COMBINING COLLECTIVE AND ARTIFICIAL INTELLIGENCE: TOWARDS A DESIGN THEORY FOR DECISION SUPPORT IN CROWDSOURCING

Research in Progress

Rhyn, Marcel, University of St. Gallen, St. Gallen, Switzerland, marcel.rhyn@unisg.ch

Blohm, Ivo, University of St. Gallen, St. Gallen, Switzerland, ivo.blohm@unisg.ch

Abstract

Crowdsourcing represents a powerful approach that seeks to harness the collective knowledge or creativity of a large and independent network of people for organizations. While the approach drastically facilitates the sourcing and aggregating of information, it represents a latent challenge for organizations to process and evaluate the vast amount of crowdsourced contributions – especially when they are submitted in an unstructured, textual format. In this study, we present an on-going design science research project that is concerned with the construction of a design theory for semi-automated information processing and decision support in crowdsourcing. The proposed concept leverages the power of crowdsourcing in combination with text mining and machine learning algorithms to make the evaluation of textual contributions more efficient and effective for decision-makers. Our work aims to provide the theoretical foundation for designing such systems in crowdsourcing. It is intended to contribute to decision support and business analytics research by outlining the capabilities of text mining and machine learning techniques in contexts that face large amounts of user-generated content. For practitioners, we provide a set of generalized design principles and design features for the implementation of these algorithms on crowdsourcing platforms.

Keywords: Crowdsourcing, Text Mining, Machine Learning, Design Science Research

1 Introduction

With the advent of digitization and the rise of new information technologies, crowdsourcing has increasingly gained traction in research and practice as a novel approach for organizations to tap the resources of the masses (Zhao and Zhu, 2014). In crowdsourcing, an organization uses an open call to engage a broad network of people over the internet and harness their collective knowledge, creativity, or expertise for value creation activities that have previously been carried out by designated employees or contractors (Afuah and Tucci, 2012; Howe, 2006). The power of this approach lies in the aggregation of contributions from a large pool of independent individuals who voluntarily provide their solutions, suggestions, or ideas to a predefined task (Blohm et al., 2013). While crowdsourcing has been found to greatly improve the efficiency and effectiveness of problem-solving in organizations (Afuah and Tucci, 2012; Jeppesen and Lakhani, 2010), it represents a latent challenge to review and evaluate the large number of user-generated contributions – especially when they are submitted in an unstructured, textual format (Barbier et al., 2012). Processing and extracting relevant information from these data is generally described as one of the most time-consuming and cost-intensive activities in crowdsourcing (Kittur et al., 2013; Zhao and Zhu, 2014; Zogaj et al., 2014). Google, for example, required almost three years and 3’000 employees to analyze the 150’000 ideas that were submitted to its Project 10 to the 100 (Blohm et al., 2013). Similarly, IBM had to employ 50 senior executives for several weeks to assess the 46’000 ideas generated during its Innovation Jam by more than 140’000 international participants (Bjelland and Wood, 2008). Beyond such exemplary cases, Piezunka and

Dahlander (2015) analyzed longitudinal data that captured how 922 organizations responded to contributions submitted by crowds and found that organizations often "fail to harness the full potential of crowdsourcing due to inadequate filtering mechanisms" (p. 876). It shows that crowdsourcing frequently yields too much information in a form that is too complex to be evaluated in an efficient and effective manner without the support of sophisticated information processing tools.

In order to cope with the magnitude and diversity of contributions in crowdsourcing, research and practice are currently aiming to use text mining and machine learning techniques for automated information processing and decision support on crowdsourcing platforms. The ability of these algorithms to recognize patterns and extract useful information from unstructured data in a fast, scalable, and repeatable way is argued to be a key factor for the (semi-)automated analysis of user-generated content (Chen et al., 2012). So far, a number of studies have already employed text mining and machine learning algorithms on crowdsourcing platforms to cluster ideas during innovation jams (Walter and Back, 2013), to prioritize defects in crowdsourced software testing (Feng et al., 2015), or to extract named entities from crowdsourced incident reports for natural disaster response (Barbier et al., 2012). While research on crowdsourcing already comprises a number of domain-specific instantiations that demonstrate the technical capabilities and applications of text mining and machine learning algorithms, it is still lacking an overarching design theory (cf. Gregor, 2006; Gregor and Jones, 2007) that guides the systematic deployment of these algorithms on crowdsourcing platforms for adequate decision support (Zhao and Zhu, 2014). Design theories make a design problem (e.g., deploying algorithms on crowdsourcing platforms for decision support) more manageable for practitioners and provide researchers with a theoretical foundation to predict and evaluate the use patterns and impacts of related instantiations (Markus et al., 2002). As shown by Arnott and Pervan (2012), only very few studies have contributed to such theory-focused design foundations and methodologies in the past decades.

We use a design science research approach based on Hevner et al. (2004) and Peffers et al. (2008) to close this gap and answer the following research question: What design theory should guide the deployment of text mining and machine learning algorithms for decision support on crowdsourcing platforms? The expected contribution of our research is twofold. For researchers, we capture the theoretical foundation for designing decision support systems in crowdsourcing. It represents an extension to existing literature on automation in crowdsourcing that already provides concepts for assembling teams (e.g., Montesliscianí et al., 2014) or allocating tasks (e.g., Geiger and Schader, 2014). Our work also intends to contribute to decision support and business analytics research by outlining potential conceptualizations for combining the power of collective and artificial intelligence in contexts that face large amounts of user-generated content. This has been frequently requested by related literature (e.g., Lycett, 2013; Sharma et al., 2014). For practitioners, we provide a set of generalized design principles and design features for the implementation of these algorithms on crowdsourcing platforms. These decision support mechanisms may serve as additional value propositions for crowdsourcing platforms or as means to increase the efficiency and effectiveness of internal data processing.

