DESIGN AND EVALUATION OF AN AI-AUGMENTED SCREENING SYSTEM FOR VENTURE CAPITALISTS

Completed Research Paper

Marc Maurer, TU Berlin, Berlin, Germany, research@maurer.fyi

Tolga Buz, Hasso Plattner Institute / University of Potsdam, Germany, tolga.buz@hpi.de

Christian Dremel, IMD Business School, Lausanne, Switzerland, christian.dremel@affiliate.imd.org

Gerard de Melo, Hasso Plattner Institute / University of Potsdam, Germany, gdm@demelo.org

Abstract

The venture capital (VC) industry has been experiencing a substantive transformation in recent years due to the rise of online sources for startup data and their utilization in identifying promising startup founders. The resulting significant increase in leads has proportionally increased the workload required for screening them along the investment process. Together with a well-known German VC firm, we have designed and implemented a machine learning pipeline trained on a large dataset of 39,114 LinkedIn profiles with their respective screening decisions aiming to enhance the efficiency of the process. Our method boosts screening productivity by auto-rejecting 23% of founders, missing at most 1% relevant ones, yielding a 663:1 accurate rejection ratio. Venture capitalists can adjust auto-rejection up to 57% at 10% miss rate. 84% of users prefer the AI-augmented workflow, while 16% pre-filter profiles for reduced workload. Its success convinced the VC firm to promptly implement the system in production.

Keywords: Venture Capital, Machine Learning, Human-AI Collaboration, Design Science.

1 Introduction

For many years, venture capital (VC) firms have significantly outperformed the market by achieving an impressive estimated 10-year net internal rate of return of 20.20% for investors (Cambridge Associates 2020), primarily through early-stage, high-growth potential companies (startups) in exchange for equity ownership. Unsurprisingly, this enormous profit potential, in combination with the recent "startup-explosion", has ignited a boom in the VC industry (Zhong et al. 2018). With the help of institutional investors, VC firms have grown their globally raised capital to $254 billion in 2022, up from $59 billion in 2012 (Pitchbook 2022). However, as startup failure rates continue to hover around 90% (Krishna et al. 2016), the ability to identify the few startups that can achieve long-term success is a key driver of the success of VC firms (Zider 1998). With a constantly growing pool of opportunities and intensifying competition between VC firms, screening and deciding on thousands of potential deals is becoming increasingly resource-intensive.

The identification of relevant startups, traditionally driven primarily by the network of venture capitalists (VCs) (Tyebjee and Bruno 1984), has shifted towards a more data-driven approach leveraging databases such as Crunchbase, PitchBook, and LinkedIn (Weibl and Hess 2019). These sources allow VCs to find startups earlier by conducting targeted searches (Teten and Farmer 2010), but increase the number of identified startups and require more resources during *screening* (Di Giannantonio et al. 2022). In response to these challenges, some efforts have been undertaken to improve the *screening* phase by developing decision support systems that aim to predict startup success and recommend promising startups for further examination in the *evaluation* phase (Davenport 2022; Retterath 2020). The existing approaches, however, may not align with the preferences of VCs, who often rely on singling out a unique outlier among a multitude of startups and might be hesitant to entrust systems with this crucial selection (Retterath 2020; Wolff et al. 2021). Additionally, some approaches focus on screening one-pager descriptions of startup companies (Retterath 2020), which require manual preparation and a certain level of startup maturity (even if still considered early stage). In contrast, our findings indicate that VCs frequently identify early-stage deals by screening profiles of founders rather than companies and tend to prioritize assistance in the elimination of low-potential founders to enhance their capacity to identify the most promising candidates manually. Recent studies share the approach of predicting underperforming startups that VCs should avoid investing in (Davenport 2022). Nevertheless, a disconnect remains between the theoretical solutions proposed versus the practical stages and techniques employed by VCs. This misalignment with VCs' interests and the lack of clarity regarding the specific applications and methods of incorporating these techniques into their daily investment processes calls for further exploration and refinement. Hence, we define our main research question:

RQ: *How can information systems be developed and tailored to assist VCs in improving their investment processes, particularly during the screening phase, and which strategies can researchers employ to better align their work with the real-world requirements and inclinations of VCs?*

We address this question by using a Design Science methodology to design and evaluate an end-to-end data-driven and AI-Augmented Screening System. This system noticeably reduces the time required for screening founders by eliminating irrelevant profiles while minimizing the risk of overlooking a high-potential founder through misclassification (false negatives). We thereby create an artifact that is highly aligned with VC interests and clearly defines its position and value proposition within the investment stages of a VC firm. To accomplish this, we establish a close collaboration with a German VC firm that is active globally (whose identity cannot be disclosed for business reasons). We focus on automating the *deal origination* (profile discovery and extraction) and *screening* (profile evaluation) phases. We utilize a unique dataset of 39,114 startup founder profiles combined with the firm's screening decisions per profile. We employ NLP techniques that are novel in this domain. By implementing a comprehensive Design Science cycle, we align with the VCs' preferences, iteratively refining our AI-driven system to effectively assess startup founders in a real-world, in-production environment, thus fulfilling our research objectives.

Employing the previously described method, we executed three iterations of the Design Science cycle to create an AI-enhanced screening system. We show that the developed artifact considerably improves screening efficiency by automatically rejecting 23% of founders while maintaining a maximum miss rate of 1%, resulting in a 663:1 accurate rejection ratio. The system offers VCs the flexibility to increase the auto-rejection rate up to 57% at the cost of a 10% miss rate. By closely aligning with the interests of VCs and emphasizing tangible process enhancements, we showcase the readiness of VCs to incorporate research-based designs into their daily operations. The AI-augmented workflow received widespread approval within the firm with an 84% preference rate, leading to a swift migration into production by the VC firm. Our study expands existing knowledge by developing an AI-enhanced model of previously undocumented, superior configuration for VCs utilizing hybrid-intelligence screening. We emphasize the importance of adopting a practitioner-focused research approach, such as Design Science research methodology, and highlight VCs' readiness to embrace innovative solutions and collaborate with researchers when research is practitioner-oriented. By highlighting the limitations of focusing solely on seemingly attractive goals, such as predicting a "unicorn startup", we emphasize the necessity of closely collaborating with practitioners (VCs) to recognize and tackle their real needs, e.g., focusing on founder profiles instead of startup company descriptions to identify early-stage opportunities. Through continuous interaction with VCs, researchers can prioritize addressing their critical challenges, ultimately creating innovative and practical solutions that genuinely support their investment decision-making processes.

2 Literature Review

Our study engages with multiple areas of literature. First, we position our research within the literature on venture capital, focusing on VC investments in entrepreneurial companies – this includes the VC investment decision-making process and its evolution due to the emergence of comprehensive and accurate data on private companies. Additionally, we contribute to the growing body of research on the automation and augmentation of VC investment decision processes through decision support systems, expanding existing knowledge in this domain. Finally, we explore the potential of automated resume assessment methods, rooted in the human resources domain, for automated founder evaluation. By doing so, we aim to streamline the process of identifying and subsequently rejecting unpromising founders.

2.1 Literature on Venture Capital (VC)

Da Rin et al. (2013) categorize research on venture capital into four areas: (1) VC investments in entrepreneurial companies, covering investment decisions, contracting, and *post-investment* topics; (2) analysis of VC firms, focusing on their organizational structure, investment strategies, and relationships with other firms and partners; (3) returns of VC investments, including data and methodological challenges, return estimates, and comparing VC returns to other private equity investment returns; and (4) the impact of VC on the economy and public policy, exploring its contributions to innovation, entry, employment, and growth. Our study is situated within the realm of VC investments in entrepreneurial companies (1), specifically addressing investment decisions as an essential component of the VC cycle. We contribute to research on making investment decisions and how a part of this process can be augmented by AI.

