-related factors. Both are very crucial for team effectiveness [West94, p. xiii] and, hence, must be considered in a matching process.
- *Situational* attributes that are dependent on the situation the attribute is related to are distinguished from those that are only tied to a person. We call the latter *independent* attributes which are, e.g., certain skills while skills applied on a certain task become competencies [Lin⁺01, p. 777] and are, thus, situational attributes. The concept of situational attributes is very important when matching human beings. Note that the situation also involves other partners. Guzzo and Bungard, e.g., state that tasks in teams do not require *co*-action between team members, but their *inter*-action and, thus, require attributes measuring these synergistic effects [Guzz96, p. 8; Bung90, p. 317].

In order to capture the first of the differentiations above, we do not treat the values of partner attributes as absolute, but relate them to the person who assessed the attribute value and to the assessment method that was used. Furthermore, we distinguish different types of assessment methods:

- *Standardized* methods evaluate an attribute based on rules and clearly defined measures. The input parameters of these methods can be other attributes or facts known about the person (e.g., measuring academic competencies by the number of publications) or the responses to a standardized test such as an ability test to evaluate mathematical skills. The input parameters are not necessarily objective as it is, e.g., the case with the personality test MBTI (Myers-Briggs Type Indicator) where the questionnaire is standardized but the responses are subjective appraisements. Hence, standardized assessment methods can be further distinguished into *objective* and *subjective* methods.
- *Individual* methods are not based on clearly defined rules to determine the attribute's value. The method only defines the type of input information and the domain of the result to make the method comparable when applied by different assessors. An example of such a method is the screening of resumes in the per-

sonnel selection process to assess a candidate's aptitude for a certain job position. This aptitude is in fact a situational attribute implicitly aggregating several competencies (the competencies could also be explicitly stated as selection criteria to provide the assessor a guideline on how to assess the aptitude). Individual methods are by nature always subjective.

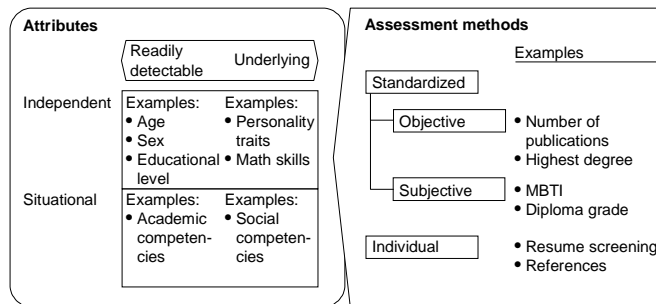


Figure 5: Classification of attributes and assessment methods

Looking at the above differentiation of personal attributes, it is obvious that the perception of subjectively assessed attributes might play a role in the partner matching process. An employer's different perception of diplomas granted from different universities is an example of this phenomenon. Hence, only attributes assessed by standardized objective methods can be considered absolute.

Assessing a candidate's aptitude for a job position by looking at his resume can be considered as an individual and, thus, subjective assessment method. The individual coherence between different recruiters' preferences and the value of the target attribute "aptitude for the job" can be expressed in the form of a rating matrix comparable to the one described in section 1.2.2 on collaborative filtering. The entries of the matrix are the assessed values in this case. As different recruiters might try to match the same candidates for different job positions, the actual target attribute can be different for each assessor (e.g., "aptitude for job A" vs. "aptitude for job B"), although they use the same method ("resume screening") and the values are of the same domain (e.g., {"qualified", "not qualified"}). Hence, the assessments are comparable.

In the following section we first introduce a standard latent aspect model for pure collaborative filtering in order to apply it to the partner matching problem or more specifically to the personnel selection problem. We then extend the model to a hybrid approach building on partner attributes as just described.

4.2 Model specification and parameters

In Chapter 3, different approaches to automated recommendations have been introduced. We now specifically look at a probabilistic approach for collaborative

filtering and hybrid systems called the latent aspect model. The main idea behind this approach is to model individual preferences as a convex combination of preference factors [HoPu99, p. 688]. In this manner, different individual reasons or aspects lying behind a rating value can be modeled. The combination of different aspects into a preference profile is a very flexible approach which has shown good results for automated recommendation [HoPu99, p. 693; Pop⁺01]. Another application of latent aspect models is probabilistic latent semantic analysis [Hofm99].