The remainder of this paper is structured as follows: First, we present the theoretical background of our work by describing crowdsourcing and outlining the foundations of decision support on which our design theory will be grounded. Second, we describe the methodology for our study and elaborate on our design science research approach. Third, we reveal preliminary results and outline the next steps.

2 Theoretical Background

2.1 Crowdsourcing

Crowdsourcing denotes an approach in which an organization broadcasts a task that has previously been carried out by dedicated employees or contractors to a diverse network of people in an open call (Blohm et al., 2013). Compared to traditional sourcing approaches that rely on only few designated agents, crowdsourcing seeks to mobilize the resources from an independent mass of contributors

(Afuah and Tucci, 2012; Schenk and Guittard, 2011). In this vein, the approach drastically facilitates the sourcing and aggregating of information and allows organizations to benefit from a wide range of different contributions and capabilities (Geiger and Schader, 2014). Given the decentralized nature of crowdsourcing, the interaction between the organizations and their crowds generally unfolds on IT-based crowdsourcing platforms (Doan et al., 2011; Zogaj et al., 2014). On the one hand, these platforms enable organizations to allocate tasks to a crowd of distributed individuals and coordinate their activities in an efficient way. On the other hand, the platforms act as focal points for organizations to aggregate and retrieve contributions provided by the crowd.

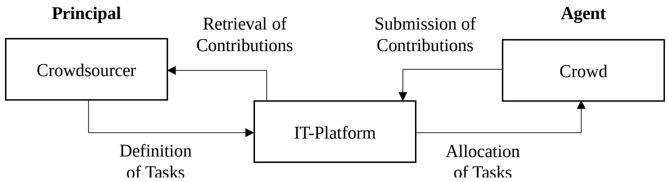

Figure 1. Basic Conceptualization of a Crowdsourcing System (based on Zogaj et al., 2014)

Crowdsourcing platforms can be hosted by the organizations themselves or by intermediaries that serve as “mediators” (Zogaj et al., 2014). In either case, the platforms represent the interface between the organizations seeking to broadcast a task and a large number of contributors willing to perform the task. Individuals working at this interface (e.g., project managers) take an important boundary-spanning role for the organizations (Tushman, 1977; Tushman and Katz, 1980). They are responsible for processing the data and *integrating* or *selecting* relevant contributions for the organizations (Geiger and Schader, 2014; Schenk and Guittard, 2011). Integration is necessary for contributions that need to be summarized (e.g., feedback) or for complementary contributions that unfold their full potential when aggregated (e.g., votes). Selection, on the other hand, is necessary when crowdsourced contributions are required to meet predefined criteria (e.g., quality requirements) or solve specific problems (e.g., “winning” ideas). However, the quantity and complexity of information in crowdsourcing make it nearly impossible for decision-makers to process the data by themselves. Especially for crowdsourcing platforms that are based on unstructured, textual contributions, it’s imperative to employ mechanisms that support decision-makers in integrating or selecting crowdsourced contributions.

2.2 Decision Support

Decision support is the area of information systems research that is concerned with supporting and improving managerial decision-making in organizations (Arnott and Pervan, 2014). The use of information technology and related decision support systems (DSS) are considered essential to this endeavor. On the one hand, they have evolved from theoretical studies offering insights on decision-making processes. Simon (1960) and Gorry and Morton (1971) contributed important groundwork for DSS by describing decision problems as existing on a continuum from programmed or structured (i.e., repetitive routine tasks) to non-programmed or unstructured (i.e., new, novel, and ill-structured tasks). On the other hand, the development of the DSS concept was influenced by technical work (e.g., Gerrity, 1971) which provided the necessary frameworks and technologies to build and understand systems capable of supporting decision-making processes (Shim et al., 2002). A classic DSS design includes components for “(1) sophisticated database management capabilities with access to internal and external data, information, and knowledge, (2) powerful modeling functions accessed by a model management system, and (3) powerful, yet simple user interface designs that enable interactive queries, reporting, and graphing functions” (Shim et al., 2002, pp. 111–112). These three major subsystems represent the basic foundations which a DSS generally comprises (Sprague, 1980). Today, there is an array of distinct types of DSS which differ with regard to their dominant technology components or drivers of decision support. They include data-driven DSS, model-driven DSS, knowledge-driven DSS, document-driven DSS, and communications-driven DSS (Power, 2008).

