A Methodology to Involve Domain Experts and Machine Learning Techniques in the Design of Human-Centered Algorithms

Tom Seymoens, Femke Ongenae, An Jacobs, Stijn Verstichel, Ann Ackaert

► To cite this version:

Tom Seymoens, Femke Ongenae, An Jacobs, Stijn Verstichel, Ann Ackaert. A Methodology to Involve Domain Experts and Machine Learning Techniques in the Design of Human-Centered Algorithms. 5th IFIP Working Conference on Human Work Interaction Design (HWID), Aug 2018, Espoo, Finland. pp.200-214, 10.1007/978-3-030-05297-3_14 . hal-02264607

HAL Id: hal-02264607

<https://hal.inria.fr/hal-02264607>

Submitted on 7 Aug 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed under a Creative Commons Attribution| 4.0 International License

A Methodology to Involve Domain Experts and Machine Learning Techniques in the Design of Human-Centered Algorithms

Tom Seymoens¹, Femke Ongenaë², An Jacobs¹, Stijn Verstichel², Ann Ackaert²

¹imec-SMIT, Vrije Universiteit Brussel, Brussels, Belgium

{tom.seymoens, an.jacobs}@vub.be

²Internet and Data Lab - imec, Universiteit Gent, Ghent, Belgium

{femke.ongenaë, stijn.verstichel, ann.ackaert}@ugent.be

Abstract. Machine learning techniques are increasingly applied in Decision Support Systems. The selection processes underlying a conclusion often become black-boxed. Thus, the decision flow is not always comprehensible by developers or end users. It is unclear what the priorities are and whether all of the relevant information is used. In order to achieve human interpretability of the created algorithms, it is recommended to include domain experts in the modelling phase. Their knowledge is elicited through a combination of machine learning and social science techniques. The idea is not new, but it remains a challenge to extract and apply the experts' experience without overburdening them. The current paper describes a methodology set to unravel, define and categorize the implicit and explicit domain knowledge in a less intense way by making use of co-creation to design human-centered algorithms, when little data is available. The methodology is applied to a case in the health domain, targeting a rheumatology triage problem. The domain knowledge is obtained through dialogue, by alternating workshops and data science exercises.

Keywords: Human-Centered Algorithms, Decision Support Systems, Knowledge Elicitation Methods, Knowledge Engineering

1 Introduction

Decision Support Systems (DSSs), a set of manual or computer-based interactive tools made to support complex decision-making and problem solving, have demonstrated applicability in a multitude of realms [1–4]. Ideally, DSSs are built through a collaboration between data scientists, who build the models on historical data, and domain experts, who communicate their proficiency in order to discern the relative importance of the data features, and to tune the model parameters [5]. Collecting expert knowledge and skills entails obtaining access to both the explicit and tacit knowledge they apply in decision-making. This necessity is underscored by precedents such as the case in which a DSS was created for a hospital in order to detect patients with sepsis. The involved physician expected the system to work on the limited set of explicit

parameters he used during his consultations. However, data scientists failed to build an effective model. Subsequently, the physician was asked to classify patients solely on data given to the model. This led to poor results which convinced him that the model required previously unexpressed parameters. Yet, the collection of relevant tacit knowledge is cumbersome and time-consuming, as experts are often unaware that they even possess it, or instinctively apply it. This makes it remarkably difficult for them to verbalize the information, and for others to collect it [6]. Furthermore, the sharing of tacit knowledge is influenced by the level of trust between the two parties [7]. These factors have led scholars of information technology systems to refer to the practice as the ‘knowledge acquisition bottleneck’ in system development [8]. Recent papers in the research field have addressed the statistical model development or the co-creation interface design of the DSS. However, a detailed description of the applied co-creation method for acquiring the underlying rules for a DSS is missing in literature, although being recognized as a strenuous task [9–11].

The current paper presents a co-creation method that aims to facilitate the troublesome task of substantiating the experts’ tacit and explicit domain knowledge by visualizing it in a graph, easily interpretable by the data scientist. An efficient procedure is proposed that requires limited investment from the domain experts. In a next step, the data scientist uses the created graph to structure the data and fine-tune the data-driven developed model. The proposed method thus makes the singular modeling decisions explicit. This eventually allows the domain expert to reflect not only on the outcomes of the applied DSS, but also on the process and choices leading to the offered recommendations, hence rendering the underpinning algorithms interpretable by humans. This is of importance as granting insight into the reasoning behind an automated decision has been known to enhance the trust and uptake of such systems [12]. An additional benefit of the involvement of domain experts is that they can act as a gateway to larger historical, annotated data sets.

