

A Recommender System For Investing In Early-Stage Enterprises

Johannes Luef*
TU Wien, Austria

Dimitris Sacharidis
TU Wien, Austria

Christian Ohrfandl*
TU Wien, Austria

Hannes Werthner
TU Wien, Austria

ABSTRACT

Selecting the appropriate enterprise to invest in can become a difficult task for prospective investors, particularly in the case of startups where limited information is available. In this work, we design, implement, and evaluate a system that makes recommendations for investing in early-stage enterprises. We first perform a qualitative and quantitative study involving prominent investors to explore their decision-making process and set out the requirements for a recommender system. Then, we consider various recommendation approaches that best convey the results of our requirements analysis. To evaluate the different methods, we simulate an offline experiment, placing real investors in a hypothetical investment decision scenario.

CCS CONCEPTS

• Information systems → Recommender systems.

1 INTRODUCTION

A sustained positive birth rate of enterprises is considered a key indicator of economic growth [14]. According to OECD and Eurostat, the global financial crisis of October 2008 resulted in the decline of new firm registrations and an increase of bankruptcies, even in countries having a high level of financial development (e.g., Germany, France) [8, 24]. However, the past few years have seen a healthy positive enterprise birth rate [9, 23]. Consequently, to support the creation of new enterprises, *startup ecosystems*, loosely defined as entrepreneurs, investors, and related organizations and activities (e.g., accelerators, incubators, coworking spaces, events, networks), are growing. In particular, investors such as *Business Angels* or *Venture Capital Funds*, who provide capital particularly needed in the early stages of the company formation and beyond, play a key role in growth. However, as a result of a continuously increasing number of enterprises entering the market, potential investors face the problem of *information overload* in their investment decisions. Therefore, the need for information filtering techniques based on computational *recommendation systems* emerges.

This paper discusses the design, implementation, and evaluation of a recommender system that assists investors looking to fund enterprises in their early stage. As this application domain is relatively new for a recommender system [31], we first need to understand the decision processes of investors, and set to investigate the first research question: *How can decision-making requirements and investor behavior be quantified?* This includes determining the investment criteria investors apply and the information on early-stage enterprises that is of interest. To answer this question, we perform a detailed requirements analysis that includes literature review of investment and valuation strategies, lengthy interviews with prominent investors across different fields, and a quantitative study of the requirements and expectations from a system supporting investment decisions. The outcome is a list of five requirements that a system should possess.

Following the results of our study, we explore the second research question: *Which recommendation algorithms are appropriate in the domain of enterprise investment?* To answer this, we need to identify which methods fit the decision process of investors, taking into account the data about enterprises that can be procured. Given the complexity of the investment decision making process and the fact that no single recommendation method can encapsulate all its nuances (as expressed in the five requirements we identifies), we chose to design various recommenders: a knowledge-based recommender, which seeks to model investment profiles, a content-based approach, as well as non-domain specific collaborative filtering, and a hybrid approach. As our requirements analysis showed that investors often have an *inner circle of trust*, i.e., they tend to follow investments of certain reputed investors, we also explore collaborative filtering with a trust mechanism.

Given this plethora of recommendation approaches, the third research question arises naturally: *Which recommendation algorithm is most effective in the domain of enterprise investment?* This prerequisites an evaluation strategy for the recommenders. As the developed system is in a prototype stage, and because collecting actual investment decisions would require a long period of time, we decided to simulate an offline evaluation strategy. Specifically, we asked several investors to fill out their profiles (all data required by the system to generate recommendations), and indicate their likelihood of investing in 10 predefined early-stage enterprises. We then split the collected data into a train and a test set, and measure the ranking accuracy of the different ranking mechanisms. Our analysis finds that, surprisingly, the knowledge-based recommender is less effective than simple collaborative filtering, and that trust-based recommendations perform the best in this domain.

The remainder of this paper is structured as follows. Section 2 introduces the necessary background in the domain of investments in early-stage enterprises. Then, Section 3 presents our requirements

*Both authors contributed equally to the paper.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SAC '20, March 30–April 3, 2020, Brno, Czech Republic

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6866-7/20/03...\$15.00

<https://doi.org/10.1145/3341105.3375767>

analysis corresponding to the first research question. Section 4 builds upon this analysis and discusses the recommendation approaches taken so as to answer the second question. Section 5 address the last research question and presents the results of our evaluation. Section 6 concludes this work.

2 BACKGROUND

Early-Stage Enterprises. An *enterprise*—company, business, or venture—is an economic entity serving as the main site of the entrepreneur’s business operations, who constitutes a natural or legal entity. The aim of an enterprise is to increase its *value* [17, 20], via value-creating processes [17] that transform certain input factors into higher-valued output factors [15]. The life cycle of an enterprise spans from the very first moment of innovating a business idea to the foundation of an enterprise until the so called *exit*, the disposal of the enterprise. The cycle consists of three financing stages, the *early-stage*, *expansion*, and *late-stage* [1]. ; see Figure ?? . The early-stage itself consists of the *seed* stage, where entrepreneurs engage in the definition of the product idea and market analysis, and the *start-up* stage, where the priority lies on the foundation of the enterprise and the development of product maturity. Early-stage enterprises commonly find themselves in the red, and must cross the *valley of death* [19] to reach profitability in the expansion stage. For this, entrepreneurs are reliant upon *funding*, which is predominantly based on the owner’s equity capital, public grants, and venture capital.