Gompers and Lerner (2004) describe the investment decision process as an essential component of the venture capital cycle, representing the structured approach taken when investing in startups. Since its inception by Wells (1974) and extension by Tyebjee and Bruno (1984), the structure of the process has remained relatively unchanged (Kollmann and Kuckertz 2010). Over the years, research has mainly delved deeper into the existing phases of the decision-making process, consequently refining and partitioning them. This has produced variations of the decision framework, each comprising distinct stages tailored to specific contexts. The refined version introduced by Tyebjee and Bruno (1984) has achieved widespread adoption within the academic community. Thus, our study will adopt their framework, which presents the following stages: (i) *deal origination* (identification of potential investments), (ii) *screening* (narrowing down opportunities to a manageable set), (iii) *evaluation* (analyzing the risk and return of each deal), (iv) *deal structuring* (involving negotiations to determine investment specifics), and (v) *post-investment* (including monitoring, consultation, and facilitating exits via mergers, acquisitions, or public offerings).

2.2 Literature on the impact of data in the VC investment process

While the structure of the VC investment process has remained consistent, the methods employed within this process have undergone significant transformation. Weibl and Hess (2019) describe how the emergence of data sources and analytics techniques have strongly affected the techniques employed in each stage, based on interviews with 13 VCs. Their findings reveal that new data and techniques predominantly affect the *deal origination* and *screening* phases of the process – they highlight the significant shift in the *deal origination* process from an inbound approach, such as receiving email applications, to an outbound orientation, where VCs proactively scan the market using data-driven methods to engage with promising startups. This includes extracting information from investment network platforms or social networks, enabling VCs to establish early relationships with entrepreneurs and increasing their ability to identify profitable investment opportunities. Wolff et al. (2021) agree and further elaborate that the *screening* stage has consequently evolved into a more data-driven phase as well. As VCs now deal with larger pools of startups, they spend an average of only six minutes per startup during the *screening* phase, resulting in a low advancement rate of 9.2%. Our own dataset shows an even lower acceptance rate of 3.8%. The study identifies the founder team's strength as a key factor for success in the investment decision process. Startups are often turned down by VCs for not matching the firm's investment focus or for not being a VC case. However, a startup with a perfect team score of 5 has a 36% chance of being presented to decision-makers during the *pitch presentation*, the final step of the *evaluation* phase. The likelihood decreases to 2% for teams with a score of 3, and it is almost zero for teams with even lower scores. Although it is not the sole criterion, team strength is a critical filter in the initial *screening* of startups, with a clear link between the initial score and progression to the *evaluation* stage. This correlation aligns with the findings of Wolff et al. (2021), who differentiate between the *screening* phase and an *initial evaluation* phase in their framework. They refer to what this paper calls the *evaluation* phase as the *due diligence* phase.

Based on these insights, we aim to develop an effective approach for identifying and filtering out founders whose team's strength is too low to be further considered in the investment process. To achieve this, we seek to investigate the feasibility of utilizing the founding team strength as a go/no-go criterion for a progression within the *screening* stage.

2.3 Literature on VC decision support systems

Previous research has aimed to support VCs in their investment decision-making processes by leveraging data-driven methods (e.g., Cao et al. 2022). Most of these studies focus on predicting the success or failure of startups by assessing potential investments through data obtained from online platforms such as Crunchbase, PitchBook, or LinkedIn, which provide broad and high-quality data on startups (Retterath and Braun 2020). Within this domain, researchers have developed various data-driven methods including machine learning (ML) techniques for predicting future success or failure. A comprehensive review of these has examined non-deep learning-based ML approaches for predicting startup success (Baskoro et al. 2022), while Cao et al. (2022) synthesized the findings of research focusing on deep learning methods for startup evaluation. Recent research benchmarks the performance of ML models in predicting startup success against VCs, e.g. Retterath (2020), who found that the success predictions of an XGBoost tree model based on early-stage company descriptions outperformed 111 venture capitalists by 29% on average, suggesting that predictive models can be effective for screening startups. Supporting this notion, Davenport (2022) found that an XGBoost model could identify poor investments. These studies advocate for the use of ML models to filter out potential bad investments, demonstrating a growing interest in research that addresses specific use cases within the investment process (Röhm et al. 2022) – shifting the focus from earlier studies, such as Krishna et al. (2016), which aimed to provide a general framework for startup predictions.

Röhm et al. (2022) explore the impact of applying ML methods on the investment decision process, concluding that they have the potential to fundamentally transform the traditional investment processes by augmenting human decision-making and fostering a hybrid intelligence approach. The author argues that the adoption of data-driven methods is essential for the continued success of VC firms. The study also distinguishes between AI-driven and data-driven VC firms based on the sophistication of the methods employed. Some AI-driven VC firms maintain close ties with research in this field, as evinced by collaborations of Retterath (2020) with German VC firm Earlybird and Cao et al. (2022) with Swedish investment firm EQT. These examples highlight the growing integration of ML and AI in the VC landscape, emphasizing the importance of leveraging data-driven methods to enhance investment decision-making and drive success in the industry. This study follows this notion by partnering with a German early-stage VC to create research aligned with the reality of the modern VC process.

Our research aligns with the study by Retterath (2020), which introduced an XGBoost model for automating the second *screening* stage in the investment process. They utilized PitchBook data to predict startups' success in the *screening* stage and prioritized the recall metric as a critical variable during the *screening* stage to minimize the risk of overlooking potentially successful investments (i.e., minimizing false negatives). Our study also focuses on the *screening* stage but emphasizes its connection to the first *deal origination* stage by utilizing a large LinkedIn dataset combined with the collaborating VC's history of screening decisions. This enables our solution to be applicable at an earlier stage and with less data preparation, as founders update their LinkedIn profiles independently, while company descriptions must be created by a service like PitchBook or a VC employee, requiring time and publicly available information about the startup – a startup in "stealth mode" can be detected via an entry in the founder's LinkedIn profile. Our research further underscores the importance of the recall metric by setting a fixed maximum value that our trained models must comply with, optimizing our model to reject the highest possible number of founders while maintaining the fixed recall value. These measures make our solution more scalable and competitive than previous approaches, including the work of Retterath, addressing and solving the limitations explained in their work.

Through this approach, we contribute to the current literature by providing a pragmatic, actionable framework for VCs to employ data-driven decision-making. This can lead to substantial resource savings without additional operational demands, seamlessly integrating artificial intelligence into formerly manual tasks. Moreover, by emphasizing the elimination of less promising founders, VCs can optimize resource allocation and focus on assessing a select group of high-potential founders and their startups, ultimately improving the efficacy and efficiency of their investment decision-making processes.

2.4 Literature on automated founder evaluation techniques

In our study, we strive to enhance the efficiency of the VC *screening* process by automatically identifying and rejecting unpromising founders at the *screening* stage. We aim to automatically analyze founder profiles of recently launched ventures, discovered during the *deal origination* phase using automated data collection from LinkedIn, a common data source in the VC domain.

A LinkedIn profile effectively functions as an online resume for founders, listing education and work experiences (Sumbaly et al. 2013). Research in this area, originating from startup success prediction, primarily extracts manual features from resumes or employs ML techniques to pinpoint the most influential features. Recently, the VC firm Earlybird conducted a more comprehensive examination of founders and their backgrounds in two projects: one investigating the most influential team characteristics of founders (Torssell 2022), and another developing a model based solely on founder characteristics to predict early-stage startup success (Wärnberg Gerdin 2022). Both studies utilized manually extracted features to assess feature importance and predict startup success, aligning with prior research working with features based on clearly defined evaluation criteria employed by manual screeners (Corea et al. 2021; Ferrati and Muffatto 2021; Sharchilev et al. 2018). Interestingly, to the best of our knowledge, scholars have not yet introduced techniques to expand the scope of information extracted from founders' backgrounds beyond the addition of manual features. While not yet applied in the VC space, these techniques are already well-established in the human resources domain for the automatic evaluation of resumes. In this study, we further investigate these techniques and discuss how they can be integrated into an automated founder evaluation pipeline to efficiently eliminate unpromising candidates, thereby enhancing the decision-making process in VC investments.