While traditional DSS still remain very important in organizations, current interest in decision support is predominantly focused on novel business analytics and data warehousing technologies that allow decision-makers to cope with increasing amounts of complex and unstructured data as encountered, for example, in crowdsourcing (Arnott and Pervan, 2012). Existing research suggests that cognitive limitations in human information processing may constrain decision-makers in their assessment of information in such settings (Todd and Benbasat, 1999; Eppler and Mengis, 2004). Increased information load has been found to reduce the ability of decision-makers to identify relevant information (Jacoby, 1977). Studies also show that high information load makes it more difficult for decision-makers to recall prior information or set priorities (Schick et al., 1990) and that their search strategies through data sets become less systematic and limited (Cook, 1993; Swain and Haka, 2000). DSS are designed to expand human information-processing capabilities and improve their decision-making in such settings of high information load (Todd and Benbasat, 1999). Typically, research has focused on two objectives: increasing the efficiency and the effectiveness in decision-making (Shim et al., 2002). A DSS can either support decision-makers by defining and ordering the necessary activities for decision-making (i.e., structuring the process) or by executing resource-intensive and standardizable information processing tasks (Häubl and Trifts, 2000; Silver, 1991). Decisional guidance in DSSs can take a purely informative form that includes pertinent information but no recommendations or a suggestive form with clear recommendations for the decision-maker (Silver, 1991). Furthermore, it is possible to distinguish between predefined guidance where the designer of a DSS prepares the recommendations, dynamic guidance with adaptive mechanisms that let the system learn as it is used, and participative guidance where the decision-maker defines the preferences (Parikh et al., 2001; Silver, 1991). Inherently, when decision-makers use a DSS, their decision-making process is restricted to the processes or strategies supported by the DSS (Silver, 1988; Wang and Benbasat, 2009).

Well-designed DSS have been found to help decision-makers in analyzing problems in greater depth and, ultimately, making effective decisions in a more efficient fashion (Häubl and Trifts, 2000; Hoch and Schkade, 1996). As reported by Arnott and Pervan (2012), however, more than 60% of related research in the past two decades have focused on examining specific instantiations of DSS. Only 8% of all publications have been concerned with contributing overarching constructs and models for designing and analyzing related DSS (Arnott and Pervan, 2012). In particular for novel text mining and machine learning technologies, theoretical foundations and conceptual frameworks to develop adequate DSS in contexts of unstructured data are largely missing.

2.3 Text Mining and Machine Learning for Decision Support

Text mining denotes the process of extracting useful information from unstructured, textual data through the exploration of meaningful patterns (Feldman and Sanger, 2007). These patterns are extracted by combining algorithms and methods from the fields of natural language processing, statistics, and machine learning (Tan, 1999). The standard procedure for text mining consists of two basic steps. First, the unstructured, user-generated text has to be preprocessed into a format that is compatible for machines (e.g., through tokenization and stemming). Afterwards, complementary machine learning techniques provide the means to structure the data, recognize patterns, or extract useful information. A number of supervised and unsupervised approaches are available for this task. Supervised approaches (e.g., classification; see Sebastiani, 2002) provide the means to assign contributions to predefined classes while unsupervised approaches (e.g., clustering; see Jain, 2010) are capable of automatically finding relationships and structures in large sets of contributions without predefined classes. A number of studies have already demonstrated the potential of text mining and machine learning algorithms for decision support in crowdsourcing. Walter and Back (2013) used text mining in combination with clustering algorithms to support decision-makers in selecting novel ideas from more than 40’000 contributions. Feng et al. (2015) applied clustering algorithms in crowdsourced software testing to support developers in prioritizing test reports during defect management. Barbier et al. (2012) employed text mining to automatically extract named entities (e.g., locations) from crowdsourced incident reports to assist organizations in distributing relief supplies during natural disasters. While such instantiations demonstrate the technical capabilities of text mining and machine learning algorithms in crowdsourcing, there is a lack of prescriptive design knowledge to guide researchers and practitioners in systematically implementing them for decision support on crowdsourcing platforms.

3 Design Science Research Approach

Design science research (DSR) represents a well-established approach in information systems research that is concerned with the creation of artifacts (Simon, 1996) seeking to extend the boundaries of human and organizational capabilities (Hevner et al., 2004). These artifacts may range from specific instantiations in the form of implemented software or algorithms to more abstract contributions in the form of complete design theories (Gregor, 2006; Gregor and Jones, 2007). In this study, we are concerned with the construction of a design theory for semi-automated information processing and decision support on crowdsourcing platforms based on text mining and machine learning algorithms that have been found to work well in domain-specific instantiations (e.g., Barbier et al., 2012; Feng et al., 2015; Walter and Back, 2013). A design theory can be defined as an integrated prescription that consists of a particular class of requirements, a set of effective development practices (e.g., principles), and a type of system solution with distinctive features (Markus et al., 2002; Meth et al., 2015; Walls et al., 1992). Design theories also include a set of justificatory and testable propositions grounded in kernel theories (e.g., decision support theory), which inform the construction of the artifact and provide empirically testable hypotheses regarding its utility and impact (Gregor and Hevner, 2013).

For our study, we follow the well-establish DSR process proposed by Peffers et al. (2008). This approach synthesizes the common phases of design science research proposed in existing literature (e.g., Hevner et al., 2004; Kuechler and Vaishnavi, 2008; Walls et al., 1992). It consists of the six phases of specifying the problem, defining the objectives of a solution, designing the solution, demonstrating the solution's feasibility, evaluating the solution, and communicating the results (Peffers et al., 2008). Figure 2 provides an overview of the design process and the final outcomes of our research.