Funding. Funding enables early-stage enterprises to generate rapid growth and is commonly classified as internal and external. Internal funding—colloquially known as the *four Fs*, an abbreviation for *founders, family, friends, fools*—mark the earliest type of funding, which however barely suffices for enterprise foundation [27]. External funding includes stakeholders such as *public funding agencies, investors*, as well as *business incubators/accelerators*. Public Funding Agencies, e.g., government, state-affiliated organizations or universities, ordinarily allocate grants, which may include monetary funding as well as intangible assets (e.g., training in entrepreneurship and economics) legally contracted under certain conditions [5]. Investors, such as *business angels (BAs)* and *venture capital firms (VCs)*, fund early-stage enterprises by investing monetary assets, referred to as *external equity capital* or *private equity* because their distribution is not organized via an official capital market (e.g., the stock market) [1, 27]. BAs are private persons, entrepreneurs themselves, that invest personal assets in privately held companies. A VC provides intangible assets in addition to capital, after a formal valuation process [27]. This includes the initial general review, when the VC checks the management team and the fit of the industry sector; initial contact, when the business plan is carefully inspected; due diligence, when an intentional agreement is drawn and the enterprise is valued at a high level of detail; and negotiation, when the VC agrees to invest. The *worth* of the enterprise is determined by the *valuation* process (in the due diligence phase).

Valuation. The process of calculating the value of an enterprise based on its future benefit to its owners is called *valuation* [20]. Early-stage enterprises, in particular, expose certain requirements to venture valuation methods. The most important is *future mindedness*, i.e., to not solely rely on historic business activity records,

which often such enterprises do not possess [1]. Valuation methods are broadly classified into situation specific and unspecific, where in the former the method depends on the specific reason for the valuation. The applicable valuation methods for early-stage enterprises, featuring future mindedness, belong to the former category, and are further subdivided into total valuation methods and rules of thumb. Especially when no historic business records exist, rules of thumb techniques, such as the Scorecard [25] and the Berkus methods [2, 4], are typically employed by business angels.

Recommender Systems. The function of *recommender systems* is the suggestion of *items* (e.g., products, news, movies) the *user* might prefer over others [26]. In general, recommendations are compiled based on *users’ profile data*, which include demographics, context, and preferences provided *directly* by the users or *indirectly* by analyzing their behavior, and *items’ content*, which describe the items. Several broad classes of techniques exist that include: *collaborative filtering* [22] where patterns in the historical interactions of users with items are exploited; *content-based* [6] recommenders that analyze the content of the items historically consumed by a user; *knowledge-based* [10] recommenders that rely on a domain-specific elicitation of users’ preferences and an appropriate specification of items’ content.

While recommenders are particularly predominant in the e-commerce domain, they have also been employed in the finance sector [31]. The most relevant systems are those for *portfolio selection*, where knowledge-based techniques are used [11, 21], and *stock recommendation*, where knowledge-based [30], content-based [28], and social network-based collaborative filtering [29] approaches are used. To the best of our knowledge, recommender systems have not been employed in the domain of investment in or funding of (early-stage) enterprises.

3 REQUIREMENTS ANALYSIS

To design a recommender system for investments in early-stage enterprise, we first need to understand the needs and expectations of the end users, the investors. This section presents the results of our requirements analysis that includes qualitative and quantitative research. More details on the study can be found in [16].

3.1 Qualitative Study

In the first phase, we conducted *qualitative expert interviews* to collect expert knowledge on decision-making criteria and enterprise valuation methods utilized by investors in the domain of early-stage enterprise investment. The results of this study, together with our literature review, inform the quantitative study.

We conducted a *strongly structured interview* [3], where a questionnaire with 39 various-typed questions is compiled beforehand following literature review and a pretest. The goal is to learn about the (i) investment decision criteria, and (ii) expectations from a recommendation platform. The interview was designed to run for 90–120 minutes, and 28 prominent investors from the DACH region (Germany, Austria, Switzerland) were invited. In the end, we conducted in-person interviews with six investors — the main reasons for the low response rate was reluctance to participate (in spite of long networking efforts) and simply lack of time for such an extended interview.

Decision Criteria for Investment in Early-Stage Enterprises.

Early-stage enterprises commonly introduce themselves to investors by showcasing a business plan and conducting a presentation of the business—which is also referred to as *pitch*. The most important aspects of the business plan is outlining the team, product idea, value proposition, and needed resources (equity capital, knowledge, etc.). Personal opinions on the meaningfulness of the *business plan* are scattered among participants. Whereas some think that business plans are not important any more, others argue that the business plan is more trustworthy than pitching slides.

Naturally, key roles in investment decisions play the *value proposition* and the *maturity level* of the product, as well as the *industry sector*, and the market’s *geographic location*, *potential*, and *competitiveness*. An important characteristic that investors tend to focus on is the experience and chemistry of the founders/management *team* behind an early-stage enterprise. Another aspect is the investor’s personal *risk awareness*: participants mention that they try to build portfolios containing multiple early-stage enterprises to reduce the risk of losing private equity capital.