The remainder of the paper is structured as follows: the first section provides an overview of the applicability of DSSs and how domain experts are traditionally involved in algorithmic design. Next, we propose the interdisciplinary method of domain expert involvement for contributing their knowledge, during which we present how it was applied in order to address a triage problem in the field of rheumatology. The paper concludes with a review of the advantages and limitations of our approach and how expert knowledge can be used to create open and interpretive DSSs, while refraining from overburdening the involved expert.

2 Related Work

2.1 Decision Support Systems

Human judgement and decision-making are often suboptimal, especially in situations where the amount of information required to attain the best solution is substantial and the outcome is expected to be precise [13]. DSSs help out by structuring the cognitive process latent to the decision-making, and by granting educated access to the various required information sources [13]. Entering the fourth decade since their inception,

DSS applicability has been demonstrated in a multitude of realms, such as business, engineering and the military [1–4, 14]. As decision-making in healthcare is growing in complexity, DSSs have also been developed specifically for clinical settings [4, 11]. Notwithstanding their suitability for multiple domains and the recent fast-paced research developments in the DSS domain, their actual adoption rate remains low. One of the causes mentioned is the lack of adequate interpretability [15]. The tools made for this purpose should be amenable to verification testing, which means that developers and end users should be able to assure that the DSS is based on correct assumptions and that it operates conform to accepted domain expertise [15].

2.2 Knowledge Acquisition for DSS

Many intelligent DSS implementations are based on expert systems, or knowledge-based systems [9, 16–19]. Their foundation is established through different knowledge acquisition techniques, such as computer modeling, case based reasoning, observations and co-creation methods [10]. The next two sections provide an overview of the machine learning and social science techniques that facilitate the process, how they can benefit from each other and the challenges they continue to encounter.

Machine Learning Techniques (ML). Computing power becomes increasingly potent through advancements in data science and machine learning. This allows for the automatic inference of the required decision rules from historical data through white-box or black-box machine learning. White-box techniques, such as decision tree induction and classification rule mining, are able to give an explanation and, as previously indicated, are applicable for expert decision support within critical domains [20]. However, their predictive performance tends to be lower than that of black-box techniques, such as artificial neural networks. The latter approaches are often able to learn features automatically with higher predictive performance, but they cannot provide an explanation for their predictions [20].

A number of challenges with the application of machine learning techniques for knowledge acquisition persist. First, a prerequisite of both white-box and black-box machine learning techniques to result in sound algorithms is access to a considerable amount of historical training data, which is not always easy to obtain. Second, it has to be noted that high predictive accuracy and consistency on historical data sets do not necessarily lead to the development of proper predictive systems, as the performance can be poor when having to deal with novel, unknown or uncertain cases [21]. Moreover, integral black-box methods have as a consequence the inaccessibility of knowledge and procedure on which decisions are based and are thus incomprehensible to their developers or eventual users. When the results of a DSS are not up to par, their designers are thus unable to know which parameters should be altered or deleted altogether. Furthermore, the application domain experts and end users of these tools cannot convey decisions on blind faith, when they are unable to explain them [12]. These considerations led us to further examine how machine learning for knowledge acquisition can be supported by input from social science techniques.

Social Science Techniques. Already in 1987, it was acknowledged by Keen that the DSS field is located at the intersection of human judgement and the power of computer technology [2]. Current social science techniques of knowledge collection for DSSs are often a combination of literature review and quantitative and qualitative data collection, such as surveys and interviews [18]. The benefits of these methods are: 1) gaining access to applied decision rules that do not necessarily come forth from the historical data; 2) the explicit account of parameters enhances the interpretability at the modeling phase, which leads to greater interpretability at run-time; and lastly, 3) the involved domain experts might act as a gateway to historical (training) data. However, it is not difficult to understand why knowledge acquisition is still often considered a bottleneck in the development of DSSs. The unwillingness to share due to mistrust, and the unavailability of experts with sufficient time are well-known obstacles [5]. Furthermore, the externalization of the tacit knowledge that domain experts gather through experience, and apply in a subconscious manner, remains a challenge as it is often a non-organized part of the experts' intuition [5, 9, 17, 22]. To this end, the use of metaphors and narratives is prescribed in order to facilitate this externalization. However, the exact course of action is not specified [9]. Moreover, literature mentions how tacit knowledge is often only revealed during observations when a prototype of the DSS is already available [22].