Techniques for extracting resume information applicable in this study primarily stem from the human resources (HR) research domain within information systems, specifically in the areas of resume classification and job recommendation systems. Resume classification aims to evaluate and categorize resumes based on predefined metrics, while job recommendation systems strive to extract as much information as possible from a resume to match it with the most appropriate job description. Overviews of the existing literature, methods, and algorithms can be found in the works of Pimpalkar et al. (2023) and Mehboob et al. (2022).

A significant portion of research in this area employs advanced Natural Language Processing (NLP) techniques to extract information, a method that is rarely utilized in the founder evaluation space. Various embedding techniques are applied to transform the textual data in resumes to enable the use of analytical methods. For instance, He et al. (2021) apply Word2Vec embeddings (Mikolov et al. 2013), Chung et al. (2023) employs GloVe embeddings (Pennington et al. 2014), Pimpalkar et al. (2023) utilize Doc2Vec embeddings (Le and Mikolov 2014), and Chung et al. (2023) explore the use of BERT-based embeddings (Devlin et al. 2018). Additionally, these studies often incorporate distance similarity measures, such as cosine similarity (Chung et al. 2023; Shovon et al. 2023), and further process the data using non-deep learning methods (Pimpalkar et al. 2023) or deep learning techniques (He et al. 2021). The overall process employed in this domain is outlined by Pimpalkar et al. (2023) in six stages: (A) text extraction from (resume) files, (B) text pre-processing, (C) feature extraction and label encoding, (D) resume classification model construction, (E) ranking, and (F) performance evaluation.

In this study, we adapt this approach for founder evaluation based on online resumes, such as LinkedIn profiles, drawing inspiration from Pimpalkar et al. (2023) and Mehboob et al. (2022). We further enhance the information density by employing a BERT-based embedding model as discussed in Chung et al. 2023. Moreover, we compare the information density by comparing the performance of BERT-based features with features based on GPT-3 embeddings (Brown et al. 2020).

By implementing these techniques, we aim to address the research gaps explained above: the novelty of our contribution lies in adapting neural network feature learning techniques to combine manually extracted features with learned features to train a classifier that filters non-relevant founders while strictly focusing on maintaining highest-possible recall. We merge this technical approach with a more precise understanding of the decision-making processes within the collaborating VC, e.g., their focus on founder profiles instead of startup company descriptions, and utilize the Design Science approach (explained below) to develop an artifact that addresses the challenges of the VC we have collaborated with and is well-accepted by its employees.

3 Research Method

3.1 Design Science

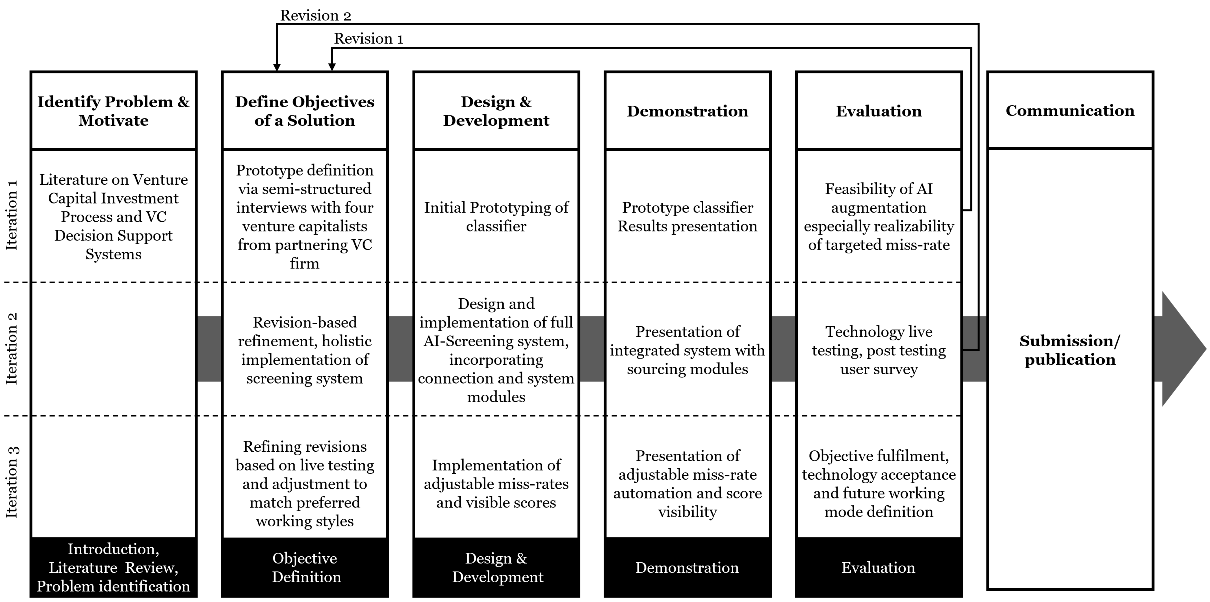

Figure 1. Research Design (based on Peffers et al. 2007 and adopted from Meier et al. 2019)

This study employs a Design Science approach as described by Peffers et al. (2007), who outline the Design Science Research (DSR) methodology as a method for conducting information systems research in six distinct stages, aiming to provide a structured process for addressing real-world problems by creating and evaluating practical artifacts. We address the real-world problem that VCs are facing and iteratively design and implement a system that utilizes techniques that are novel in this domain. Our research follows a problem-centered approach, identifying the issue through a comprehensive literature review and practical challenges faced by a partnering VC firm. Due to its practical focus, the DSR methodology is well suited for this work and in alignment with similar approaches, such as *Engaged Scholarship* (Van de Ven 2007), which advocates for embracing the engagement of research with practitioners. The six stages of the DSR are shown in Figure 1 and implemented in our study as follows: (1) problem identification, acknowledging the increased competition, growing time requirements for startup screening, and the risk of missing high-potential deals due to limited screening capacity; (2) objectives of a solution, establishing clear goals for our system, including the reduction of screening time, earlier detection of deals through AI augmentation combined with data-driven sourcing, and explicit risk management through the ability to control the maximum miss rate (false negatives) of high-potential deals; (3) design & development, delineating system architecture, methods and ML pipeline, and unique learned features by employing techniques derived from other relevant domains; (4) demonstration, verifying the system's effectiveness with a VC through implementation and operation in a real-world scenario; (5) evaluation, appraising the artifact's performance using quantitative and qualitative metrics, demonstrating reduced screening time and VC preference for augmented systems; and (6) communication, showcasing the utility, design, and innovation of our work through this paper. We detail the main research phases of our Design Science research in the following sections.

4 Designing an AI-Based Screening Augmentation System

4.1 Problem identification

The rapid growth and heightened visibility of startups, coupled with an increased influx of capital, have expanded the pool of potential deals for VCs. This has led to fierce competition in swiftly identifying the most promising startups. However, as the number of startups to be screened grows, VCs face a significant time constraint, with up to 20% of their weekly working hours dedicated to this highly manual task, as indicated by our survey results discussed below. This time commitment detracts from other essential VC activities, and the limited time capacity of screeners imposes constraints on the data-driven sourcing techniques employed by VC firms. Consequently, this may result in missed potential deals.