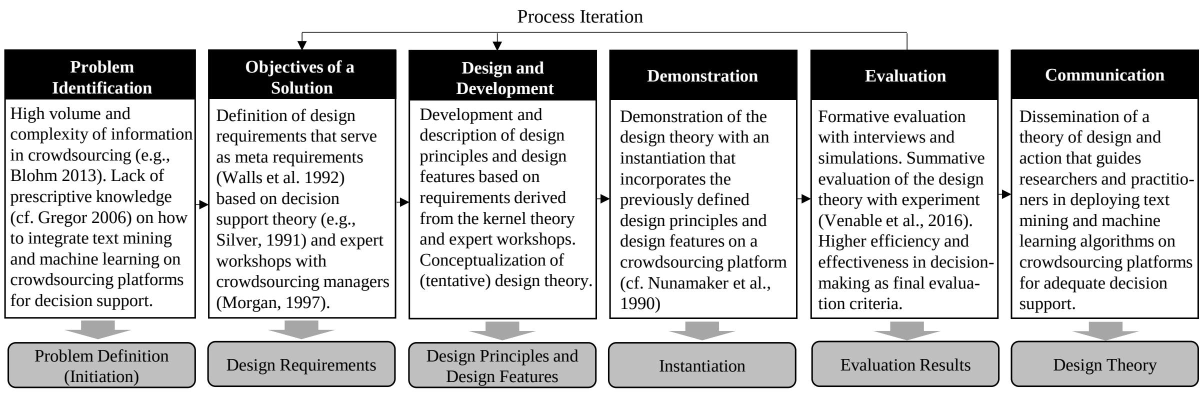

Figure 2. Design Science Research Approach (based on Peffers et al., 2008)

As design science research aims to bridge theory and practice (Holmström et al., 2009) and is inherently an iterative and incremental approach (Hevner et al., 2004), we plan to conduct at least three iterations of the design-and-evaluate cycle, which we explain in detail below, with a cross-industry research consortium (cf. Österle and Otto, 2010) comprising 8 organizations. The consortium includes 2 financial institutes, 2 insurance companies, 2 large industrial corporations, 1 multinational retailer, and 1 public transportation provider. All companies actively use crowdsourcing for software testing activities. This setting was chosen as crowdsourced software testing exhibits two characteristics which make it especially well-suited for constructing an overarching design theory in crowdsourcing. First, crowdsourced software testing comprises a broad range of tasks that yield different types of contributions. For example, functional software testing requires the crowd to contribute short and technical bug reports with ground truth. Usability testing, on the other hand, aims to elicit generative feedback and innovate ideas for new software features with no ground truth. Second, these different types of testing activities require very distinct decision-making activities from crowdsourcing managers. In functional software testing, for example, decision-makers are required to judge the severity of bug reports and prioritize them. In usability testing, they need to be able to aggregate feedback and select the most requested features for change requests. In general, crowdsourced software testing can be regarded as a "microcosm" (Leicht et al., 2016) for crowdsourcing insofar that it integrates a large variety of contributions and decision-making tasks. This ensures that our design theory is framed as generalizable as possible and may be applied in different crowdsourcing contexts that deal with textual data, such as innovation platforms (e.g., Leimeister et al., 2009) or in humanitarian aid (e.g., Barbier et al., 2012).

4 Artifact Description and Preliminary Results

In this research-in-progress paper, we present preliminary results from the first iteration of our DSR project. The first iteration of our design science research study is dedicated to an initial conceptualization of the design theory with a set of design requirements, design principles, and design features (cf. Meth et al., 2015). The design requirements serve as meta-requirements for our artifact and describe the general objectives of the design theory (Baskerville and Pries-Heje, 2010; Walls et al., 1992). They are informed by an expert workshop (cf. Morgan, 1997) that we conducted with crowdsourcing managers from our research consortium, and grounded in decision support theory, which represents the kernel theory of our work. Design principles can be defined as actionable statements that prescribe how artifacts instantiated from a design theory should be built in order to meet its requirements (Chandra et al., 2015; Meth et al., 2015). Finally, design features represent specific ways to implement design principles in an actual artifact and close the last step of the conceptualization (Meth et al., 2015). In order to validate the initial conceptualization of our design theory, we aim for a formative evaluation by conducting interviews with experts from our research consortium (Venable et al., 2016).

4.1 Objectives of the Design Theory

To elicit the problems when dealing with crowdsourced data and define the requirements for the design theory, we draw upon insights from an expert workshop that we conducted with 20 crowdsourcing managers of our research consortium. We employed a moderated focus group discussion (Morgan, 1997) over the course of 2 hours and covered three topics. We asked the participants to (1) describe the crowdsourcing process, (2) explicate the challenges they faced, and (3) outline potential avenues for improvements. We took notes during the session and clustered the responses of the crowdsourcing managers. Amongst the most frequently mentioned problems are the high amount of data generated during crowdsourcing initiatives, the diversity of information submitted by the crowd, the large share of duplicates in the data sets, as well as low quality contributions or unreliable labels. One of the insurance companies' crowdsourcing managers, for example, reviewed a total of 221 contributions over the course of 2 working days before forwarding only 21 bug reports to the developers. A large number of duplicates, issues that are out of scope, or issues that cannot be reproduced had to be manually filtered out. The crowdsourcing managers further found that such problems are aggravated by the fact that members of a large crowd often use different terms and descriptions to refer to similar issues, which made it more difficult for them to recognize duplicates or categorize usability feedback. They also emphasized that the quality of the crowdsourced contributions is of high variance. For example, a crowdtest for a website revealed that user-generated severity ratings for bug reports were incorrect in 33.6% of the cases and had to be manually verified. In general, the descriptions of the crowdsourcing managers refer to two major problems: the quantity and complexity of information when dealing with crowdsourced contributions. The former makes the evaluation time-consuming and resource-intensive. The latter induces a high information load and makes the evaluation an error-prone process. These problem descriptions resonate with the two fundamental objectives discussed in existing decision support research (see section 2.2): *increasing the efficiency* and *increasing the effectiveness* in decision-making (Shim et al., 2002). Thus, we propose these two objectives as meta-requirements for our design theory. As processing crowdsourced contributions with automated text mining and machine learning algorithms is not yet well-established and restricts the decision-making process, it is also crucial for crowdsourcing managers to *maintain the control over the degree of automation* (cf. Silver, 1988).