A noteworthy point raised was the *trust relationships* between investors. All participants agreed that they consider other investors’ opinions when making their own decisions. This manifests in several ways, such as consulting within an inner circle of trust, and by collective investments (driven by a lead investor, or by the distribution of risk among several investors).

Expectations from a Recommendation Platform. In general, participants were a bit sceptical about the idea of automatic recommendations of early-stage enterprises to investors. Nonetheless, answers about general requirements to such a system were rather concrete: ability to describe the profile of the investor in great detail; ability to rate teams and, subsequently, rank early-stage enterprises depending on the responsible team; calculation of metrics for valuation purposes; analysis of the competitive landscape (other similar enterprises in the given industry sector).

A major benefit of the recommendation system would be to produce suitable *matches* between an investor’s profile and early-stage enterprises. According to the participants, the matching should consider the enterprise’s *industry sector*, life cycle *stage*, *needed capital*, or *similarity* to other early-stage enterprises. Another expectation is the ability to get recommendations from other investors.

The participants also stated the desire for the system to present information and various statistics about all early-stage enterprises with specified market locations. Collecting the appropriate data, however, is recognized as the greatest challenge.

3.2 Quantitative Study

In the second phase of our requirements analysis, based on the aforementioned findings, we conducted a quantitative study in order to guide us in the design of the recommender system. Specifically, we crafted a questionnaire containing seven questions, six of which where to be answered on a Likert 1–5 scale, and one was dichotomous (yes/no). Overall, we contacted 47 investors and got responses from 25. We next present the results of the three most important questions; for a complete discussion refer to [16].

Upon which criteria do you decide whether to invest in an early-stage enterprise? Participants were asked to rate 10 criteria




Figure 1: Investment Criteria




Figure 2: Enterprise Characteristics

options on a 5-point scale, ranging from unimportant to important. Figure 1 plots the distribution of ratings per each option. To assess the practicability of the results, we perform the Wilcoxon Signed Rank test to investigate whether the opinion that an option is Rather Important or Important is statistically significant (at the 97.5% level) compared to the null hypothesis of it being Neutral. The test is positive for the following 5 options: Recommendations (e.g. by investors); Industry sector; Experience of entrepreneur(s); Return on investment vs. risk; Market research.

Which characteristics of early-stage enterprises are important for venture valuations? In this question, participants were asked to rate 9 characteristic options on the 5-point scale; Figure 2 plots the distribution of ratings. Among the options, the following 6 were found to be statistically significant (based on the Wilcoxon Signed Rank test at the 97.5% level): Opportunity (market situation, revenue in 5 years); Maturity level of the product idea; Customer acceptance of the product idea; Industry structure (market entry barriers, market growth); Competition; Experience of the founder team.

According to which criteria should recommendations be generated? Participants were asked to rate 6 criteria options on the 5-point scale. Figure 4 plots the distributions of ratings. Among them, the following 2 criteria are statistically significant: Recommendations based on your former investment decisions; Recommendations based on an investor’s profile.

Which additional functionality should a recommendation platform provide to investors? As before, participants were asked

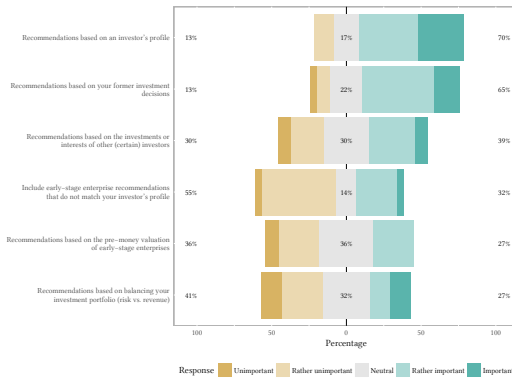


Figure 3: Recommendation Criteria

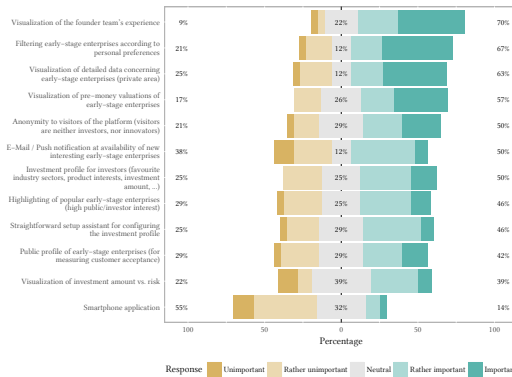


Figure 4: Recommendation Functionality

to rate 12 functionality options on the 5-point scale; Figure 4 plots the distributions of ratings. Among all options, 4 are statistically significant: Visualization of detailed data concerning early-stage enterprises (private area); Filtering early-stage enterprises according to personal preferences; Visualization of pre-money valuations of early-stage enterprises; Visualization of the founder team's experience.

3.3 Outcomes

In what follows, we discuss five specific requirements that we identify from our qualitative and quantitative studies.