Rising startups' visibility intensifies VC competition. The rapid growth in the number of new startups and a significant increase in available capital, combined with enhanced visibility facilitated by advancements in technology and digital platforms, have led to an ever-expanding pool of potential deals for VCs to screen (Pitchbook 2022; Röhm et al. 2022; Zhong et al. 2018), intensifying competition among VCs to become faster in identifying and closing the most promising deals (Retterath and Braun 2020; Röhm et al. 2022). Furthermore, VCs are shifting screening tasks to junior team members and interns due to time constraints, which is a potential cause for overburdening and a loss of decision quality (Wolff et al. 2021).

Increasing expenditure of time for startup screening. This study identifies a key issue concerning the substantial amount of time VCs must devote to screening startups, which detracts from their ability to concentrate on other critical activities, such as closing deals and supporting portfolio companies. Our survey conducted at the collaborating VC reveals that employees, particularly juniors, spend a significant portion of their week (up to 20%) on screening. This finding aligns with similar results by Wolff et al. (2021). The problem is expected to worsen due to the growing adoption of advanced data-driven sourcing techniques in the industry (Roy et al. 2020), resulting in a steadily growing deal funnel. As efficient screening becomes increasingly important, it is crucial to develop automated strategies that align with VCs' interests, optimize their time usage, and streamline the investment process.

Missed deals due to constraints on data-driven sourcing. Further, an issue identified in this study is the limitation on data-driven sourcing due to constrained screening capacity, potentially leading to missed deals. Presently, VC firms primarily create their data-driven sourcing tools in-house to yield a manageable number of startups for manual screening (Röhm et al. 2022). However, this approach does not automatically prioritize all identified startups and founders, but usually only provides a long list of hundreds or thousands of profiles. Hence, high-potential leads may be overlooked if screening capacities are limited, as screeners cannot effectively prioritize or eliminate profiles based on quality and must screen all profiles uniformly. This emphasizes the necessity for a more efficient method for screening to leverage advanced data collection methods while reducing the risk of missed opportunities.

4.2 Objective definition

The primary objective of this Design Science study is to devise an AI-powered screening system that integrates with a data-driven sourcing tool using LinkedIn API access for scoring founders and filtering out non-relevant individuals during the *screening* phase of the VC investment process. Our goal is to achieve a reduction of time expenditure and an increase in flexibility for the *screening* process, while minimizing the risk of overlooking promising startup founders.

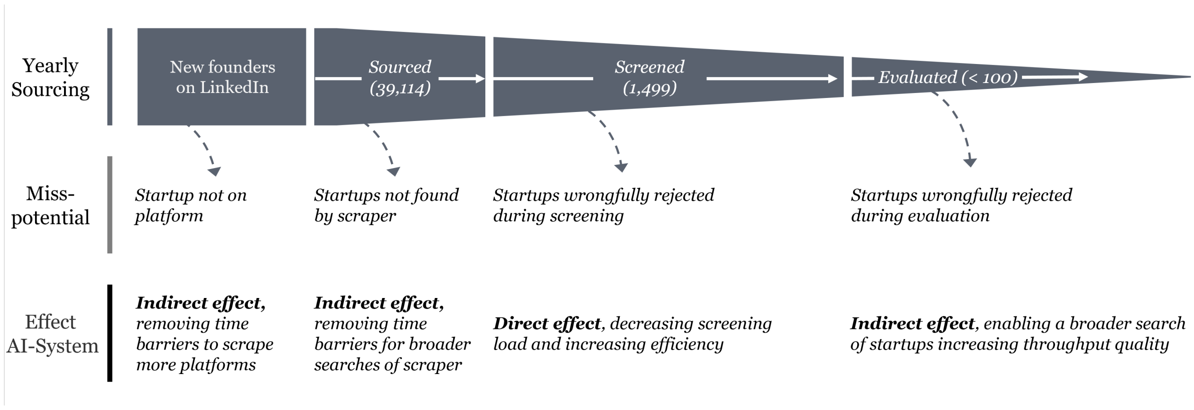

Figure 2. Funnel progression and potential missed deals of LinkedIn data-based sourcing.

We aim to extend the existing literature by employing a Design Science approach to provide practitioners with a well-defined roadmap for establishing their own AI-Augmented Screening Systems, thereby contributing to the broader field of VC decision support systems. Additionally, we introduce NLP techniques for evaluating startup founders exclusively using their LinkedIn profiles, which has not been previously explored within the context of VC decision support systems. We opt for LinkedIn profiles as the evaluation basis, since founders frequently announce their new ventures on this platform, often while the venture remains in so-called "stealth mode" and is not discoverable on other platforms. This allows VCs to identify and engage with promising founders ahead of their competitors (especially if the latter only review company descriptions or pitch decks), enhancing their prospects for future collaboration during funding rounds, which in turn enables them to invest sooner and secure better terms. Ultimately, this leads to an improved competitive position for the VC firm.

The sought system should exclusively depend on the input data, eliminating the need for manual input. It must effectively discard unsuitable founders while adhering to a predefined maximum miss rate, calculated as the ratio of relevant founders missed to the total number of relevant founders. Hence, the two fundamental objectives of the AI-enhanced screening system are: the automatic rejection of unsuitable founders, and a decoupling of the sourcing and screening time through augmented screening.

Reduce screening time: automatically reject unpromising founders. The first objective of this Design Science research paper is to develop an automated system that accurately identifies and excludes non-relevant founders, thereby enhancing the screening efficiency without the risk of missing high-potential deals. Our goal is to define an acceptable rate of missed high-potential founders with the cooperating VC and subsequently train a classifier that rejects as many non-relevant founders as possible while staying under the predefined miss rate. A miss rate of 1% has transpired to be acceptable from our initial survey with employees of the considered VC, meaning the system may reject at most one percent of high-potential founders.

See more deals earlier: decouple sourcing and screening time via AI-augmentation. The second objective of our study is to remove the constraint on data-driven sourcing techniques by decoupling the output of the data-driven sourcing and the actual required screening time. We achieve this by developing an augmentation for the existing *screening* process that scores each incoming founder and enables screeners to only screen either a set percentage of the top-scoring founders (discarding a flexible percentage of low-scoring founders based on time capacity) or rank founders and only screen for a fixed time budget in each week the best scoring profiles among all retrieved profiles. By incorporating an adjustable threshold into the algorithm, we aim to provide VCs with the ability to customize the system according to their screening capacity, thereby achieving a balance between minimizing missed strong founders and maximizing time savings for VCs through prioritization.

Definition of success: how do we measure performance? We measure the success of the augmented screening system by: (1) the potential reduction in time spent per screener by automating the rejection of founder profiles with a maximum miss rate of 1%, and (2) the willingness of VCs to adopt the algorithm's scores in their weekly *screening* process. We assess the time reduction by comparing the outcomes of VCs and the algorithm in a live screening environment, calculating the achievable time savings per screener per week. Additionally, we gauge the VCs' interest in utilizing the algorithm through a post-testing survey, with successful adoption by a single VC firm indicating potential generalizability to other firms.