The methodology disclosed in the current paper seeks to overcome these drawbacks by providing an efficient means to capture tacit knowledge before a prototype is developed. It is part of a larger evolution investigating how machine learning and social science techniques can complement each other for knowledge acquisition in the development of DSSs [10]. The method consists of a succession of generative workshops with explicit probing and project techniques, designed to obtain latent and tacit knowledge [23]. Involvement of domain experts in this manner has been recognized to reduce the chance of overlooking crucial data [18]. More specifically, this paper aims to address the general lack of an explicit presentation of the detailed protocol used for knowledge acquisition in DSS design in literature [10].

3 A Methodology for Knowledge Acquisition for DSS Design

The present paper proposes a methodology for the materialization of expert knowledge through a succession of qualitative workshops. The method is particularly applicable towards the development of the knowledge base in a variety of domains for which the decision-making process is susceptible to automated support. Each step of the protocol is rigorously described with an immediate illustration of how it is implemented in a health use case, regarding the support of triage in the field of rheumatology. Figure 1 provides an overview of the protocol.

Fig. 1. Overview of the developed methodology

3.1 Step 1: Mapping of Context and Motivation for DSS Development

The need for a DSS can be propagated in different ways: the research team can observe the benefits themselves and take it up in subsequent explorations with relevant experts (top-down), or domain experts can voice their concerns and contact the research team with their request (bottom-up). In both cases, it cannot be assumed that the researchers fully understand the context, frustrations and motivations in the specific domain.

After an introductory meeting with domain experts, a preliminary literature review and context mapping is performed. However, an exhaustive understanding is unattainable, and can even be counterproductive for the course of the development process, as it undermines the need for detailed descriptions from experts. Nevertheless, a basic understanding is imperative in order to decide upon 1) the minimum number of workshops, 2) the required domain expertise and characteristics of the attending participants, 3) the number of participants in each workshop and, 4) the mix of participants in each workshop. The participants should have different experiences and should preferably not be part of the same team, as this would result in shared implicit knowledge that would not be clearly articulated.

Application in the Rheumatology Use Case. A rheumatologist who witnessed the potential utility of a DSS contacted the research team and briefly explained his motivation and the context in which he operates as follows.

Rheumatologists in Flanders, Belgium have to cope with an ever-increasing workload. This is due to the growing number of patients per expert, as a result of an aging population and a decrease in active rheumatologists, as the older generation is retiring and new doctors specializing in the field are scarce. The rheumatologist had strong beliefs that a digital triage system, supporting General Practitioners (GPs) in their diagnosis, could remove some of the burden from the medical specialist, if able to reduce the number of misguided patient referrals. As a first step, he enables a number of GPs to fill in a digital checklist for patients who were suspected of suffering from rheumatism. Data for 127 patients was collected. As a second step, the rheumatologist

went through the data to determine if the specific patients were eligible to get an urgent appointment. Although successful, the post-analysis by the professional was too labor-intensive to be practical in a real-life setting.

In a first meeting with the research team, it was established that the logical step forward was to extract the knowledge rheumatologists applied in evaluating such digital checklists. It was decided that three workshops would be conducted: one with rheumatologists affiliated with a university, one with private rheumatologists, and one with general practitioners. The objective of the workshops was to identify the patients who were eligible for referral according to each party involved, which rules they apply in their diagnosis, and which data sources should be accessible. A fourth workshop, containing a mix of the different participants, would subsequently be conducted in order to validate the extracted knowledge.

3.2 Step 2: Gathering Historical Data and Explicit Knowledge

The inclusion of domain experts from early on in the creation of a DSS, before a prototype is conceived, is a valuable asset as they can act as a gateway to annotated historical data. During the initial meeting, it is discussed if they know or are in the possession of annotated data. Moreover, the experts can refer to easily collectable, explicit knowledge they apply in the course of the decision-making process. This consists of the decision rules and data sources experts can effortlessly identify when inquired about the information they base their day-to-day judgements upon. These are in subsequent steps used for both the creation of a baseline DSS and the further development of the co-creation workshops.