R1. Valuations of early-stage enterprises. Qualitative expert interviews indicate that classical venture valuation methods are largely irrelevant. Instead, investors tend to primarily use the Scorecard method, while also employ techniques similar to those of the Berkus method. Moreover, “visualization of pre-money valuations of early-stage enterprises” was one of the statistically significant desired functionalities discovered in the quantitative study. *Therefore, the system should support the aforementioned valuation methods and also allow for the retrieval of the valuations of an enterprise.*

R2. Information on early-stage enterprises and the target market. Both qualitative and quantitative studies highlight the importance of certain characteristics of the enterprises to support investment decisions. These include information about (i) the team, namely its size and past experience of their members, (ii) the product, i.e., its maturity level, whether a prototype and/or potential customers exist. Moreover, an important factor guiding investment

decisions is the state of the target market, which includes information indicating whether it is still in its growth phase and not saturated, and about the level of competition. *Therefore, for each enterprise the system should present information about its team, product, and market sector.*

R3. Investor profiles. One of the statistically significant desired recommendation functionalities was “filtering early-stage enterprises according to personal preferences”. In the qualitative study, it was made clear that in their investment decisions investors employ multiple criteria, which they often weigh differently. Decisions are essentially made on a balanced combination of the aforementioned enterprise characteristics and the market situation. *Therefore, the system should allow investors to build a detailed profile that captures their interests and their relative importance, and support the retrieval of best matching enterprises.*

R4. Historical investment decisions. Another statistically significant desired functionality was to make “recommendations based on an investor's former investment decisions”. Moreover, in the qualitative study, the investors emphasized how one's experience affects one's decisions. *Therefore, the system should draw upon historical investment behavior when making recommendations.*

R5. Trust among investors. Although “recommendations based on the investments or interests of other (certain) investors” was not among the statistically significant desired functionalities, it acquired 70% non-negative opinions (Neutral, Rather Important, Important). Moreover, the qualitative study showed that “trust relationships” among investors was broadly considered an important decision criterion, and that *circles of trust* often exist. *Therefore, the system should provide the option to consider the opinions of other investors when making recommendations.*

4 RECOMMENDER SYSTEM

This section presents the recommender system designed as per the requirements outlined in Section 3. We first present the main entities of the system, how they interact, and show that they comply with requirements R1–R4. Then we discuss various recommendation approaches and how they support requirements R3–R5.

4.1 Data Entities

Enterprise (Item). Early-stage enterprises are to be recommended and thus correspond to items in the recommender systems parlance.

Enterprise Description. Enterprises are described by their *content*, certain attributes that were found to be highly important in our requirements analysis. Therefore to comply with requirement R2, the content includes the name, location and establishment date of the enterprise, life cycle stage of enterprise, description of the product (keyword-based), a set of investment options (including offer and cost), a set of valuations (including amount and method), the market sector (among a predefined set), team members and roles.

Entrepreneur. An enterprise is managed by a team of entrepreneurs. This user role is also tasked with maintaining the entry in the system.

Valuator. Another user role is that of a valuator that creates valuations for enterprises, using either the Scorecard or the Berkus method. This complies with requirement R1.

Investor (User). Investors are the entities receiving recommendations, and are thus also called users. The system allows users to build their *investor profile*, but also collects user feedback in terms of *user-item interactions*.

Investor Profile. To comply with requirement R3, the investor profile contains the following information: Market Sectors (set of predefined options); Product Description (set of predefined keywords); Investment Amount (monetary range); Investment Stake (percentage range); Valuation (monetary range); Valuation Type (among predefined types); Life Cycle (among predefined options); Establishment Date (range); Enterprise Location (predefined set of geographic regions); Team Size. The investor can also set the *relative importance* (specified as percentages) of these attributes in her investment decision process. For example, an investor who only looks to invest within a specific market sector at a specific geographic region, might set the importance of these two attributes to 100% and of all others to 0%.

User-Item Interactions. Requirement R4 mandates that past investor behavior is an important source of information. As investors interact with the system, they provide *feedback* that a recommender can build upon. We collect three types of signals, presented in increasing strength. An investor may *click* on (view) an enterprise, seeing all information available in the system. An investor may express interest on an enterprise by *liking* it. Of course, an investor may also specify that she has already *invested* in an enterprise.

4.2 Recommendation Algorithms

Our requirements analysis found that there are different ways in which investors decide to choose an enterprise to invest in. For example, the three last requirements R3–R5 essentially correspond to distinct types of recommendations that the system should support. Therefore, instead of proposing a single recommendation list, we choose to go with a Netflix-style presentation of recommendations, where there are multiple recommendation lists (rows), each produced by a different recommendation engine. In what follows, we describe the five engines employed.

Knowledge-Based (KB). This recommendation algorithm is based solely on the investor profile and the enterprise description alone, and thus complies with requirement R3. The algorithm essentially matches enterprises (based on their description) to investors using a combination of *filters* and *preferences*. Filters come directly from the attribute values specified in the investor profile, and act as *hard constraints* that an enterprise needs to satisfy in order to be recommendable/relevant. Preferences are specified on the same attributes, but act as *soft constraints* that induce an ordering on the attribute domain; e.g., investors might prefer enterprises with high valuation amounts.