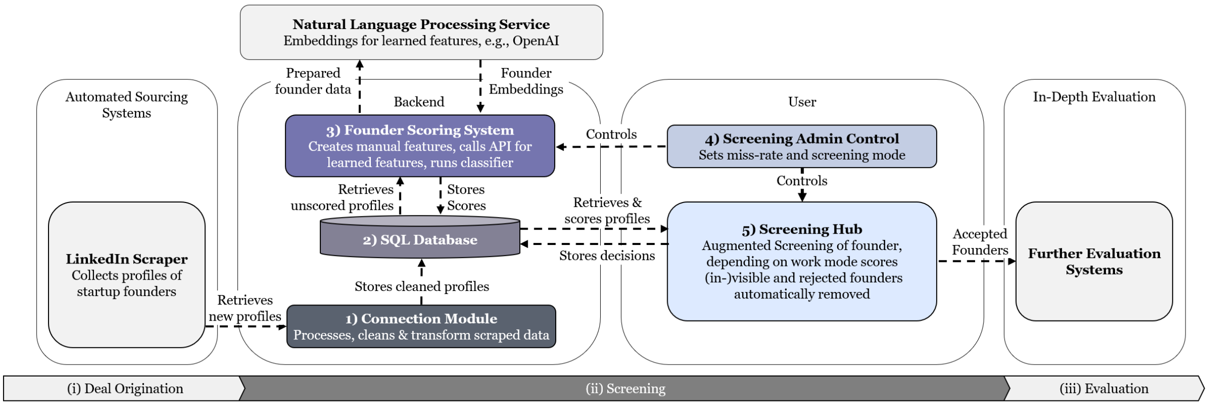

Figure 3. AI-Augmented Screening System Design

4.3 Design and development

The Design and Development closely follows the defined objectives. The following sections describe the overall system as well as the key modules and techniques leveraged to create the AI-Augmented Screening System.

AI-augmented screening system sesign. The screening system comprises five distinct modules as illustrated in Figure 3. The Connection Module (1) connects to data-driven sourcing, specifically a LinkedIn data collector. The SQL Database Module (2) organizes and stores the collected data, including founder scores and screener decisions. The Founder Scoring System (3) invokes a ML model to continuously evaluate new startup founders based on both manually extracted features and learned features derived from a NLP service that generates embedding vectors. The Screening Admin Control (4) manages the weekly screening process and configures the classifier for the VC. The Screening Hub (5) enables venture capitalists to screen the retained and scored profiles.

Founder classification and scoring system. We present a novel founder classification and scoring system by combining venture capital decision support methods (prior studies) and NLP-based feature learning techniques from HR research. We leverage the following methods to create a classifier as artifact to predict if a founder would be rejected during the *screening* stage of the venture capital investment decision process (Tyebjee and Bruno 1984). The method is built upon the framework by Pimpalkar et al. (2023), which proposes the following steps for selecting resumes with Machine Learning: (A) extracting texts from (resume) files, (B) pre-processing the texts, (C) feature extraction and label encoding, (D) resume classification model construction, (E) ranking, (F) performance evaluation. We modify this framework by adjusting the stages to fit the evaluation of founders based on their LinkedIn profile, which results in the following five steps:

A: Extracting founder information. In this step, targeted searches are performed on LinkedIn Recruiter to discover new startup founders. The available information on each founder is then collected and saved for further analysis. Our dataset consists of 39,114 founder profiles including the standard information available on a LinkedIn profile, such as academic background, professional experience, and general information. Each profile further has a binary label indicating whether a profile was approved in *screening* (true for 1,499, i.e., 3.8% of founders). The exploratory analysis of feature distributions between the collected data and approved profiles reveals interesting patterns: Most founders are in North America with 32% (collected) and 42% (approved) founders, significantly outpacing Asia Pacific (16%

and 15%) and Europe (11% and 10%). Regarding the educational background, there is a noticeable uptick in MBA (12% and 17%) and PhD (5% and 9%) holders among approved profiles. Stanford University emerges as the top institution (3% of collected and 8% of approved profiles), with a substantially increased representation of Computer Science graduates from 9% to 20% in approved profiles. Regarding work experience, approved founders have 5.3 years of experience (as opposed to 4.5 years in the dataset), with an increased presence of Google as the most common previous employer (2% pre-screening, 6% post-screening) and Software Engineer as the most common previous job title (5% in collected dataset, 13% in approved dataset). Some features appear to be strong indicators for approval, regarding the ratio of occurrence in the approved set to the full set: BCG, Facebook, and McKinsey & Co. as previous employers, computer science, electrical engineering, and mathematics as field of study, software engineers, product managers, and advisors as previous role, Carnegie Mellon University, Stanford University Graduate School of Business, and Stanford University as alma mater.

B: Structuring and pre-processing of data. The raw data needs to be cleaned, pre-processed, and structured. This involves extracting proper date, location, and numeric values from respective data points and removing corrupted samples. To ensure repeatability, an automated pipeline termed the Connection Module is established to perform these tasks and store the processed data in an SQL Database, specifically utilizing PostgreSQL.

C: Feature extraction and label encoding. Features are chosen based on discussions with the partnering VC on what they deem the most relevant features as well as existing literature (Torssell 2022; Wärnberg Gerdin 2022; Corea et al. 2021; Ferrati and Muffatto 2021; Retterath 2020). Overall, we extract manual categorical and learned features from venture capitalist profile reviews, obtaining a 235-dimensional feature-vector with 31 manual and 204 learned dimensions. The learned features are derived from pre-trained models such as Sentence-BERT (SBERT) and OpenAI GPT-3. The manually extracted features, represented as numeric, binary, or one-hot encoded values, encompass four data types: 9 dimensions of academic background (degree type, top-tier university, years of academic experience), 9 dimensions of professional experience (serial founder, top-tier company, years of professional experience), and 13 dimensions of general features (actual founder, fund-relation, LinkedIn connections, region, startup hub location). Learned features embed and cluster information from the four profile sections, with 51 dimensions per section. We combine these learned and manual features into the final feature vector and create two feature vectors of the same dimensionality for further testing, one using SBERT and the other using GPT-3. This is described below in further detail.

D: LinkedIn profile classifier construction. We train four classifiers: Logistic Regression, Random Forest, Gradient Boosted Trees, and XGBoost on the datasets from Step C, optimizing for recall (i.e., $\frac{TP}{TP+FN}$ where TP = true positives and FN = false negatives) to minimize overlooking strong founders with a targeted recall of 99%. Hyperparameters are tuned using five iterations of five-fold Cross Validation, adjusting from standard parameters to enhance performance. The final hyperparameters for the model implementations are as follows: *penalty='l2', c=1.0, max_iter=10000* for Logistic Regression; *n_estimators=200, max_depth=3* for Random Forest; *n_estimators=100, max_depth=5, learning_rate=0.1, loss='deviance'* for Gradient Boosted Trees; *colsample_bytree=0.8, learning_rate=0.01, max_depth=3, n_estimators=222, subsample=0.8, objective='binary:logistic', eval_metric='logloss'* for XGBoost. All models share the same class weights (0:1, 1:20) to balance the class distribution, an 80/20 train and test data split, a target false negative rate of 1% (adjusted by threshold), and are trained in Python with the scikit-learn library.

E: Augmented screening of profiles and evaluation. Model performance is evaluated in a live setting, with the primary objective of minimizing the startup funnel size while maintaining a constant false negative rate of 1% (corresponding to a recall of 99%). The false negative rate, defined as the ratio of false negatives to the sum of false negatives and true positives, is targeted at 1% ($FNR = \frac{FN}{FN+TP}$). The model's threshold is adjusted to achieve this rate when further training ceases to improve performance. All models are calibrated to meet the 1% criterion, and their performance is assessed by comparing the number of profiles rejected at this threshold. Our findings reveal that logistic regression exhibits the highest performance in this specific use case, discarding 23.47% (17.86%) of founders during training based on GPT-3 (SBERT) derived features. This is followed by Random Forest, Gradient Boosted Trees, and XGBoost, which discard 21.63% (13.65%), 18.25% (13.59%), and 12.34% (7.86%) of founders, respectively. All models demonstrate superior performance when utilizing GPT-3 based features compared to SBERT based learned features, indicating a higher information density in the GPT-3 learned features. It is important to note that, as our models are optimized solely for recall at a strict rate, other metrics such as precision are deprioritized.