The basic principles of the use of latent aspect models for collaborative filtering are depicted by Hofmann and Puzicha [HoPu99]. Variations can be found, e.g., in Popescul et al. and Schein et al. [Pop⁺01; Sche⁺02]. Using a basic model for pure collaborative filtering, we look at observations of user/object pairs (x,y) with $x \in X = \{x_1, \dots, x_n\}$ and $y \in Y = \{y_1, \dots, y_m\}$ where X is a set of users and Y is a set of objects. For the basic model, observations are just co-occurrences of users and objects representing events like "user x has accessed object y ", i.e., preference values are not considered. The aspect model can then be represented as a *latent variable model* using a latent aspect variable $z \in Z = \{z_1, \dots, z_k\}$ which is associated with each observation (x,y) , assuming that x and y are independent conditioned on z . The model can then be depicted as shown in Figure 6(a) and the probability model can be written as:

$$P(x, y) = P(x) \sum_{z \in Z} P(z|x) \cdot P(y|z) \quad (5)$$

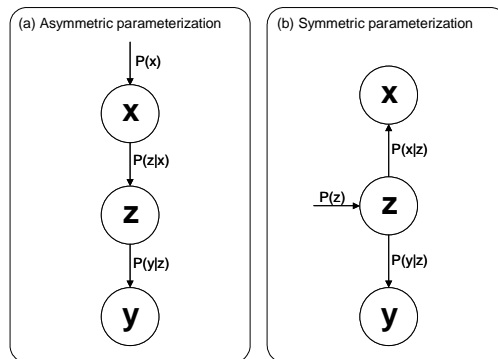


Figure 6: Basic latent aspect model

While this representation is intuitively very appealing, a symmetric formulation is used for estimating the parameters (see Figure 6(b)). To re-parameterize the model, the identity $P(z) \cdot P(x|z) = P(x, z) = P(x) \cdot P(z|x)$ is used leading to the following formulation:

$$P(x, y) = \sum_{z \in Z} P(z) \cdot P(x|z) \cdot P(y|z) \quad (6)$$

While the basic model described above shows the basic principle of this approach, it does not consider preference values or in our case assessed values of attributes. Therefore, we use an extended version as proposed by Hofmann and Puzicha [HoPu99, pp. 689]. They introduce an additional variable v into the model representing the rating value. The model can now be specified in different ways depending on the structure of the rating or assessment process. We choose a variant in which the value depends indirectly on x and directly on y . The selection of the rating object y is not part of the assessment process and, hence, is independent of x (see Figure 7). In our case, the assessor and the specifics of the target attribute to be assessed (e.g., the job position) are both represented by variable x .

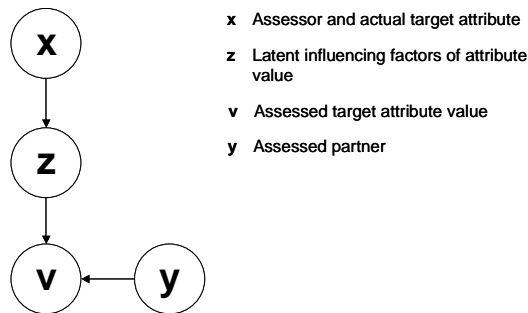


Figure 7: Basic structure of an aspect model for assessing attributes

With a set of observed values v for an attribute assessed by x and assigned to y , we are able to estimate the model parameters using the Expectation Maximization (EM) algorithm [Dem⁺77]. Very good introductions to the EM algorithm for latent aspect models can be found in Hofmann and Puzicha and Popescu et al. [HoPu99, p. 689; Pop⁺01, p. 439]. Figure 8 shows possible parameter estimates in a scenario with 3 assessors, 3 candidates, 2 latent variable values, and 2 possible target attribute values. Using the parameters, we can derive the predictions.