4.2 Design and Development

For the conceptualization of our design theory (see Figure 3), we followed Meth et al. (2015) and derived a set of design principles with corresponding design features to address the design requirements. The components emerged from the expert workshop and are grounded in existing literature.

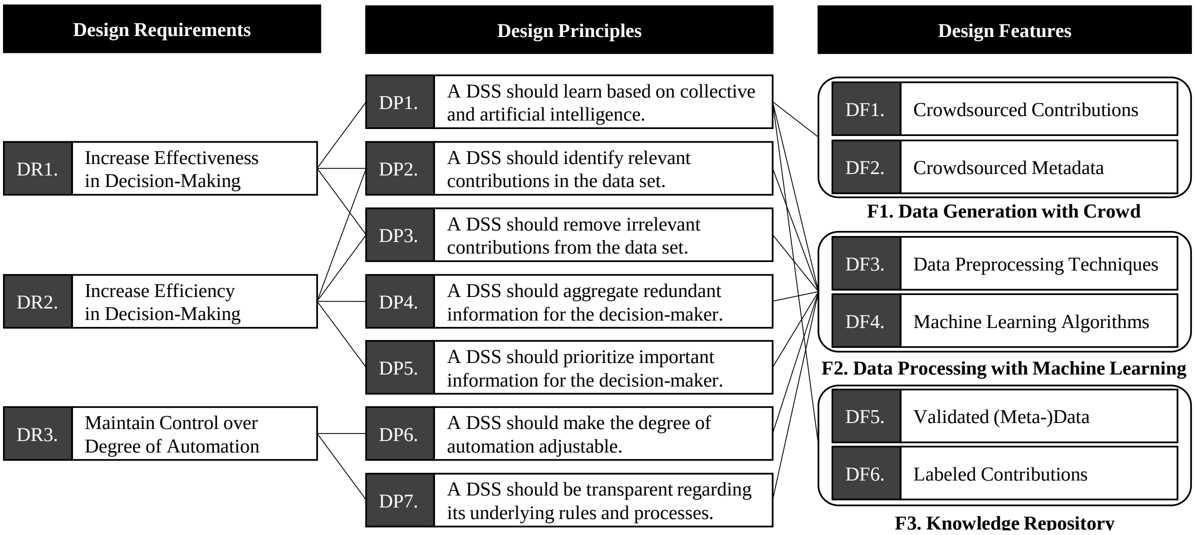

Figure 3. Preliminary Conceptualization of Design Requirements, Principles, and Features

The first three design principles (DP) for the construction of DSSs on crowdsourcing platforms aim to increase the effectiveness in decision-making. They were elicited during our expert workshop and correspond to the crowdsourcing managers' need to deal with high information load and complexity when evaluating crowdsourced data. However, given the high quantity and diversity of user-generated content in crowdsourcing, it is difficult to employ predefined recommendations that lead to better decisions for crowdsourcing managers. Instead, as specified by DP1, a DSS on a crowdsourcing platform should be designed as an adaptive system (cf. Silver, 1991) with machine learning algorithms that learn based on data generated by both the crowd (e.g., the contributions) and the decision-makers (e.g., the validated metadata). Thus, the basic concept of our design theory revolves around a combination of collective intelligence (for the generation and validation of contributions) and artificial intelligence (for the automated recognition of underlying patterns to provide decision support). As the crowdsourcing managers explained that they need support in filtering crowdsourced data and do not want to have to deal with data that induce unnecessary workload, DP2 and DP3 describe design principles to assist crowdsourcing managers in identifying potentially relevant contributions and removing irrelevant contributions with text mining and machine learning algorithms. This could be implemented, for example, with classification algorithms that serve as quality filters (e.g., Zhang and Varadarajan, 2006). DP4 and DP5 specifically deal with the high complexity experienced by the crowdsourcing managers that participated in our workshop. They need to be able to capture an overall picture of the crowdsourced contributions and select the most important ones for the organizations, which is difficult and time-consuming considering the magnitude of information in crowdsourcing. Thus, as stated by DP4-DP5, a DSS in crowdsourcing should be able to aggregate and prioritize relevant contributions to support efficient decision-making. Exemplary implementations of such mechanisms are clustering approaches capable of grouping contributions (e.g., Walter and Back, 2013). Finally, DP6 and DP7 are concerned with providing decision-makers the ability to maintain the control over the degree of automation, for example, by allowing them to choose whether the system only filters data or also suggest priorities.