Feature extraction and label encoding. In our system design, we integrate manually-extracted features, as demonstrated in Retterath (2020), and learned features, as shown in Pimpalkar et al. (2023). We select features based on semi-structured interviews with investors at the VC firm in which the investors explained how they approach the manual *screening*. These findings were consolidated with a guide by the VC firm that outlined the best-practice approach to screen founders. The combination of these inputs resulted in a long list of criteria that should be incorporated explicitly via manual features (e.g., number of LinkedIn connections) or explicitly via learned features (e.g., what kind of work experience does the founder have in comparison to other founders?). To facilitate feature learning, we employ pre-trained Transformer neural network models, namely SBERT and GPT-3. These models generate text embeddings (i.e., high-dimensional numerical vector representations) for specific resume data segments, which are subsequently clustered for aggregating the information, and the distances for each feature are computed. The feature learning process is defined as follows:

1. For each input string $\boldsymbol{x_i}$, create embedding vectors using two models: *all-MiniLML-L6-v2* (SBERT) and *text-embedding-ada-002* (GPT3), with 768 and 1,536 dimensions, respectively. Denote the embeddings as $\boldsymbol{e_i^{LM}} \in \mathbb{R}^d$ with $LM \in \{SBERT, GPT3\}$ and $d \in \{d_{SBERT} = 768, d_{GPT3} = 1536\}$.
$$\boldsymbol{e_i^{LM}} = LM(\boldsymbol{x_i})$$
2. Perform K-Means clustering on all embedding vectors:
$$\text{K-Means}(\{\dots\})$$
3. Let $\boldsymbol{c_k^{LM}}$ be the centroid of the *k*-th cluster. Measure the cosine similarity distance between each centroid and each embedded vector (with X_j denoting the j^{th} component of a vector $\mathbf{X}$):
$$d_{ik} = S_C\left(e_i^{LM}, c_k^{LM}\right), \text{ where } S_C(\boldsymbol{A}, \boldsymbol{B}) = \frac{\sum_{j=0}^{n=d} A_j B_j}{\sqrt{\sum_{j=0}^{n=d} A_j^2}\sqrt{\sum_{j=0}^{n=d} B_j^2}}$$
4. Use the calculated distance to each K-Means centroid as learned feature input for a classifier: Let $\boldsymbol{F_i^{LM}} = [d_{i1}^{LM}, d_{i2}^{LM}, \dots, d_{iK}^{LM}]$ be the feature vectors for the i-th input string, where K is the number of clusters. Then, combine this feature vector with manually created features and use the resulting vector as input for the classifiers scoring the founder profiles.

4.4 Demonstration

Initially, we demonstrated the AI-Augmented Screening System exclusively in the form of the Founder Scoring System, presenting its potential to score and reject founders at the targeted miss rate of 1%. Our demonstration was based on an asynchronous setting, wherein collected founders were evaluated with a time delay. In the second demonstration, we implemented the fully-developed system for a duration of 10 weeks in a live setting, facilitating a concurrent screening of 5,807 founders in tandem with venture capitalists. This system comprises the automated pipelines, seamless integration with the LinkedIn data collector, the natural language processing module, as well as the internal mechanisms for forwarding profiles to the screening platform, ultimately advancing them to the *evaluation* phase. This setup enabled us to demonstrate the system's functionality and performance within a realistic environment. Lastly, we fine-tuned the pipeline by incorporating the feedback procured following the live testing phase. The implementation of visible scores and AI-augmented screening at a general level embodies the long-term integration of the system.

4.5 Evaluation

The presented AI-Augmented Screening System is the result of three design cycle iterations as follows.

I. Evaluation cycle: feasibility. In the initial evaluation phase, our objective was to ascertain whether extracting sufficient information solely from a founder's LinkedIn profile can reliably eliminate unpromising candidates. Following the development of our prototype, we discovered that a logistic regression classifier, utilizing the manually-selected and GPT-3-derived features as defined in 4.3C, yields the highest amount of rejections while maintaining a 1% maximum miss rate. This approach filters out 23.47% of unpromising founders across numerous training and testing iterations. Upon demonstrating its efficacy and evaluating the results, we proceeded to the next evaluation cycle.

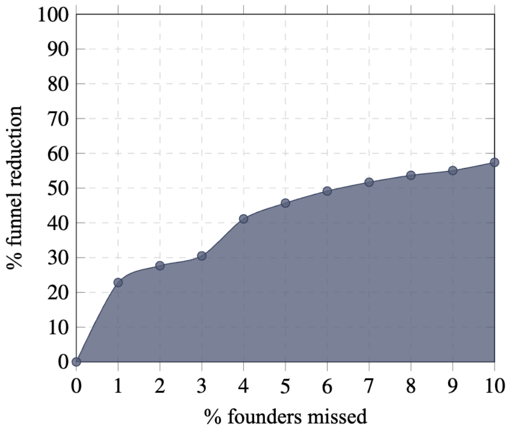

Figure 4a Automated Rejection Potential

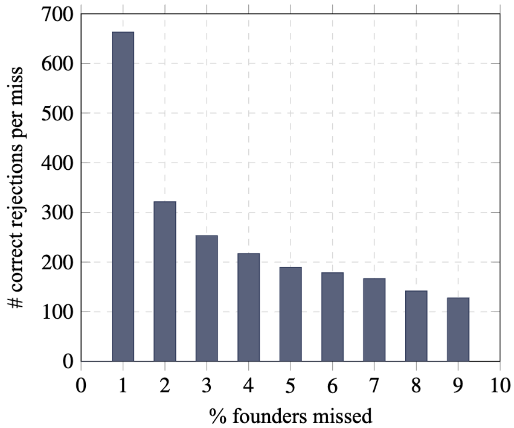

Figure 4b Rejection Trade-Off

II. Evaluation cycle: technology live testing. In the second evaluation cycle, we conducted a ten-week live test of the AI-Augmented Screening System. This assessment incorporated an adjustable miss rate, enabling the VC firm to modify screening loads based on weekly availability and data collector output. We established a dedicated backend, control hub, and interconnected screening hub for the entire system. Our findings revealed that 22.28% of founders can be accurately rejected while maintaining a miss rate below 1% (See Figure 4a). Nine investors used the system at varying frequencies with junior members using the system at higher intensity due to their higher share of time dedicated to screening founders. Further analysis showed that the system could effectively discard over 57.4% of non-relevant founders when a higher miss rate of 10% was allowed (See Figure 4a). This demonstrates that as the initial *screening* funnel expands, the system can maintain a manageable volume while only excluding a small percentage of potentially strong founders. Furthermore, we evaluated the ratio of accurately rejected founders to incorrect rejections at different miss rates. Our findings revealed that at a 1% miss rate, for each erroneous rejection (high-potential founder), 663 founders are correctly rejected (non-relevant founder), while at a 10% miss rate, 128 founders are accurately rejected for every wrongful rejection (See Figure 4b). This indicates the efficiency of our assessment in accurately identifying and rejecting founders at various levels of miss rates.

III. Evaluation cycle: objective fulfilment and technology acceptance. In the third evaluation cycle, we precisely assessed the proportion of founders rejected and compared it with the survey results from venture capitalists involved in the *screening* process, determining the potential time savings for each screener per week. Our survey indicated that junior team members spend a larger portion of their work time (up to 20%) on screening compared to senior members, who allocate about 3 hours per week, if at all (time based on self-declared hours in survey). We demonstrate the time-saving benefits across varying screening intensities (Figure 5a). Moreover, we incorporate an additional feature that not only rejects unsuitable founders but also maintains scores for all founders, even those not rejected following the screeners' survey results (Figure 5b). This enhances the overall experience beyond the mere dismissal of unpromising candidates. By enabling AI-augmented screening for non-rejected founders, we further improve the screening experience, as screeners can now consider an initial score for all profiles and adjust their process accordingly, e.g., by prioritizing profiles through ordering by score. All investors confirmed the efficiency gains from the augmented screening process.