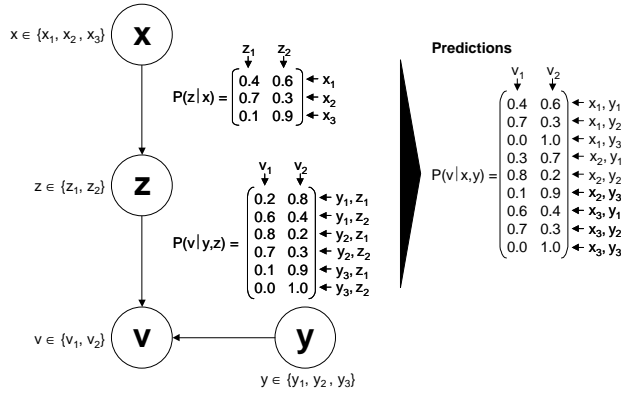


Figure 8: Numeric example of model parameters (asymmetric parameterization)

In another example we applied the model of Figure 7 on a dataset with 5 assessors and 10 partners of which a certain attribute has to be assessed. We chose a dimensionality of 3 for the latent aspect variable. The EM algorithm then estimated the model parameters $P(z)$, $P(x|z)$, and $P(v|y, z)$, which we used to predict the rating values.

Original value matrix

	x001	x002	x003	x004	x005
y001	0,80/0,20	0,00/0,00	0,00/1,00	0,20/0,80	0,00/1,00
y002	0,50/0,50	0,50/0,50	1,00/0,00	1,00/0,00	0,00/0,00
y003	1,00/0,00	1,00/0,00	0,50/0,50	0,40/0,60	0,10/0,90
y004	0,50/0,50	0,50/0,50	0,50/0,50	0,90/0,10	0,90/0,10
y005	0,30/0,70	0,30/0,70	0,40/0,60	0,50/0,50	0,50/0,50
y006	1,00/0,00	1,00/0,00	0,80/0,20	1,00/0,00	1,00/0,00
y007	0,00/1,00	0,00/1,00	0,40/0,60	1,00/0,00	1,00/0,00
y008	0,60/0,40	0,60/0,40	0,00/0,00	1,00/0,00	1,00/0,00
y009	1,00/0,00	1,00/0,00	1,00/0,00	1,00/0,00	1,00/0,00
y010	0,00/0,00	0,00/0,00	0,00/0,00	0,00/0,00	0,00/0,00

Predicted value matrix

	x001	x002	x003	x004	x005
y001	0,76/0,24	0,71/0,29	0,00/1,00	0,09/0,91	0,00/1,00
y002	0,49/0,51	0,53/0,47	1,00/0,00	0,94/0,06	1,00/0,00
y003	1,00/0,00	1,00/0,00	0,62/0,38	0,25/0,75	0,17/0,83
y004	0,52/0,48	0,49/0,51	0,49/0,51	0,89/0,11	0,91/0,09
y005	0,30/0,70	0,30/0,70	0,40/0,60	0,49/0,51	0,51/0,49
y006	0,95/0,05	0,94/0,06	0,90/0,10	1,00/0,00	1,00/0,00
y007	0,00/1,00	0,00/1,00	0,46/0,54	0,91/0,09	1,00/0,00
y008	0,64/0,36	0,60/0,40	0,48/0,52	0,99/0,01	1,00/0,00
y009	1,00/0,00	1,00/0,00	1,00/0,00	1,00/0,00	1,00/0,00
y010	0,00/0,00	0,00/0,00	0,00/0,00	0,00/0,00	0,00/0,00

Figure 9: Original and predicted value matrix

Figure 9 shows the original and the predicted values with the columns being the assessors and the lines being the assessed partners. Each column has two parts

which either sum up to 1 or both equal 0. The left part represents the relative frequency the respective assessor has assessed the partner of the according line as qualified, while the right part represents the relative frequency the assessor has assessed the partner as not qualified. In a realistic personnel selection case, the frequency is either 1 or 0 since each candidate is only assessed once. Note that we defined the domain of the assessed value v being {"qualified", "not qualified"}. In cases where both frequencies are 0, the assessor has not assessed the respective partner.

In the original value matrix in Figure 9, we can see that the assessors 1 and 2 (the first two columns) provided similar assessments except for the fact that assessor 2 has not assessed partner 1. Both assessors obviously have a very similar preference structure. The latent aspect model is not able to capture the very small difference between the two assessors since the number of values of the latent variable z is smaller than the number of assessors. This leads to the result we can see in the lower part of Figure 9 showing the predicted value matrix. We can see that both, assessor 1 and assessor 2, have a very strong tendency towards considering partner 1 as qualified.