The KB recommender operates as follows. First, the filters are applied to identify the set of relevant enterprises. Then, each specified preference creates a separate ranking of (relevant) enterprises. For a preference specified on a numerical attribute (e.g., valuation, investment amount, date) the ranking is obvious, either ascending

or descending on that attribute’s values. For a preference on an ordinal attribute (e.g., market sector, product description) the Jaccard similarity of the investment’s attribute with the preference is computed to act as a score based on which to rank descending. Thus, for each preference, we end up with an enterprise ranking. To obtain a single ranking as the output, the system performs *rank aggregation* [7]. Specifically, we choose a variant of the Borda count method [12], where each ranked list carries a weight defined by the relative importance of the attribute in the preference, as specified in the investor profile. For instance if the preference on market sectors is specified to be twice as important as that for the valuation amount, the Borda count of an enterprise based on its rank by market sector has twice the voting power than that based on rank by valuation amount. So, the ranking score of an enterprise is a weighted sum (instead of a simple sum) of the Borda counts in each list. A final note concerns the case when the filters are too restrictive, returning few or no enterprises. In this case, to offer more options to the investor, after ranking all filtered-in enterprises (based on preferences), we then rank the filtered-out enterprises.

Content-Based (CB). This recommendation algorithm is based on the enterprise descriptions and the investor’s historical behavior, and thus complies with requirement R4. Briefly, the algorithm creates a “virtual” enterprise based on the investor’s history and then matches every enterprise to it. Specifically, the enterprise description attributes Market Sectors, Product Description, and Life Cycle of each enterprise that the investor has liked, clicked, or invested in are used to define the “virtual” enterprise. Note that all said attributes are set-based. To compute a matching score for a particular enterprise, the following process is taken. For each attribute, we compute its Jaccard similarity with the corresponding attribute of the “virtual” enterprise, and then define the matching score as the average Jaccard similarity across attributes. This matching score is ultimately used to rank enterprises.

Collaborative Filtering (CF). This recommendation approach is based solely on the historical behavior of all investors, and thus satisfies requirement R5. We implement neighborhood methods, based on either user/investor or item/enterprise similarities.

In user-based CF (CF-UB), each investor is represented as a vector encoding the enterprises she has clicked on, liked, and invested in. Similarity between two investors is given by the cosine similarity of their vectors. The neighborhood of an investor contains other investors that have cosine similarity above a fixed threshold (e.g., 0.7). Based on the neighborhood, CF-UB computes the chance $s(u, i)$ of a target investor u clicking on (or liking or investing in) a target enterprise i simply as:

$$s(u, i) = \frac{\sum_{u' \in N(u)} \text{sim}(u, u') \cdot r(u', i)}{\sum_{u' \in N(u)} \text{sim}(u, u')}, \quad (1)$$

where $\text{sim}()$ is the cosine similarity among investor vectors, and $N(u)$ is the neighborhood of u . The final list of enterprises to recommend is compiled by computing for each enterprise the weighed sum of its chance of an investor clicking, liking, investing (with weights ratio 1:2:4), and then ranking by these weighted sums.

A similar approach is adopted for item-based CF (CF-IB). The main difference is that for each enterprise three separate vectors are computed, one for the investors who clicked, another for those

who liked, and another for those who invested. Each type of vector is used to create different types of neighborhoods and thus predict separate scores that act as chances of a target investor interacting with a target enterprise. As in CF-UB, a weighted sum is used to combine predictions from the different types of signals.

Social-Based (SB). This recommendation algorithm is also based on the historical behavior of all investors, and offers a different interpretation of requirement R5. Specifically it corresponds to a trust-aware strategy [18], and materializes the *inner-circle* of trust that investors mentioned in our qualitative study. For this recommendation strategy, an investor needs to state which other investors she trusts/follows. Then, SB is essentially a user-based CF method, where the neighborhood in Eq. 1 is explicitly given, rather than extracted from the similarities of investors.

Hybrid (HB). The final recommendation strategy attempts to reconcile requirements R4 and R5, making use of the historical behavior and the profiles of all investors. HB is also a variant of user-based collaborative filtering, with the key difference lying in the similarity function employed: two users are similar if they have similar investor profiles. Specifically, to assess the similarity between two investors we compare their personalized knowledge-based rankings as determined by their profiles. We compute similarity indirectly on ranked lists rather than directly on investor profiles, for two reasons. First, it is not clear how to compare profiles that may specify different filters and preferences. Second, we wish for the similarity to be conditioned on the available enterprises. If two investor profiles differ in attributes that do not matter, i.e., there are no enterprises that would rank differently, there is no need to actually treat the investors as dissimilar.

Rankings are compared with Kendall's rank correlation that essentially counts the number of concordant item pairs (a pair is concordant if the two rankings are consistent, i.e., both rank the same item higher) and subtracts the number of discordant pairs. Then, based on this similarity metric, we form the neighborhood and compute predicted interest to an enterprise as in Eq. 1.

5 EVALUATION

In this section, we conduct an evaluation of the proposed recommendation algorithms. Specifically, we seek to answer two questions: how different are the recommendations produced by the various algorithms, and which recommendation algorithm is the most effective. We chose to perform an offline experiment [13], as the system is still at a prototype stage. Given the limited availability of end-users, the investors, we chose to compile a small dataset to run experiments on by collecting all necessary information from investors using a questionnaire. Section 5.1 describes and analyzes the data collected and the evaluation methodology followed. Then, Section 5.2 presents the results of our evaluation.

5.1 Setup

Data Collection. An important step before collecting investor information, is to create a small, manageable list of candidate enterprises for investment. We considered several existing early-stage enterprises that are active in the DACH region, and select 10 of




Figure 5: Distribution of answers for Likert-type questions.

them so as to have a balanced mix of different products, life cycles, market sectors, known or unknown valuations, management team size, and establishment date. More details on the selected enterprises can be found in [16].