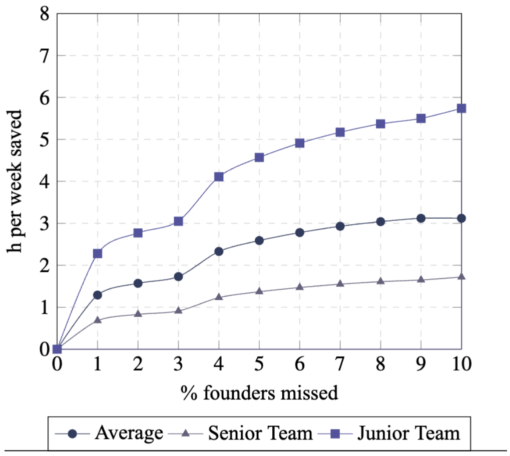

Figure 5a. Weekly Time Savings per Screener

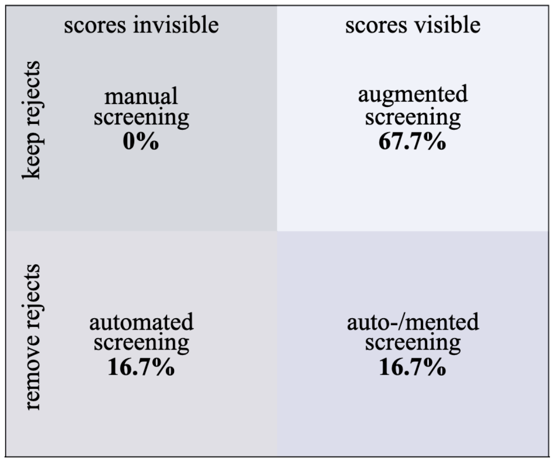

Figure 5b. Work Preference Matrix

Future working mode and handover. In partnership with the VC firm, we have the following future working mode for the AI-Augmented Screening System after three iterative evaluation cycles, encompassing three primary attributes: (1) dynamic weekly adjustments of the targeted maximum miss-rate, allowing the team to efficiently allocate resources based on their capacity; (2) recommendations for rejecting founders according to the targeted miss-rate, without automatically eliminating them – this removes non-relevant founders from the *screening* process while keeping them visible; and (3) integration of scores from non-rejected profiles into the screening hub for an initial assessment. This method employs AI-augmentation for all profiles, not only rejected ones, improving decision-making.

In conclusion, our research has effectively incorporated AI-augmented screening alongside human screeners, achieving a harmonious integration with the LinkedIn data collector. This synergy between artificial and human intelligence facilitated an optimal screening approach for future applications. By tailoring the process to the venture capitalist's time availability, we significantly reduced their weekly workload while eliminating constraints on data-driven sourcing. As a result, venture capitalists can confidently expand their data-driven sourcing initiatives without overwhelming their investment teams.

4.6 Comparison with existing methods

There are existing solutions for automating or accelerating this task: AI-driven tools such as WALE (accessible via wale.ai), Sircular (sircular.io), and Kaiku (kaiku.co) promise efficiency gains in deal sourcing. Other tools like ArkiFi (arkifi.ai) provide process automation. In addition, there are data-driven portfolio management tools like Aumni (aumni.fund) and Affinity (affinity.co). Unfortunately, these are all proprietary tools that require paid access before they can be evaluated. A comparative study could be conducted in future work to compare properties such as functionality, benefits, performance, scalability, and cost.

4.7 Implications for research on Information Systems

Historically, the VC industry has not been a major focus for information systems research, despite its involvement with cutting-edge firms. This is primarily due to the traditionally manual nature of the industry's operations. However, we demonstrate that VC firms have begun incorporating information systems into their investment processes, thus opening new avenues for information systems research to focus on real-world use cases within the venture capital domain. Our findings suggest growing interest among VCs in collaborating with researchers, as evidenced by various studies (Cao et al. 2022; Retterath 2020; Torssell 2022; Wolff et al. 2021). This paper, which resulted from a collaboration with a leading German VC fund, further explores novel techniques to enhance investment processes through the application of the Design Science approach and machine learning techniques, producing a highly-accepted and novel solution designed to tackle the VC firm's challenges. Consequently, we encourage researchers to seek partnerships with VC firms, as both parties stand to benefit from the exploration of this emerging area of research: the potential of information systems in the VC industry.

4.8 Implications for research on Venture Capital

First, our study showcases the applicability of Design Science in the domain of VC decision support, as initially pioneered by Wolff et al. We focus on the context of startup screening systems, a previously unexplored area. Wolff et al. introduced the concept in the form of KPI dashboards for the overall VC decision support context. However, we delve deeper into the integration of Design Science into the investment process, elucidating how research can effectively address the challenges faced by venture capitalists in a practical and actionable manner. Second, we demonstrate the implementation of AI augmentation in real-world startup screening processes, designed to enhance and support the decision-making process for VCs. Third, we incorporate NLP techniques, commonly used in the Human Resources sector for evaluating resumes, into the context of VC decision support systems. This integration allows for a more efficient and comprehensive analysis of startup (founder) potential and investment opportunities.

4.9 Practical implications

From the perspective of VCs, this paper provides a guide on enhancing the screening phase using AI augmentation. We have conducted a case study involving an integrated augmentation system that connects *deal origination* with data-driven sourcing from LinkedIn. This system persistently processes and autonomously evaluates the acquired data through AI augmentation, resulting in substantial time savings and dissociating the screening time from the volume of data collected. With it, VCs can screen potential investments more efficiently while having the safety that the risk of overlooking high-potential founders is minimized. Our improved *screening* process lets VCs focus on the most promising deals, thereby increasing the probability of successful investments and enhancing portfolio performance.

From the perspective of entrepreneurs, this paper underscores the importance of a founder's background in the venture capitalist *screening* process. Although not the primary factor during initial evaluations, the founder's background serves as a crucial criterion during VC assessments. As VC firms increasingly use data-driven methods to refine *screening* processes, founders ought to keep their LinkedIn profiles up-to-date and comprehensive, or proactively reach out and engage with VCs.

4.10 Limitations and ethical considerations

Firstly, the research is based on collaboration with a single VC from Germany, which provides unique insights but could limit the study's exposure to different perspectives. We recommend collaboration with additional VC firms in future work to cover alternative investment strategies, screening processes, and user preferences. However, we believe that the presented approach is generalizable to any other VC that is scouting founders in a similarly early stage – provided they can obtain a comparable dataset of annotated CVs. Secondly, there may be limitations due to focusing on LinkedIn as our only data source. Exploring additional data could prove valuable. Finally, there is potential for latent biases in the system: Although we have removed founder names from our classifier to lessen bias, the VC's past decisions may have been biased. Features like attending top U.S. universities may reflect underlying biases, as they are not accessible to everyone. Practitioners should be aware of this when screening, especially for profiles from minority groups, and consider separate reviews for such cases. We must also recognize that the system could reinforce stereotypes or lead founders to alter their profiles to gain investor attention. It is crucial for users of this research to consider these points and look for ways to minimize these sorts of undesirable effects. We recommend two approaches to identify and reduce bias: practitioners should (1) review differences in the statistical distributions of features in the raw dataset versus the CVs with a positive annotation and investigate why those differences occur, and (2) revisit CVs with a negative annotation and check whether the founders have been successful nonetheless.

5 Conclusion

In summary, we have shown that we have successfully designed and implemented an AI-Augmented Screening System for a VC firm. Following the Design Science methodology, we have identified the relevant problems and objectives, which we utilized to adapt our system to the requirements of the users. While maintaining a strict maximum miss rate of 1% for high-potential founders, our system is able to reduce the number of deals by more than 23% by filtering non-relevant candidates – a significant increase in *screening* productivity. When a maximum miss rate of 10% is allowed, the system can save up to 6 hours of work per week per user. Our approach is a novel contribution that combines the potential of machine learning with the requirements of the system's users in order to enable a well-accepted Human-AI collaboration tool.