The implementation of such systems on crowdsourcing platforms, as outlined by design features DF1-DF6, requires three core components that correspond to the fundamental subsystems of a DSS design (Sprague, 1980; Shim et al. 2002): a frontend that allows for the collection of crowdsourced contributions and metadata, a backend with preprocessing techniques and machine learning algorithms, and a knowledge repository with validated metadata and labeled contributions from decision-makers.

4.3 Planned Demonstration and Evaluation

For the demonstration and evaluation of our preliminary design theory, we follow the evaluation framework proposed by Venable et al. (2016). In the first iteration, we are aiming for an early validation of our design principles and design features by conducting expert interviews with crowdsourcing managers from our research consortium and external machine learning experts. These expert interviews are intended to provide a formative evaluation. The objective is to verify whether the preliminary components of our conceptualization are complete and appropriately aligned with practical needs or whether they need to be adapted. The evaluation is currently in progress and will complete the first iteration. Five semi-structured interviews with crowdsourcing managers have already been conducted.

5 Next Steps and Expected Contributions

To complete our research, further iterations of the design-and-evaluate cycle are planned. In the second iteration, we will return back to the design and development phase to integrate the results of the evaluation in our design theory. The adapted and revised conceptualization will then be used as the foundation to demonstrate and evaluate the technical feasibility of our design theory with prototypical instantiations. We will use state-of-the-art text mining and machine learning algorithms in R and Python to perform simulations (Venable et al., 2016) and experiment with different models that are intended to support crowdsourcing managers in evaluating the contributions. For classification tasks, for example, we've already successfully trained and evaluated a Random Forest algorithm to predict the quality of crowdsourced contributions and filter them based on a set of textual features (see Rhyn and Blohm, 2017). The training data for these simulations stem from crowdsourcing projects by organizations of our research consortium. We've already received platform data that include more than 4'500 labeled contributions and information on 11'000 contributors to train, validate, and test our models. We will use statistical measures, such as the precision or recall (Fawcett, 2006) as evaluation criteria.

The third iteration then aims to finalize our design theory by incorporating the results retrieved from the expert interviews and the technical simulations with our prototypes. Based on the adapted and validated design principles and design features, we build a full-fledged expository instantiation of the design theory on a crowdsourcing platform that will be used by organizations of the research consortium. This instantiation is intended to facilitate the processing of crowdsourced contributions for the organizations' decision-makers and demonstrate the theory's usefulness for constructing decision support systems in crowdsourcing. We aim for a summative evaluation of our final design theory with an experiment in naturalistic setting (Venable et al., 2016). By comparing the performance of decision-makers who are asked to manually process crowdsourced contributions with the performance of decision-makers who are supported by DSS instantiated from our design theory, the experiment aims to validate two testable prepositions (i.e., decision-makers supported by instantiated DSSs are significantly more efficient and more effective in evaluating crowdsourced data than non-supported peers).

For research on decision support and business analytics, our complete design theory is intended to capture the theoretical foundation for designing systems that combine collective and artificial intelligence in contexts that require decision-makers to evaluate large amounts of user-generated contributions. Our work also aims to serve as an extension to existing literature on automation in crowdsourcing that already provides concepts for assembling teams or allocating tasks. Finally, we aim to offer generalized design principles and design features to help practitioners in systematically deploying these algorithms on crowdsourcing platforms and leveraging the potential of this approach to its fullest extent.

References

Afuah, A. and C. L. Tucci (2012). "Crowdsourcing as a Solution to Distance Search." *Academy of Management Review* 37 (3), 355–375.

Arnott, D. and G. Pervan (2012). "Design Science in Decision Support Systems." *Journal of the Association for Information Systems* 13 (11), 923–949.

Arnott, D. and G. Pervan (2014). "A Critical Analysis of Decision Support Systems Research Revisited: The Rise of Design Science." *Journal of Information Technology* Vol. 29 No. 4, 269–293.

Barbier, G., Zafarani, R., Gao, H., Fung, G. and H. Liu (2012). "Maximizing Benefits from Crowdsourced Data." *Computational and Mathematical Organization Theory* 18 (3), 257–279.

Baskerville, R. and J. Pries-Heje (2010). "Explanatory Design Theory." *Business & Information Systems Engineering* 2 (5), 271–282.

Bjelland, O. M. and R. C. Wood (2008). "An Inside View of IBM's 'Innovation Jam'." *MIT Sloan Management Review* 50 (1), 32–40.

Blohm, I., Leimeister, J. M. and H. Krcmar (2013). "Crowdsourcing: How to Benefit from (Too) Many Great Ideas." *MIS Quarterly Executive* 12 (4), 199–211.

Chandra, L., Seidel, S. and S. Gregor (2015). "Prescriptive Knowledge in IS Research: Conceptualizing Design Principles in Terms of Materiality, Action, and Boundary Conditions." *Proceedings of the 48th Hawaii International Conference on System Sciences*, 4039–4048.

Chen, H., Chaing, R. H. L. and V. C. Storey (2012). "Business Intelligence and Analytics: From Big Data to Big Impact." *MIS Quarterly* 36 (4), 1165–1188.

Cook, G.J. (1993). "An Empirical Investigation of Information Search Strategies with Implications for Decision Support System Design." *Decision Sciences* 24 (3), 683–698.