The next step is to collect from investors all necessary information, i.e., investor profiles and historical behavior that is required by the system in order to generate investment recommendations based on the various algorithms. We design a questionnaire with eleven questions that requires a maximum participation time of 15 minutes. In the end, 35 out of 68 investors participated. In what follows, we present the questions asked and some descriptive statistics. The first nine questions concern the investors' profile, question 10 concerns their inner circle of trust, while the last question is used to create their historical investment record.

Questions 1, 2, 4, 6, 7, 8, and 9 are on a 5-point Likert scale, ranging from Unimportant to Important. Figure 5 presents the distribution of answers for all questions on a Likert scale. Questions 3 and 5 are categorical.

Q1: A recommendation based on the innovations/products of early-stage enterprises is important to me.

Q2: A recommendation based on the current stage (life cycle) of the early-stage enterprise is important to me.

Q3: Which life cycles of an early-stage enterprise are of interest?

Q4: A recommendation based on the market sector is important to me.

Q5: Which market sectors¹ would you be interested in?

Q6: An already public acceptance / interest of the product of the early stage enterprise is important to me?

Q7: The management team of an early stage enterprise is important to me.

Q8: An already existing valuation is important to me.

Q9: An early-stage enterprise established earlier is more important to me than later ones.

Q10: Would you take into account the opinions of well-known investors / business angels in the start-up scene in your decision-making process or would you follow them?

¹The market sectors are those that appear in the selected 10 early-stage enterprises.




Figure 6: Investor trust network from Q10. Red nodes are investors and participants; green are only investors; yellow are only participants. Node size indicates degree.




Figure 7: Distribution of ranks assigned to the 10 specified early-stage enterprises in Q11.

This question determines the inner circle of trust of investors. A list of well-known investors from the DACH region was presented, and each participant was asked to select those whose recommendations they would trust. As a result, a directed graph representing the social network of trust among the participants and the well-known investors is built. This network is depicted in Figure 6. Yellow and green nodes correspond to the questionnaire participants, who provided their trust relationships and thus have outgoing edges. Green nodes in particular are both participants and well-known investors, i.e., they were listed and filled out the questionnaire, thus have both in- and outgoing edges. Yellow nodes are participants that were not included in the list of well-known investors, and thus only have outgoing edges. Red nodes correspond to the well-known investors who did not participate in the questionnaire, i.e., they only have ingoing edges. It appears that a few well-known investors received the most trust relationships.

Q11: Please select 5 out of 10 early-stage enterprises from the fact sheet and rank them in the order in which you would invest.

In Q11, short information about ten early-stage enterprises was shared with the participants, who have to identify their top-5 choices for possible investment. Figure 7 presents their choices.

Data Preparation. The raw data collected by the questionnaire has to be processed for training and evaluating the recommender algorithms. We use the answers to Q1–Q9 to create the investor profiles (required by KB and HB), Q10 to define the follow relationships (required by SB), and Q11 to define the historical investor-enterprise interactions (required by CB, CF, SB), which we then partition into a train and a test dataset.

We note that the questions on the Likert scale (depicted in Figure 5) determine how much an investor is interested in her profile attributes. We map the 5-point Likert scale to the integer scale 0%, 25%, 50%, 75%, and 100% of relative importance (see KB in Section 4).

Table 1: Ranking Correlation

	CF-UB	CF-IB	CB	SB	KB	HR
CF-IB	0.1736	—	—	—	—	—
CB	-0.2048	-0.0312	—	—	—	—
SB	0.0886	0.0715	0.0841	—	—	—
KB	0.0569	0.0856	0.2521	0.0709	—	—
HR	0.0345	0.1773	0.1540	0.1264	0.4439	—

5.2 Results

Evaluation Methodology. We evaluate the six recommendation algorithms from Section 4, namely neighborhood-based collaborative filtering CF-UB, CF-IB, content-based CB, knowledge-based KB, social-based SB, and hybrid HR recommenders. We compare them using the dataset constructed from Q11. Specifically, we treat the answers of Q11 as “likes” from the investors to early-stage enterprises. So every investor has liked 5 (out of 10) enterprises. We perform multiple splits into train and test data via leave-one-out method. There are a total of 5 folds, where in each fold a liked enterprise is placed in the test set and the other four in the train set. For each user and each fold, a recommender algorithm produces a ranking of the 6 remaining (not in the train set) enterprises.

We measure (i) the correlation across the recommendation lists, and (ii) the ranking accuracy of the recommendation algorithms. Correlation is measured by *Spearman rank correlation* that measures the strength of association between two lists and its direction. It takes values between -1 and 1, where the higher the values, the more similar the lists are. Ranking accuracy is measured by: *Precision@k* ($P@k$), which simply counts the number of relevant items (early-stage enterprises) among the top- k , and by *Mean Average Precision@k* ($MAP@k$) and *Normalized Discounted Cumulative Gain@k* ($nDCG@k$), where the position of the relevant items in the ranked list discounts their relevance. We vary k from 1 to 3. All ranking accuracy metrics range from 0 to 1, with higher values preferred.