Application in the Rheumatology Use Case. As the rheumatologist declared in the initial meeting that he was already in possession of a digital checklist and referral data for 127 patients, the research team inquired for an anonymized version of this labeled training data. This file was based on the already-existing screening sheet with patient data and the evaluation of the expert on whether they were urgently to be referred to a rheumatologist.

3.3 Step 3: Preparation and Application of the Labeled Training Data in Machine Learning (ML)

The information on knowledge application that was gathered in the second step allowed for the creation of a baseline DSS. The objective was to discover how successful the DSS is when created by purely using different data-mining algorithms on the available historical annotated data. Gathered training data is often not directly useable to its full extent as input for machine learning, e.g. due to inconsistent input of parameters such as variable data formats, open text fields that allow unstructured input and missing values. Therefore, an initial data cleaning is often required. Even when the results of this initial ML exercise are not convincing, the examination of the data and the creation of a baseline DSS remain important: not only does this give a baseline performance

level by which future enhancements through the adoption of expert knowledge can be measured, but it also serves as an incentive for the domain experts to accept that an interdisciplinary approach, such as the one presented in this protocol, is beneficial.

Application in the Rheumatology Use Case. Based on the anonymized patient data file provided by the rheumatologist, a baseline DSS was built. The expert emphasized the need for an interpretable DSS which meant that it needed to be capable of giving an explanation regarding which process had led to its conclusions. This was required to gain the trust of the domain experts, i.e., the rheumatologists and GPs, in the designed system. The physicians will only be comfortable to use the DSS and follow its recommendation, if they fully understand how the model works. Moreover, it allows the rheumatologist to understand why the patient was referred, so that the urgency of the case can be assessed in deciding upon an appointment date. Therefore, white-box machine learning algorithms were put forward to design the DSS. The most well-known white-box algorithms are decision tree induction and rule learning [24, 25]. Based on its wide successful adoption in practice, the See5 toolkit was used for the rheumatology case [26–28]. It uses the C5.0 algorithm to manufacture decision trees or collections of if-then rules. Several classifiers (either decision trees or rule sets) are generated rather than just one. When a new case is to be classified, each classifier votes for its predicted class and the votes are counted to determine the final class.

As a first step, the data file provided by the rheumatologist needed to be cleaned. This was performed based on the input gathered during the initial meeting with the expert. As a result, 215 different, sometimes overlapping, items on the checklist were derived as input for the DSS. These consisted, for example, of an indication of all the joints and whether they were swollen or painful to the patient, the medical history of the patient, his/her medication usage, results of laboratory tests and medical imaging, etc. An initial filtering of unimportant items, e.g., administrative features, was performed to reach a set of 188 items. See5 generated DSS algorithms with an error rate ranging between 29.36% and 53.8%, when using two separate disjoint datasets for training and evaluation, and an error rate between 7.1% and 14.19% when the complete set was used for training and validation. This was clearly due to the fact that the set of 188 items per case was too large compared to the number of available cases, i.e. 127 patients. Therefore, unconvincing results were obtained, even when trained and executed on the same set. Additional domain knowledge was required in order to reduce the number of items.

3.4 Step 4: Design of the Co-Creation Workshops

As an intermediate step, the social scientists involved perform an investigation on probing material and adapting generative techniques aimed at extracting domain expertise. Through a set of workshops, the implicit knowledge the domain experts possessed is extracted and tested. The challenge lies in allowing the participants to share their experiences and intermittently ask them pertinent questions as a non-specialist in the field. This inquiry by an inexperienced party may determine the participants to tell

their stories in a precise manner. Different tools can be used to help convey the descriptions, such as text, narration and visualizations. Each workshop has its specific set-up and lasts for a maximum of two hours. Two different types of workshops can be distinguished: a first set where the knowledge of the experts is gathered and systematized (acquisition workshops), and a second set in which the resulting flowcharts are validated (validation workshop). In the following sections we will elaborate on these, as well as the sequence of different exercises that constitute each session. In general, a preparatory step is required in order to discuss and structure the workshop format. In general, they should contain the different elements enumerated in the following sections. The material used can vary according to the different object of expertise. The specific workshops that are discussed below were designed in close collaboration with the data scientist.