Doan, A., Ramakrishnan, R. and A. Y. Halevy (2011). "Crowdsourcing Systems on the World-Wide Web." *Communications of the ACM* 54 (4), 86–96.

Eppler, M. J. and J. Mengis (2004). "The Concept of Information Overload: A Review of Literature from Organization Science, Accounting, Marketing, MIS, and Related Disciplines." *The Information Society* 20 (5), 325–344.

Fawcett, T. (2006). "An Introduction to ROC Analysis." *Pattern Recognition Letters* 27 (8), 861–874.

Feldman, R. and J. Sanger (2007). *The Text Mining Handbook: Advanced Approaches in Analyzing Unstructured Data*. Cambridge: Cambridge University Press.

Feng, Y., Chen, Z., Jones, J.A., Fang, C. and B. Xu (2015). "Test Report Prioritization to Assist Crowdsourced Testing." In: *Proceedings of the 10th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2015*, ACM, Lombardy, pp. 225–236.

Geiger, D. and M. Schader (2014). "Personalized Task Recommendation in Crowdsourcing Information Systems - Current State of the Art." *Decision Support Systems* 65 (2014), 3–16.

Gerrity, T. P. (1971). "Design of Man-Machine Decision Systems: An Application to Portfolio Management" *Sloan Management Review* 12 (2), 59–75.

Gorry, G. A. and M. Morton (1971) "A Framework for Management Information Systems", *Sloan Management Review* 13 (1), 55–70.

Gregor, S. (2006). "The Nature of Theory in Information Systems." *MIS Quarterly* 30 (3), 611–642.

Gregor, S. and A. R. Hevner, (2013). "Positioning and Presenting Design Science Research for Maximum Impact." *MIS Quarterly* 37 (2), 337–355.

Gregor, S. and D. Jones (2007). "The Anatomy of a Design Theory." *Journal of the Association for Information Systems* 8 (5), 312–335.

Häubl, G. and V. Trifts (2000). "Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids." *Marketing Science* 19 (1), 4–21.

Hevner, A. R., March, S. T., Park, J. and S. Ram (2004). "Design Science in Information Systems Research." *MIS Quarterly* 28 (1), 75–105.

Hoch, S. J. and D. A. Schkade (1996). "A Psychological Approach to Decision Support Systems." *Management Science* 42 (1), 51–64.

Holmström, J., Ketokivi, M. and A.-P. Hameri (2009). "Bridging Practice and Theory: A Design Science Approach." *Decision Science* 40 (1), 65–87.

Howe, J. (2006). "The Rise of Crowdsourcing." Wired Magazine 14 (6), 1–5.

Jacoby, J. (1977). "Information Load and Decision Quality: Some Contested Issues." *Journal of Marketing Research* 14 (4), 569–573.

Jain, A. K. (2010). "Data Clustering: 50 Years Beyond K-Means." *Pattern Recognition Letters* 31 (8), 651–666.

Jeppesen, L. B. and K. R. Lakhani (2010). "Marginality and Problem-Solving Effectiveness in Broadcast Search." *Organization Science* 21 (5), 1016–1033.

Kittur, A., Nickerson, J. V., Bernstein, Michael, S., Gerber, E.M., Shaw, A., Zimmermann, J., Lease, M. and J. J. Horton (2013). "The Future of Crowd Work." In: *Proceedings of the 16th ACM Conference on Computer Supported Cooperative Work, CSCW 2013*, ACM, San Antonio, pp. 1–17.

Kuechler, B. and V. Vaishnavi (2008). "On Theory Development in Design Science Research: Anatomy of a Research Project." *European Journal of Information Systems* 17 (5), 489–504.

Leicht, N., Rhyn, M. and G. Hansbauer (2016). "Can Laymen Outperform Experts? The Effects of User Expertise and Task Design in Crowdsourced Software Testing." In: *Proceedings of the 24th European Conference on Information Systems (ECIS)*. Istanbul, Turkey, pp. 1–11.

Leimeister, J. M., Huber, M., Bretschneider, U. and H. Krcmar (2009). "Leveraging Crowdsourcing: Activation-Supporting Components for IT-Based Ideas Competition.*" Journal of Management Information Systems* 26 (1), 197–224.

Lycett, M. (2013). "'Datafication': Making Sense of (Big) Data in a Complex world", *European Journal of Information Systems*, 22 (4), 381–386.

Markus, M. L., Majchrzak, A. and L. Gasser (2002). "A Design Theory for Systems That Support Emergent Knowledge Processes." *MIS Quarterly* 26 (3), 179–212.

Meth, H., Mueller, B. and A. Maedche (2015). "Designing a Requirement Mining System." *Journal of the Association for Information Systems* 16 (9), 799–837.

Monteslisciani, G., Gabelloni, D., Tazzini, G. and G. Fantoni (2014). "Skills and Wills: The Keys to Identify the Right Team in Collaborative Innovation Platforms." *Technology Analysis & Strategic Management* 26 (6), 687–702.

Morgan, D. L. (1997). *Focus Groups as Qualitative Research*. 2nd Edition. Thousand Oaks, CA: Sage Publications.

Nunamaker, J. F., Chen, M. and T. Purdin (1990). "Systems Development in Information Systems Research." *Journal of Management Information Systems*, 7 (3), 89–106.