Ranking Correlation. Table 1 reports the Spearman rank correlation for every pair of recommendation algorithms. Overall we report very weak correlations except two cases. The content-based (CB) and knowledge-based (KB) recommenders exhibit weak rank correlation, given that they both depend on the content of the items (enterprise description). We also identify a moderate correlation between the knowledge-based (KB) and the hybrid (HR) recommenders, which is explained by the fact that HR is based on KB, as it explicitly uses the rankings that KB produces to define similarities between investors.

Ranking Accuracy. Table 2 presents the ranking accuracy metrics of all methods. The following observations are of interest. First, item-based approaches, such as collaborative filtering (CF-IB) and content-based recommendations (CB), perform badly. This implies that knowledge of the investor’s past investment decisions is probably not a good indicator for her future actions. Second, the knowledge-based (KB) method is not able to produce consistently more accurate recommendations compared to the content-based (CB) method. Put differently, explicitly capturing the investor’s preferences does not help more than observing her decisions. Third, user-based approaches, such as collaborative filtering (CF-UB) and social-based recommendations (SB), perform the best in terms of ranking accuracy, with the latter outperforming the former. This observation essentially corroborates the fact that investors tend to

Table 2: Ranking Accuracy

	P@1	P@2	P@3	MAP@2	MAP@3	nDCG@2	nDCG@3
CF-UB	0.2176	0.2529	0.2235	0.3618	0.4167	0.2450	0.2261
CF-IB	0.1471	0.1500	0.1490	0.2235	0.2725	0.1493	0.1488
CB	0.1235	0.1500	0.1216	0.2118	0.2333	0.1440	0.1254
KB	0.1529	0.1176	0.1176	0.1941	0.2333	0.1256	0.1238
SB	0.2571	0.2607	0.2238	0.3893	0.4393	0.2599	0.2341
HR	0.1235	0.1088	0.1078	0.1706	0.2059	0.1122	0.1107

think alike. Moreover explicit trust relationships among investors are more trustworthy than implicit investor similarity relationships extracted from past behavior. Fourth, we note that the hybrid recommender (HR) performs worse than KB and CF-UB upon which it is based. This implies that investor similarities based on their profiles (in KB) are less useful than similarities based on investors past behavior (in CF-UB).

Discussion. An important conclusion of this study is that there seem to exist circles of trust among investors, and capturing them explicitly (as in SR) or implicitly (as in CF-UB) plays a key role in improving the ranking accuracy of the recommendation lists.

We also want to discuss some limitations of the current work. Due to general lack of availability from the investors, we could only simulate (via a questionnaire) an offline experimental evaluation with a rather limited number of participants. This in turn dramatically reduces the number of user-item interactions we can collect. We also want to note that the low ranking accuracy of the knowledge-based (KB) recommender—and thus of HR as well—could be attributed to potential discrepancies in the collected data. For example, it may be possible that participants in our study did not accurately or completely express their preferences in their profiles, and/or their selection of early-stage enterprises does not reflect their stated preferences.

Moreover, another limitation is that the investors that participated in our study are all from a single geographical area. Similarly, the selected early-stage enterprises all come from the same geographical area. Thus there is the risk of selection bias in the collected data, as investors may have been more familiar with, or already invested in, some of the enterprises.

6 CONCLUSION

This paper presented a design study of recommendation approaches in the domain of investing in early-stage enterprises. A detailed requirements analysis found specific desiderata for a system, including the need to respect investor profiles, learn from past decisions, and propagate trust among investors. An evaluation scenario demonstrates that simple collaborative filtering approaches, and especially a trust-based one, is more effective than knowledge- and content-based approaches.