3.5 Step 5: Acquisition Workshops

First Section: Introduction. In the first part of the Acquisition Workshops (AWs), the structure and objectives are introduced to the participants. This component does not differ from that of other interpretative social science research methods. The essential steps required are:

- The organizers present themselves, as well as their roles in the workshop as presented in Table 1.
- The project is shortly presented as well as the contribution of the current workshop to its goal.
- The structure and objectives of the session is explained.
- If the session is being recorded, the participants are notified of this, as well as of what will happen with the recordings.
- An informed consent is provided and signed by all participants. This document describes the aims of the session, what will happen with its results and the possible risks for the participants.
- The different participants and organizers present themselves and their background.

Table 1. The different roles of the organizers during the workshop

Expertise	Roles	Applied Reasoning
Data Scientist	Architect of the DSS	Deduction
	Schematizes the articulated knowledge of the domain expert	
Social Scientist #1	Facilitator of the workshops Inquires until all relevant knowledge surfaces	Induction
Social Scientist #2	Note taker	N/A
	Brings attention to missing information	

Second Section: Experience Case Description. Each participant is asked to write down in a few lines a specific case from their experience, for which the decision-making process was intriguing or peculiar. Every expert subsequently shares the case with the other participants. The objective is to steer the discussion away from an abstract level towards more practical and experience-based examples. After each attendant discloses their case to the group, one of them is selected through a short discussion led by the social scientists. This example will lead the rest of the expert consultation, as a common ground on which the participants can share their implicit knowledge. The particularity of the case can for example be attained by pursuing cases where the decision was difficult to make, a bad decision was eventually made or the urgency of the case was underestimated. Past experience has shown that looking for extreme examples contributes to a more animated discussion.

Application in the Rheumatology Use Case. The rheumatologists were asked to write down three different patient cases: a patient that was wrongfully referred to them, a patient that was referred too late, and a patient that was correctly and timely referred. In the alternate workshop, the general practitioners were asked to write down one case in which they had difficulty in deciding if the patient should see a rheumatologist. Afterwards the participants presented their cases to the group with an emphasis on the particularity of each case. After a short discussion, the most thought-provoking patient case was selected, which would be used throughout the workshop.

Third Section: Playing the Omnipotent DSS. In this phase, the experts have to work together and take up the role of an all-knowing and powerful DSS. This fictional system has access to every conceivable information source. The objective of this exercise is to define what information and knowledge are crucial to be considered and in what order it should be requested by the system. The questions that should be repetitively asked by the social scientist facilitating the process are:

- What knowledge do you need (next) to make an educated decision?
- Where do you get this information from?
- Why do you feel this information is pertinent?
- Can you give examples of the knowledge? How do they relate to one another?
- Does everyone agree with this?

Simultaneously, the data scientist takes on the role of the system architect who visualizes the thought processes on large papers. The different information components indicated by the participants are compartmentalized, related to each other and given an hierarchy based on importance. Imperative to this process is a mutual recognition and reciprocal communication flow between facilitator and DSS architect. The facilitator aims to assist the DSS architect by abstracting from the situational and working towards a consensus. When saturation is achieved as no new information components are distinguished, the exercise enters a new phase. The other experience cases, described by the participants but not yet used, are run through the designed omnipotent DSS. This

allows to question the structure, components and sequence as provided by the DSS architect.

Application in the Rheumatology Use Case. In the use case, the facilitator sketched the situation as follows: a patient arrived at a general practitioner, being supported by a DSS. Together the participants had to take up the role of this DSS and decide on whether the patient needed to be immediately referred to a rheumatologist. They had to indicate their information needs during the different stages of the patient visit. The rheumatologist or GP that presented the case needed to provide the information required by her/his colleagues. After each request for information, the facilitator enquired about its importance and at what time the DSS may need this particular piece of information. During this process, the data scientist schematized the whole process. If something was not possible to schematize or a conflict appeared on the existing schema, this was mentioned. In case something was missed, the note-taker brought this to the attention of both the facilitator and the data scientist. After each discussion, the facilitator asked what the follow-up question would be that the system needed in order to make a meaningful referral. This process was repeated until the participants, and thus the DSS, had gathered all of the information they needed.