This simple example shows the basic function of the latent aspect model of Figure 7. While this might work very well when recommending movies or books that have been rated positively or negatively by the users, a realistic partner matching scenario such as candidate selection in recruitment raises some specific challenges. First of all, the pool of potential candidates is frequently changing, i.e., new candidates enter the pool and candidates that have been matched leave the pool. Therefore, a large number of potential candidates is only rated by a few assessors, namely those that are looking for candidates just in the same time period the candidates are looking for jobs. Additionally, the same assessors trying to match different job positions have to be treated as different variables x in the model, since a different preference structure might be applied. This might also lead to a very large dimensionality of the value matrix.

We approach this problem by adapting the pure collaborative filtering model of Figure 7 using the idea of a hybrid method. We treat the assessments of partners as assessments of the partners' attributes. Hence, we substitute the model variable y of Figure 7 representing the candidates in our example with the candidates' "features", namely their attributes. As opposed to hybrid approaches to recommending documents where the documents' features, the words, are easily observable, we have to consider the complications arising with the different determination and individual perception of human attributes as described in section 4.1.

The structure of the hybrid model (see Figure 10) is similar to the pure collaborative filtering model. At a first glance, the partner to be assessed is just substituted by a single partner attribute and, thus, target attribute values assigned to a certain partner are actually assigned to all his attributes. However, variable a stands for a quadruple consisting of attribute name, assessment method, assessed attribute

value and assessor (e.g., a =(*"mathematical skills"*, *"diploma grade"*, *"1.0"*, *"University of Frankfurt"*)) taking the specifics of human attributes into account. Therefore, a positive influence of a diploma granted from University of Frankfurt on the target attribute value does not necessarily mean that a diploma granted from another university has the same positive influence.

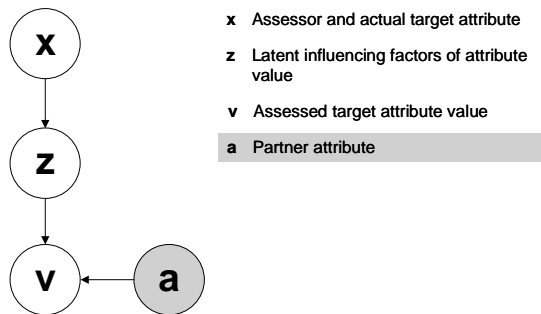


Figure 10: Hybrid assessment model

Another effect of the hybrid approach is that the size of the original value matrix is limited to the number of attribute quadruples (although the total number can be very large due to the heterogeneity of human attributes) while the number of potential partners in the pure collaborative filtering model is constantly increasing with new partners entering the system.

Although the hybrid approach partially addresses the sparsity problem of the original value matrix which is typical for recommender systems based on collaborative filtering there are still too few matrix entries to ensure good predictions. The following section presents an estimation and matching process that incorporates a possible way to decrease the sparsity of the original matrix.

4.3 Estimation and prediction process

In order to estimate a robust model, the sparsity problem described above has to be further reduced. Therefore, we propose a process in section 4.3.1 that complements the original value matrix using linear interpolation. Secondly, we need to define a method how to predict target attribute values for profiles from the predicted value matrix that assigns values to single attributes instead of complete profiles (see section 4.3.2).

4.3.1 Estimation of model parameters

The model estimation process (see Figure 11) requires a set of existing partner profiles that are already assessed. This results into a value matrix R that assigns assessed values to profiles:

$$R = (r_{x,y,v}) \text{ with } r_{x,y,v} = \begin{cases} 1 & \text{if assessor } x \text{ assesses the target attribute} \\ & \text{of partner } y \text{ with value } v \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

In order to derive the value matrix on the attribute level, the attributes assigned to each partner have to be identified. These are then used to derive the original value matrix R' on attribute level:

$$R' = (r'_{x,a,v}) \text{ with } r'_{x,a,v} = \begin{cases} 1 & \text{if } r_{x,y,v} = 1 \text{ and } a \text{ is an attribute of partner } y \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

As many attributes are assigned to several profiles and, hence, might be observed several times with different values v , the entries of matrix R' are actually not either 0 or 1 but take values in the interval $[0;1]$ according to the relative frequency of value v being assigned to attribute a by assessor x .