Österle, H. and B. Otto (2010). "Consortium Research." *Business & Information Systems Engineering* Vol. 2 No. 5, 283–293.

Parikh, M., Fazlollahi, B. and S. Verma (2001). "The Effectiveness of Decisional Guidance: An Empirical Evaluation." *Decision Sciences* 32 (2), 303–332.

Peffers, K., Tuunanen, T., Rothenberger, M. A. and S. Chatterjee (2008). "A Design Science Research Methodology for Information Systems Research." *Journal of Management Information Systems* 24 (3), 45–77.

Piezunka, H. and L. Dahlander (2015). "Distant Search, Narrow Attention: How Crowding Alters Organizations' Filtering of Suggestions in Crowdsourcing." *Academy of Management Journal* 58 (3), 856–880.

Power, D. J. (2008). "Decision Support Systems: A Historical Overview" *Handbook on Decision Support Systems*. Berlin, Heidelberg: Springer, 121–140.

Rhyn, M. and I. Blohm (2017). "A Machine Learning Approach for Classifying Textual Data in Crowdsourcing" In: *Proceedings of the 13th International Conference on Wirtschaftsinformatik (WI),* St. Gallen, Switzerland.

Schenk, E. and C. Guittard (2011). “Towards a Characterization of Crowdsourcing Practices.” *Journal of Innovation Economics & Management* 7 (1), 93–107.

Schick, A. G., Gordon, L.A. and S. Haka (1990). “Information Overload: A Temporal Approach.” *Accounting, Organizations and Society* 15 (3), 199–220.

Sebastiani, F. (2002). “Machine Learning in Automated Text Categorization.” *ACM Computing Surveys* 34 (1), 1–47.

Sharma, R., Mithas, S. and A. Kankanhalli (2014). “Transforming Decision-Making Processes: A Research Agenda for Understanding the Impact of Business Analytics on Organisations” *European Journal of Information Systems* 23 (4), 433–441.

Shim, J. P., Warkentin, M., Courtney, J. F., Power, D. J., Sharda, R. and C. Carlsson (2002). “Past, Present, and Future of Decision Support Technology.” *Decision Support Systems* 33 (2), 111–126.

Silver, M. (1991). “Decisional Guidance for Computer-Based Decision Support.” *MIS Quarterly* 15 (1), 105–122.

Silver, M. S. (1988). “User Perceptions of Decision Support System Restrictiveness: An Experiment.” *Journal of Management Information Systems* 5 (1), 51–65.

Simon, H. A. (1960). *The New Science of Management Decision*. Prentice Hall: New Jersey.

Simon, H. A. (1996). *The Sciences of the Artificial*. 3rd Edition. Cambridge, MA: MIT Press.

Sprague, R.H. (1980). “A Framework for the Development of Decision Support Systems” *MIS Quarterly* 4 (4), 1–26.

Swain, M. R. and S. F. Haka (2000). “Effects of Information Load on Capital Budgeting Decisions.” *Behavioral Research in Accounting* 12 (2), 171–198.

Tan, A.-H. (1999). “Text Mining: The State of the Art and the Challenges.” In: *Proceedings of the 3rd Pacific-Asia Conference on Knowledge Discovery and Data Mining, PAKDD’99*, Beijing, China, pp. 65–70.

Todd, P. and I. Benbasat (1999). “Evaluating the Impact of DSS, Cognitive Effort, and Incentives on Strategy Selection.” *Information Systems Research* 10 (4), 356–374.

Tushman, M. L. (1977). “Special Boundary Roles in the Innovation Process.” *Administrative Science Quarterly* 22 (4), 587–605.

Tushman, M. L. and R. Katz (1980). “External Communication and Project Performance: An Investigation into the Role of Gatekeepers.” *Management Science* 26 (11), 1071–1085.

Venable, J., Pries-Heje, J. and R. Baskerville (2016). “FEDS: A Framework for Evaluation in Design Science Research.” *European Journal of Information Systems*, Nature Publishing Group, 25 (1), 77–89.

Walls, J. G., Widmeyer, G. R. and O. A. El Sawy (1992). “Building an Information System Design Theory for Vigilant EIS.” *Information Systems Research* 3 (1), 36–59.

Walter, T. P. and A. Back (2013). “A Text Mining Approach to Evaluate Submissions to Crowdsourcing Contests.” In: *Proceedings of the 46th Hawaii International Conference on System Sciences, HICSS*, IEEE, Waikoloa, Hawaii, pp. 3109–3118.

Wang, W. and I. Benbasat (2009). “Interactive Decision Aids for Consumer Decision Making in E-Commerce: The Influence of Perceived Strategy Restrictiveness.” *MIS Quarterly* 33 (2), 293–320.

Zhang, Z. and B. Varadarajan (2006). “Utility Scoring of Product Reviews.” In: *Proceedings of the 15th ACM International Conference on Information and Knowledge Management*, ACM, Arlington, pp. 51–57.

Zhao, Y. and Q. Zhu (2014). “Evaluation on Crowdsourcing Research: Current Status and Future Direction.” *Information Systems Frontiers* 16 (3), 417–434.

Zogaj, S., Bretschneider, U. and J. M. Leimeister (2014). “Managing Crowdsourced Software Testing: A Case Study Based Insight on the Challenges of a Crowdsourcing Intermediary.” *Journal of Business Economics* 84 (3), 375–405.