Fourth Section: Synthesis and Evaluation. The process of the workshop is reviewed in a joint effort with the participants. This starts with a recapitulation of the explicit knowledge available at the start of the session and how this is supplemented by the tacit knowledge that became available. At this moment, an inquiry is made of the information and procedures that are still lacking. Some time is also devoted to the evaluation of the process itself: do the contributors feel that they got to voice their concerns and that their participation was worthwhile?

Application in the Rheumatology Use Case. This final part of the AW started with running through the diagram of different patient cases from the start, in order to see whether all required information needs had materialized. Once everything was included, the original digital screening sheet was brought to the table in order to help identify the differences with the new diagram and to look for an explanation for these differences. Afterwards some time was devoted to the evaluation of the workshops.

3.6 Step 6: Formalization and Issue Detection

After each acquisition workshop, the resulting flowcharts are formalized and structured, with extra input from the transcripts of the note-taker, as some information may have been overlooked during the session itself.

After the finalization of all AWs, the data scientist compares the different flowcharts and integrates them into one consolidated diagram, containing all the different concepts and flows that had been previously indicated (See Fig. 2). Ambiguities and experienced missing links that should be taken into account in the validation workshop (VW) are noted down in the course of this process.

Fig. 2. Part of the integrated diagram of the rheumatology case.

3.7 Step 7: Validation Workshop

Objective of the Validation Workshop. The objective of the validation workshop (VW) is to run through the integrated flowchart and to take away any existing ambiguity that still resides over the priority and necessity of the different questions that the DSS needs to take into account. It also re-evaluates the information sources it needs to address in order to formulate a successful outcome. If the AWs were conducted with different types of stakeholders, a sub-selection with equal representation of each party needs to be present in the VW. The different steps that are undertaken in the VW are summarized below.

Application in the Rheumatology Use Case. The participants of the VW were two rheumatologists and two general practitioners.

First Section: Introduction. The intent and procedure of the workshop is introduced in this section. This process is similar to the introduction of the AW.

Second Section: Prioritization of Components. The result of the AWs is an integrated flowchart, which lists the main information components that are necessary to formulate an informed decision. As preparation for the VW, the different components are listed. In this first part of the workshop, the participants are asked to go through the list individually and give a priority rating (Low/Neutral/High) to each of the information needs. This exercise is repeated until all components have been rated. The sheets of the different responders are gathered, processed and integrated by the data scientist and note-taker. The result of this process is presented to participants and discussed in group until a consensus about the priority level of each of the sub-questions is reached.

Application in the Rheumatology Use Case. In the rheumatology case, it became apparent that two main information needs existed: inquiries about pain and swollen joints. These main components contained several sub-components which were listed on two different files that were shared with the participants. All participants went through these questions individually and ranked them according to their importance (Low/Neutral/High). Afterwards, the results of this exercise were discussed publicly. This exercise resulted in the deletion of some of the sub-questions.

Third Section: Experience Case Description and Playing the Omnipotent DSS. Similar to the AW, the stakeholders are asked to describe a case for which the decision-making process was peculiar. These cases are then used to jointly run through the integrated flowchart, making use of the priority levels that were obtained during the previous exercise. In an ideal situation, the use of the flowchart should guide the participants towards the right conclusion for each of the cases. If this is not the case, the possible reasons for the drawbacks are explored until a new consensus is found.

Application in the Rheumatology Use Case. In the rheumatology case, the participants were asked to write down an experience case. Similar to the AWs, the rheumatologist had to describe a case that was wrongfully or tardily referred to them. The attending GPs had to describe a difficult case to diagnose. One of them was chosen to run through the integrated flowchart. New pitfalls in the flowchart were discovered and discussed in the group.

Fourth Section: Synthesis and Evaluation. The validation workshop ends similar to the acquisition workshops, with an overview of the results and an open discussion about the followed procedure.

3.8 Step 8: Optimized Machine Learning based on Expert Involvement

Not all data engineering enhancements done in step 3 resulted in better performance of the DSS. Therefore, it is of utmost importance to apply the correct re-engineering principles in the correct situation, something which can be achieved with the application of domain expertise. The presented protocol allows for this information to be extracted, shared and correctly applied. Based on the expert knowledge, data engineering principles can be applied to reduce the number of items into more practical input features in order to achieve a more accurate result. Afterwards, the new results are compared with the baseline DSS created in step 3.