Figure 11: Process for estimating model parameters

In a realistic scenario with many attributes, the matrix R' is still rather sparse. However, the structure of many attributes (e.g., ordinal scale) allows the inference of additional entries in the value matrix using intuitive assumptions. For example, if the target attribute value v ="qualified" is assigned to the attribute "very good math skills" in 80% of the occurrences of this attribute and to the attribute "intermediate math skills" in 40% of the occurrences of this attribute by a certain assessor, we may hypothesize an assignment of this value to the attribute "good math skills" with a frequency of 60% if the assessor has not actually assessed this attribute. Hence, we apply simple linear interpolation as a first attempt to overcome the sparsity problem. With the complemented original value matrix, the model parameters are finally estimated using a standard EM algorithm.

4.3.2 Prediction

When looking at a new partner (or a new potential candidate in our personnel selection example) that has not been assessed by a certain assessor, we are now able to predict an assessment even if this partner has not been assessed by any other assessor in the system. First, predictions of the target attribute value can be assigned to the individual attributes (e.g., those derived from a resume) of the partner to be assessed. Assuming a robust estimation of the parameters, attributes a

that are important to determine the target attribute should have a very high probability $P(v|a, x)$ for a value v given the assessor x . Attributes a that are less important will have a more or less equal probability distribution over all possible target attribute values. The final prediction is made by averaging over all available partner attributes with A_y being the set of attributes of partner y and $n_y = |A_y|$:

$$P(v|y, x) = \frac{1}{n_y} \sum_{a \in A_y} P(v|a, x) \quad (9)$$

5 Conclusions and further research

Based on the question on how methods from information systems research are capable of improving the matching quality of partner matching platforms, we presented a possible approach using automated recommendation methods and building on findings from team building and work psychology. Examples as well as a survey on modern internet-based recruitment practices that we conducted with the largest 1,000 companies in Germany showed a lack of intelligent support to search for potential partners. Usually, only Boolean search mechanisms are employed that do not embrace the underlying aspects that might lead to a good match.

We use a probabilistic approach applying a latent aspect model which is capable of learning underlying matching criteria. The model automatically predicts candidates with a high probability to be qualified for a specific job. It also considers the specific characteristics of human attributes and their importance in matching processes. Therefore, this model-based method potentially represents a significant improvement compared to matching methods used on existing internet platforms.

The main challenge in implementing this concept lies in robustly estimating the model parameters as the training data might not be sufficient. Therefore, we proposed a possible way to generate additional training data. In our future work we will consider additional methods to further improve matching quality. We will specifically look into relationships among assessors and into relational attributes such as trust to be used as an additional source of information for preference similarity.

We gratefully acknowledge the support of the European Union under the Fifth Framework Programme Information Society Technologies (contract number: IST-2000-28295).