References

Baskoro, H., Prabowo, H., Meyliana, M., and Gaol, F. L. (2022). "Predicting Startup Success, a Literature Review," in: *International Conference on Information Science and Technology Innovation (ICoSTEC)*, vol. 1. 1, pp. 123–129.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). "Language models are few-shot learners," *Advances in Neural Information Processing Systems* (33), pp. 1877–1901.

Cambridge Associates (2020). "*US Venture Capital Index and Selected Benchmark Statistics*," Accessed: 2023-04-28. Accessible via cambridgeassociates.com.

Cao, L., von Ehrenheim, V., Krakowski, S., Li, X., and Lutz, A. (2022). *Using Deep Learning to Find the Next Unicorn: A Practical Synthesis*, arXiv preprint arXiv:2210.14195.

Chung, Y.-C., Kuo, R., et al. (2023). "A domain adaptation approach for resume classification using graph attention networks and natural language processing," *Knowledge-Based Systems* (266).

Corea, F., Bertinetti, G., and Cervellati, E. M. (2021). "Hacking the venture industry: An Early-stage Startups Investment framework for data-driven investors," *Machine Learning with Applications* (5).

Da Rin, M., Hellmann, T., and Puri, M. (2013). "A survey of venture capital research," in: *Handbook of the Economics of Finance*, vol. 2. Elsevier, pp. 573–648.

Davenport, D. (2022). *Predictably Bad Investments: Evidence from Venture Capitalists*. SSRN 4135861.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). *Bert: Pre-training of deep bidirectional transformers for language understanding*. arXiv preprint arXiv:1810.04805.

Di Giannantonio, R., Murawski, M., and Bick, M. (2022). "The Impact of Machine Learning-Based Techniques on the Scouting and Screening Processes of Early-Stage Venture Capital Firms," in: *The Role of Digital Technologies in Shaping the Post-Pandemic World: 21st IFIP WG 6.11 Conference on e-Business, e-Services and e-Society*, I3E 2022, Newcastle upon Tyne, UK, September 13–14, 2022, Proceedings, Springer, pp. 136–147.

Ferrati, F. and Muffatto, M. (2021). "Reviewing equity investors' funding criteria: A comprehensive classification and research agenda," *Venture Capital* (23:2), pp. 157–178.

Gompers, P. A. and Lerner, J. (2004). *The venture capital cycle*, MIT press.

He, M., Zhan, X., Shen, D., Zhu, Y., Zhao, H., and He, R. (2021). "What about your next job? predicting professional career trajectory using neural networks," in: *The 4th International Conference on Machine Learning and Machine Intelligence*, pp. 184–189.

Kollmann, T. and Kuckertz, A. (2010). "Evaluation uncertainty of venture capitalists' investment criteria," *Journal of Business Research* (63:7), pp. 741–747.

Krishna, A., Agrawal, A., and Choudhary, A. (2016). "Predicting the outcome of startups: less failure, more success," in: *IEEE 16th International Conference on Data Mining Workshops (ICDMW)*, IEEE, pp. 798–805.

Le, Q. and Mikolov, T. (2014). "Distributed representations of sentences and documents," in: *International Conference on Machine Learning*, PMLR, pp. 1188–1196.

Mehboob, M., Ali, M. S., Islam, S. U., and Ali, S. S. (2022). “Evaluating Automatic CV Shortlisting Tool For Job Recruitment Based On Machine Learning Techniques,” in: *Mohammad Ali Jinnah University International Conference on Computing (MAJICC)*, IEEE, pp. 1–4.

Meier, P., Beinke, J. H., Fitte, C., Behne, A., and Teuteberg, F. (2019). “FeelFit-Design and Evaluation of a Conversational Agent to Enhance Health Awareness.” In: *40th International Conference on Information Systems, ICIS 2019,* Association for Information Systems.

Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). *Efficient Estimation of Word Representations in Vector Space*. arXiv preprint arXiv:1301.3781.

Peffers, K., Tuunanen, T., Rothenberger, M. A., and Chatterjee, S. (2007). “A design science research methodology for information systems research,” *Journal of management information systems* (24:3), pp. 45–77.

Pennington, J., Socher, R., and Manning, C. D. (2014). “Glove: Global vectors for word representation,” in: *Proceedings of the 2014 conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1532–1543.

Pimpalkar, A., Lalwani, A., Chaudhari, R., Inshall, M., Dalwani, M., and Saluja, T. (2023). “Job Applications Selection and Identification: Study of Resumes with Natural Language Processing and Machine Learning,” in: *IEEE International Students' Conference on Electrical, Electronics and Computer Science (SCEECS)*, IEEE, pp. 1–5.

Pitchbook (2022). *Global Private Market Fundraising Report*. [Online; accessed 18 April 2023].

Retterath, A. (2020). *Human versus computer: benchmarking venture capitalists and machine learning algorithms for investment screening*. SSRN 3706119.

Retterath, A. and Braun, R. (2020). *Benchmarking venture capital databases*. SSRN 3706108.

Röhm, S., Bick, M., and Boeckle, M. (2022). “The Impact of Artificial Intelligence on the Investment Decision Process in Venture Capital Firms,” in: *Artificial Intelligence in HCI: 3rd International Conference, AI-HCI 2022, Held as Part of the 24th HCI International Conference, HCII 2022*, Virtual Event, June 26–July 1, 2022, Proceedings, Springer, pp. 420–435.

Roy, P. K., Chowdhary, S. S., and Bhatia, R. (2020). “A machine learning approach for automation of resume recommendation system,” *Procedia Computer Science* (167), pp. 2318–2327.

Sharchilev, B., Roizner, M., Rumyantsev, A., Ozornin, D., Serdyukov, P., and de Rijke, M. (2018). “Web-based startup success prediction,” in: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pp. 2283–2291.

Shovon, S. S. F., Mohsin, M. M. A. B., Tama, K. T. J., Ferdaous, J., and Momen, S. (2023). “CVR: An Automated CV Recommender System Using Machine Learning Techniques,” in: *Data Science and Algorithms in Systems: Proceedings of 6th Computational Methods in Systems and Software*, Vol. 2, Springer, pp. 312–325.

Sumbaly, R., Kreps, J., and Shah, S. (2013). “The Big Data Ecosystem at LinkedIn,” in: *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, pp. 1125–1134.

Teten, D. and Farmer, C. (2010). “Where are the deals? Private equity and venture capital funds' best practices in sourcing new investments,” *The Journal of Private Equity* (14:1), pp. 32–52.

Torssell, J. (2022). *The Most Influential Team Attributes When Predicting Start-up Success: A quantitative study of 25430 European new ventures*, DiVA diva2:1697289.

Tyebjee, T. T. and Bruno, A. V. (1984). “A Model of Venture Capitalist Investment Activity,” *Management Science* (30:9), pp. 1051–1066.

Van de Ven, A. H. (2007). “Engaged scholarship: A guide for organizational and social research,” Oxford University Press, USA.

Wärnberg Gerdin, L. (2022). *Predicting Success in Early-Stage Start-ups using Founding and Executive Team Characteristics*, DiVA diva2:1736961.

Weibl, J. and Hess, T. (2019). “Finding the Next Unicorn: When Big Data Meets Venture Capital,” in: *Proceedings of the 52nd Hawaii International Conference on System Sciences (HICSS 2019)*, pp. 1075–1084.

Wells, W. A. (1974). *Venture Capital Decision-Making*. Carnegie Mellon University.

Wolff, Y., Hess, M., Grichnik, D., and Wincent, J. (2021). *Looking behind closed Doors at Venture Capital Investments,* University of St. Gallen, Eprints ID 263140.