Application in the Rheumatology Use Case. First, some continuously valued items (e.g. blood pressure values) were converted into value portioned ones, based on boundaries indicated during the workshops. Second, labeled data that was never mentioned during the workshops was discarded (e.g. the medication history). Third, some items were too detailed, according to the domain experts, and these could be clustered into a single input item. An example here is the localization of joints that could be clustered into six categories: left/right hand, left/right foot and left/right side of the body. Different weights were also attached to the components (e.g. morning stiffness was given a higher importance). Similarly, misclassification costs were given a weight. Under-diagnosis was for instance deemed much worse than over-diagnosis. However, it has to be noted that some of the parameters that were deemed important by the domain experts were not included in the labeled data set. This resulted in the newly developed DSS not having access to all of the required information.

The previously discussed baseline DSS for the rheumatology case achieved error rates between 29,36% and 53,8%. After expert involvement, a best-case DSS was achieved with an error percentage of 22.8%. As mentioned, the DSS is a combination of various decision trees or rule sets that achieve the final result through majority voting. Through a manual check-up of all these trees and sets, it was possible to find a DSS algorithm which performed 100% correct. However, it has to be kept in mind that this was achieved with only a limited dataset at hand and that the inclusion of more diverse data would probably lead to errors in that particular DSS algorithm as well.

4 Discussion

Since the inception of DSSs, the importance of domain expert inclusion in the constitution of a knowledge base has been underwritten [2]. However, to our knowledge, there is a lack of information regarding the specifics of co-creation protocols guiding knowledge acquisition [10]. In general, the externalization of tacit knowledge is said to be enabled by the application of narratives or metaphors, but a description of an approved method is missing [9].

The current paper addresses these issues by clearly presenting a methodology to unravel, define and categorize the implicit and explicit domain knowledge, through an alternation of co-creation workshops. The objective is to enable the discovery of

underlying skills and know-how as they are provided through dialogue by professionals in the field. This interpretability at the modeling phase is more likely to create interpretability at use time in comparison to black-box machine learning methods. Unlike previous attempts, the proposed methodology captures the domain knowledge before a prototype of the DSS is developed. This allows experts to exert a higher influence on the design choices.

The methodology was applied in a health case targeting a triage problem in the domain of rheumatology. The preliminary results have been encouraging as both the tacit and explicit knowledge of domain experts could be extracted and formalized. However, it should be stated that it is impossible to know if all relevant tacit knowledge was extracted during the succession of workshops. During the workshops, no mistrust towards the researchers was noticed. The participants stayed engaged throughout the process as long as it did not last longer than the envisioned two hours. Although some participants showed some initial reserve, this changed throughout their involvement in the group discussions and as they witnessed the rules materialize in the drawings of the data scientist. The results strengthen the authors' conviction that the use of narration through the experience case description exercise as well as the simultaneous diagram creation can be used to extract knowledge, also in other domains than DSS creation, e.g. organizational information sharing.

The domain expert involvement guided the subsequent data engineering. An initial comparison of the baseline DSS with the one created after the involvement of the domain expert had some promising results. However, due to the data set's small size, as only data on 127 patients was present, and because not all required information was available in the initial labeled training set, the study requires more future validation.

This method is only one example; we would like to stimulate the community to develop other ways of involving the domain expert in knowledge capture processes in a pragmatic and efficient way.

Further validations of the applied methodology in alternative application domains are necessary. Additional research is also required in order to see how the aim of interpretability at run-time can be implemented so that the end-user receives insight into why specific decisions are made on their behalf. Moreover, if such a system is used by more people, it might become beneficial to add a feedback loop where users can indicate themselves which parameters should hold more weight in the decision-making process, so that the knowledge base on which the system operates is adapted to new findings.

5 Acknowledgement

We would like to thank Klaas Vandevyvere, MD, for his support in this research, both as project initiator and rheumatologist.