References

- [Albe98] Albert, G.: Betriebliche Personalwirtschaft. Kiehl: Ludwigshafen a.R. 1998.
- [BaSh97] Balabanović, M.; Shoham, Y.: Fab: content-based, collaborative recommendation. *Communications of the ACM* 40 (3), 1997: pp. 66-72.
- [BeCr92] Belkin, N.; Croft, W.B.: Information filtering and information retrieval: two sides of the same coin? *Communications of the ACM* 35 (12), 1992: pp. 29-38.
- [Bre⁺98] Breese, J. S.; Heckerman, D.; Kadie, K.: Empirical analysis of predictive algorithms for collaborative filtering. *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence (UAI)*, 1998: pp. 43-52.
- [Bung90] Bungard, W.: Team- und Kooperationsfähigkeit. In: Sarges, W. *Management-Diagnostik*. Hogrefe Verlag für Psychologie: Göttingen. 1990: pp. 315-325.
- [Dem⁺77] Dempster, A.P.; Laird, N.M.; Rubin, D.B. Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Statist. Soc. B* 39, 1977: pp. 1-38.
- [FoDu92] Foltz, P.W.; Dumais, S.T.: Personalized information delivery: an analysis of information filtering methods. *Communications of the ACM* 35 (12), 1992: pp. 51-60.
- [GaWa01] Galla, M.; Wagner, M.: Partnersuche im E-Business. *Proceedings Workshop Gemeinschaften in Neuen Medien (GeNeMe)*, 2001: pp. 19-40.
- [Good⁺99] Good, N.; Schafer, J.; Konstan, J.; Borchers, A.; Sarwar, B.; Herlocker, J.; Riedl, J.: Combining collaborative filtering with personal agents for better recommendations. *Proceedings of the 16th National Conference on Artificial Intelligence*, 1999: pp. 439-446.
- [Guzz96] Guzzo, R.A. Fundamental considerations about work groups. In: West, M.A. (Ed.) *Handbook of workgroup psychology*. John Wiley & Sons: Chichester, UK. 1996: pp. 3-21.
- [Hofm99] Hofmann, T.: Probabilistic latent semantic analysis. *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence (UAI)*, 1999: pp. 289-296.
- [HoPu99] Hofmann, T.; Puzicha, J.: Latent class models for collaborative filtering. *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, 1999: pp. 688-693.
- [Jack96] Jackson, S.E.: The consequences of diversity in multidisciplinary work teams. In: West, M.A. (Ed.) *Handbook of workgroup psychology*. John Wiley & Sons: Chichester, UK. 1996: pp. 53-76.
- [Kön⁺03] König, W.; Wendt, O.; Färber, F.; Keim, T.; Weitzel T.; von Westarp, F.: *Recruiting Trends 2003 - Eine empirische Untersuchung der Top-1.000-Unternehmen in Deutschland*. Research report, University of Frankfurt, 2003.
- [LaPi99] Lang, A.; Pigneur Y.: Digital trade of human competencies. *Proceedings of the 32nd Annual Hawaii International Conference on System Sciences (HICSS)*, 1999: pp. 165-173.

- [Lin⁺01] Lindgren, R.; Stenmark, D.; Bergquist, M.; Ljungberg, J.: Rethinking Competence Systems for Innovative Organizations. Proceedings of the 9th European Conference on Information Systems, 2001: pp. 775-786.
- [MaLa98] Malone, T.W.; Laubacher, R.J.: The dawn of the e-lance economy. *Harvard Business Review* 76 (5), 1998: pp. 144-152.
- [Mel⁺02] Melville, P.; Mooney, R.J.; Nagarajan, R.: Content-boosted collaborative filtering for improved recommendations. Proceedings of the 18th National Conference on Artificial Intelligence, 2002.: pp. 187-192.
- [Pop⁺01] Popescul, A.; Ungar, L.H.; Pennock, D.M.; Lawrence, S.: Probabilistic models for unified collaborative and content-based recommendation in sparse-data environments. Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence (UAI), 2001: pp. 437-444.
- [Res⁺94] Resnick, P.; Iacovou, N.; Suchak, M.; Bergstrom, P.; Riedl, J.: Grouplens: an open architecture for collaborative filtering of netnews. Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW), 1994: pp. 175-186.
- [ReVa97] Resnick, P.; Varian, H.R.: Recommender systems. *Communications of the ACM* 40 (3), 1997: pp. 56-58.
- [Sar⁺00] Sarwar, B.; Karypis, G.; Konstan, J.; Riedl, J.: Analysis of recommendation algorithms for e-commerce. Proceedings of the ACM Conference on Electronic Commerce (EC), 2000: pp. 158-167.
- [Sar⁺01] Sarwar, B.; Karypis, G.; Konstan, J.; Riedl, J.: Item-based collaborative filtering recommender algorithms. Proceedings of the 10th International World Wide Web Conference, 2001: pp. 285-295.
- [Sche⁺02] Schein, A.; Popescul, A.; Ungar, L.H.; Pennock, D.M.: Methods and metrics for cold-start recommendations. Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR), 2002: pp. 253-260.
- [Schn95] Schneider, B.: *Personalbeschaffung*. Peter Lang Europäischer Verlag der Wissenschaften: Frankfurt a.M. 1995.
- [UnFo98] Ungar, L.H.; Foster, D.P.: Clustering methods for collaborative filtering. Workshop on Recommendation Systems at the 15th National Conference on Artificial Intelligence, 1998.
- [West94] West, M.A.: *Effective teamwork*. BPS Books: Leicester, UK. 1994.
- [WeWo00] Wells, N.; Wolfers, J.: Finance with a personalized touch. *Communications of the ACM* 43 (8), 2000: pp. 31-34.