REFERENCES

- [1] A.K. Achleitner and E. Nathusius. 2004. *Venture valuation - Bewertung von Wachstumsunternehmen: klassische und neue Bewertungsverfahren ; mit Übungsaufgaben und Fallstudien*. Schäffer-Poeschel. <https://books.google.at/books?id=YTZOWAACAAJ>
- [2] David Amis and Howard Stevenson. 2001. *Winning Angels: The 7 Fundamentals of Early Stage Investing*. FT Press.
- [3] Peter Atteslander. 2006. *Methoden der empirischen Sozialforschung* (11., neu bearb. u. erw. Aufl. ed.). Schmidt, Berlin.
- [4] David W. Berkus. 2016. *After 20 years: Updating the Berkus Method of valuation*. <https://berkonomics.com/?p=2752> As consulted online on 12 October 2017.
- [5] Dirk Czarnitzki and Andreas Fier. 2002. Do Innovation Subsidies Crowd Out Private Investment? Evidence from the German Service Sector. <http://ub-madoc.bib.uni-mannheim.de/440/>
- [6] Marco de Gemmis, Pasquale Lops, Cataldo Musto, Fedelucio Narducci, and Giovanni Semeraro. 2015. Semantics-Aware Content-Based Recommender Systems. In *Recommender Systems Handbook*. Springer, 119–159.
- [7] Cynthia Dwork, Ravi Kumar, Moni Naor, and D. Sivakumar. 2001. Rank aggregation methods for the Web. In *WWW*. ACM, 613–622.
- [8] Eurostat. 2016. *Business demography by size class (from 2004 onwards, NACE Rev. 2)*. <http://appsso.eurostat.ec.europa.eu/nui/submitViewTableAction.do> As consulted online on 24 July 2016.
- [9] Eurostat. 2016. *Business demography statistics*. http://ec.europa.eu/eurostat/statistics-explained/index.php/Business_demography_statistics As consulted online on 24 July 2016.
- [10] Alexander Felfernig, Gerhard Friedrich, Dietmar Jannach, and Markus Zanker. 2015. Constraint-Based Recommender Systems. In *Recommender Systems Handbook*. Springer, 161–190.
- [11] Ángel García-Crespo, José Luis López-Cuadrado, Israel González-Carrasco, Ricardo Colomo-Palacios, and Belén Ruiz-Mezcua. 2012. SINVLIO: Using semantics and fuzzy logic to provide individual investment portfolio recommendations. *Knowledge-Based Systems* 27 (2012), 103–118.
- [12] Jose Luis Garcia-Lapresta and Miguel Martinez-Panero. 2002. Borda Count versus Approval Voting: A Fuzzy Approach. *Public Choice* 112, 1-2 (7 2002), 167–184. <https://ideas.repec.org/a/kap/pubcho/v112y2002i1-2p167-84.html>
- [13] Dietmar Jannach, Markus Zanker, Alexander Felfernig, and Gerhard Friedrich. 2010. *Recommender Systems: An Introduction* (1st ed.). Cambridge University Press, New York, NY, USA.
- [14] Leora Klapper and Inessa Love. 2011. The impact of the financial crisis on new firm registration. *Economics Letters* 113, 1 (2011), 1–4. <https://doi.org/10.1016/j.econlet.2011.05.048> As consulted online on 24 July 2016.
- [15] Richard Kühler, Hans-Ulrich Küpper, and Andreas Pflugsten. 2007. *Handwörterbuch der Betriebswirtschaft*. Schäffer-Poeschel. <https://books.google.at/books?id=TO17GQAACAAJ>
- [16] Johannes Luef and Christian Ohrfandl. 2018. Recommender Systems in the Domain of Early-stage Enterprise Investment.
- [17] K. Macharzina and J. Wolf. 2008. *Unternehmensführung: das internationale Managementwissen ; Konzepte, Methoden, Praxis*. Gabler Verlag. <https://books.google.at/books?id=TF99QiEm1mwC>
- [18] Paolo Massa and Paolo Avesani. 2007. Trust-aware recommender systems. In *RecSys*. ACM, 17–24.
- [19] W. D. Maughan. 2010. Crossing the “Valley of Death”: Transitioning Research into Commercial Products: A Personal Perspective. In *2010 IEEE Symposium on Security and Privacy*. 21–26. <https://doi.org/10.1109/SP.2010.40>
- [20] B.H. Meyer and H.C.J. Bloech. 2006. *Stochastische Unternehmensbewertung: Der Wertbeitrag Von Realoptionen*. Deutscher Universitätsverlag. https://books.google.at/books?id=66L_md9C0qoC
- [21] Cataldo Musto, Giovanni Semeraro, Pasquale Lops, Marco de Gemmis, and Georgios Lekkas. 2015. Personalized finance advisory through case-based recommender systems and diversification strategies. *Decision Support Systems* 77 (2015), 100–111.
- [22] Xia Ning, Christian Desrosiers, and George Karypis. 2015. A Comprehensive Survey of Neighborhood-Based Recommendation Methods. In *Recommender Systems Handbook*. Springer, 37–76.
- [23] OECD. 2012. *Business start-up rates diverging across OECD economies*. <http://www.oecd.org/std/business-stats/businessstart-upratesdivergingacrossoececonomies.htm> As consulted online on 24 July 2016.
- [24] OECD. 2016. *Timely Indicators of Entrepreneurship (ISIC4) – Number of enterprise entries (Q1-2007 - Q1-2016)*. http://stats.oecd.org/Index.aspx?DataSetCode=TIMELY_BDS_ISIC4# As consulted online on 24 July 2016.
- [25] William H. (Bill) Payne. 2011. SCORECARD VALUATION METHODOLOGY - Establishing the Valuation of Pre-revenue, Start-up Companies. <http://billpayne.com/wp-content/uploads/2011/01/Scorecard-Valuation-Methodology-Jan111.pdf>
- [26] Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor. 2010. *Recommender Systems Handbook* (1st ed.). Springer-Verlag New York, Inc., New York, NY, USA.
- [27] Markus Rudolf and Peter Witt. 2002. *Bewertung von Wachstumsunternehmen: traditionelle und innovative Methoden im Vergleich*. Springer-Verlag.
- [28] Mona Taghavi, Kaveh Bakhtiyari, and Edgar Scavino. 2013. Agent-based computational investing recommender system. In *RecSys*. ACM, 455–458.
- [29] Wenting Tu, Min Yang, David W. Cheung, and Nikos Mamoulis. 2018. Investment recommendation by discovering high-quality opinions in investor based social networks. *Inf. Syst.* 78 (2018), 189–198.
- [30] Jungsoon P. Yoo, Melinda T. Gervasio, and Pat Langley. 2003. An adaptive stock tracker for personalized trading advice. In *IUI*. ACM, 197–203.
- [31] David Zibriczky. 2016. Recommender Systems meet Finance: a Literature Review. In *FINREC (CEUR Workshop Proceedings)*, Vol. 1606. CEUR-WS.org, 3–10.