6 References

1. Arnott, D., Pervan, G.: Eight key issues for the decision support systems

- discipline. *Decis. Support Syst.* 44, 657–672 (2008).
2. Keen, P.G.W.: Decision support systems: The next decade. *Decis. Support Syst.* 3, 253–265 (1987).
3. Shim, J.P., Warkentin, M., Courtney, J.F., Power, D.J., Sharda, R., Carlsson, C.: Past, present, and future of decision support technology. *Decis. Support Syst.* 33, 111–126 (2002).
4. Stivaros, S.M., Gledson, A., Nenadic, G., Zeng, X.J., Keane, J., Jackson, A.: Decision support systems for clinical radiological practice - Towards the next generation. *Br. J. Radiol.* 83, 904–914 (2010).
5. Gebus, S., Leiviskä, K.: Knowledge acquisition for decision support systems on an electronic assembly line. *Expert Syst. Appl.* 36, 93–101 (2009).
6. Haase, T., Termath, W., Martsch, M.: How to save expert knowledge for the organization: Methods for collecting and documenting expert knowledge using virtual reality based learning environments. *Procedia Comput. Sci.* 25, 236–246 (2013).
7. Holste, J.S., Fields, D.: Trust and tacit knowledge sharing and use. *J. Knowl. Manag.* 14, 128–140 (2010).
8. Hoffman, R.R., Shadbolt, N.R., Burton, A.M., Klein, G.: Eliciting knowledge from experts: A methodological analysis, (1995).
9. Becerra-Fernandez, I., Rajiv Sabherwal: *Knowledge Management Systems and Processes.* (2010).
10. Wagner, W.P.: Trends in expert system development: A longitudinal content analysis of over thirty years of expert system case studies. *Expert Syst. Appl.* 76, 85–96 (2017).
11. Jeffery, A.D., Novak, L.L., Kennedy, B., Dietrich, M.S., Mion, L.C.: Participatory design of probability-based decision support tools for in-hospital nurses. 0, 1–10 (2017).
12. Ribeiro, M.T., Singh, S., Guestrin, C.: “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. (2016).
13. Druzdzel, M.J., Flynn, R.R.: Decision Support Systems. *Encycl. Libr. Inf. Sci.* 1–15 (2002).
14. Dalinger, E.: A framework for design of an integrated system for decision support and training. *Proc. 31st Eur. Conf. Cogn. Ergon. - ECCE '13.* 1 (2013).
15. Lisboa, P.J.G.: Interpretability in Machine Learning - Principles and Practice. In: Masulli, F., Pasi, G., and Yager, R. (eds.) *Fuzzy Logic and Applications. 10th International Workshop, WILF 2013.* pp. 15–21. Springer, Genoa, Italy (2013).
16. Power, D.J.: A Brief History of Decision Support Systems, <http://dssresources.com/history/dsshhistoryv28.html>, (2007).
17. Garcia-Taylor, M.C.: Development of a Knowledge-Based Framework for Demand Management for Refrigerated and Shelf-life Constrained Food Supply Chains, (2016).
18. Padma, T., Balasubramanie, P.: Domain experts’ knowledge-based intelligent decision support system in occupational shoulder and neck pain therapy. *Appl. Soft Comput. J.* 11, 1762–1769 (2011).

19. Gai, Y., Dang, Y., Xu, Z.: A Methodology for Problem-driven Knowledge Acquisition and its Application. In: *Communications in Computer and Information Science*. pp. 15–29. Springer, Singapore (2016).
20. Vandewiele, G.: Enhancing white-box machine learning processes by incorporating semantic background knowledge. *Lect. Notes Comput. Sci.* (including Subser. *Lect. Notes Artif. Intell. Lect. Notes Bioinformatics*). 10250 LNCS, 267–278 (2017).
21. Turban, E., Volonio, L., Mclean, E., Wetherbe, J.: *Information technology for management: Transforming organizations in the digital economy*. John Wiley & Sons, New York (2009).
22. Lindgren, H.: Decision Support System Supporting Clinical Reasoning Process – an Evaluation Study in Dementia Care. *eHealth Beyond Horiz. – Get IT There*. 136, 315–320 (2008).
23. Sanders, L., Stappers, P.J.: *Convivial Toolbox: Generative Research for the Front End of Design*. BIS Publishers, Amsterdam (2013).
24. Kotsiantis, S.B.: Supervised Machine Learning: A review of Classification Techniques. *Informatica*. 249–268 (2007).
25. Kuhn, M., Johnson, K.: *Applied Predictive Modeling*. Springer (2013).
26. Quinlan, J.R.: Improved Use of Continuous Attributes in c4.5. *J. Artificial Intell.* 77–90 (1996).
27. Quinlan, J.R.: *C4.5: Programs for Machine Learning*. Kaufmann Publishers (1993).
28. *Is See5/C5.0 Better Than C4.5?* Springer (2013).