ONTOLOGIES AND KNOWLEDGE BASES: A NEW WAY TO REPRESENT AND COMMUNICATE VALUES IN TECHNOLOGY DESIGN

Kathrin Bednar, Till Winkler

Institute for Information Systems and Society,
Vienna University of Economics and Business (Austria)

kbednar@wu.ac.at; till.winkler@wu.ac.at

ABSTRACT

A promising pathway towards an ethically aligned design of technology is to consider human values such as “privacy” and “autonomy” in the design process. Value-oriented approaches have inspired the development of several unique methods for identifying stakeholders, values, and design requirements. However, none of these methods focus solely on the representation of values and the communication of value knowledge. In this paper, we outline a methodology that is inspired by techniques from the semantic web community. Based on a case study, we demonstrate how building an ontology can be used to represent and visualize value knowledge. The underlying empirical data comes from a sample of students who applied Value-based Engineering to elicit and analyse values for a telemedicine communication system. In spite of the specifics of this case study, the techniques for representing and communicating the resulting value data are compatible with any value-oriented approach. Furthermore, they support quality criteria that are essential when dealing with values both in research and design. The formal representation of value knowledge in form of semantic data ensures a high level of detail while respecting context-specificity. The underlying ontology helps to represent key concepts and their relations and supports the transparency of data analysis, including the initial coding of the value-related data. The resulting knowledge base can be shared with stakeholders and researchers, supporting the joint evolution of value-oriented approaches and technology.

KEYWORDS: values, design, engineering, ontology, knowledge base, semantic web.

1. INTRODUCTION

Since the advent of the internet, the variety of devices and applications has kept increasing. Information technology helps us to structure and organize our everyday life, but also shapes our work and social lives. Traditionally designed technological products have been optimized mainly for functionality, ignoring that functionality also depends on non-functional characteristics (Chung & do Prado Leite, 2009). Ignoring high-level non-functional characteristics during the design process can lead to harmful effects with ethical implications, such as information distortion in the form of search engine manipulations (Epstein & Robertson, 2015), filter bubbles (Pariser, 2011), and algorithm biases (O’Neil, 2016). Consequences often play out at the societal level, as in the example of social media’s impact on democracy (Cadwalladr, 2017), but there are also physiological effects for individuals, which can be observed in the form of symptoms of stress and depression (Barley, Meyerson, & Grodal, 2011). Technology is a mediator for biases and human values, it moulds its use context and changes the perceptions and actions of people to the point where it creates new practices and forms of living (Verbeek, 2008). In this view, negative or positive effects do not emerge as a result of technology use, but are triggered by the affordances *inherent* in technological artefacts, systems, and infrastructures (van den Hoven, 2017).

For this reason, designers, researchers and engineers need to address the potential effects and consequences of a technology throughout its design and development process. This is especially important, as many engineers are willing to go beyond traditional functional requirements but do not have the necessary time or autonomy to implement them (Bednar, Spiekermann, & Langheinrich, 2019; Spiekermann, Korunovska, & Langheinrich, 2018). This situation might change if the consideration of values is incorporated into the design and development process of technologies as for instance Value-based Engineering requires developers to take the necessary time to think about values (Spiekermann & Winkler, 2020). A promising pathway towards an ethically aligned design of technology is to consider human values such as “wellbeing”, “privacy”, “security”, and “autonomy”.

Considering human values during system development minimizes biases by making the system more accessible to a greater diversity of users, leads to more desirable software, and increases the likeliness of new technology to be adopted (Friedman & Nissenbaum, 1996; Isomursu, Ervasti, Kinnula, & Isomursu, 2011; Spiekermann, 2016). While there is a diverse landscape of value-oriented approaches and methods (e.g. Friedman & Hendry, 2019; Spiekermann, 2016), a common framework that focuses on their commonalities is still missing. However, a common way to represent value knowledge (e.g. in form of lists or networks) could lead to a better understanding of value knowledge gained in a value-orient project and make it easier to share this knowledge *among* team members as well as *across* different value paradigms. Ideally, such a framework would support quality criteria such as transparency and preserve context-specificity, e.g. by including information on the technology under investigation or the affected stakeholders. Considering challenges of value-oriented design processes such as keeping a high level of detail when representing original data throughout the design process, such a framework could benefit the further development of value-oriented approaches. Improving the capability to maintain quality criteria could even increase the recognition of value-oriented design outside academia (Detweiler & Harbers, 2014; Miller, Friedman, Jancke, & Gill, 2007).

To the best of our knowledge, there is no method that focuses solely on the representation and communication of values independent from the underlying theoretical background (e.g. the definition of values) or the specific method used (e.g. for eliciting values). In this paper, we propose that ontology-building and the development of a knowledge base, techniques from the semantic web community, can step in here. In the following, we take a closer look at the challenges of value-oriented design processes. Then, we explore the different phases that a value-oriented project would need to run through to develop an ontology and build a knowledge base that can be shared with project members and other researchers. We present and discuss a case study on a telemedicine communication system to illustrate these steps and discuss benefits and future potentials of this method. While the case study dataset resulted from the elicitation and analysis of values in accordance with Value-based Engineering, the proposed methodology for value representation and communication is compatible with any value-oriented approach.

2. CHALLENGES OF VALUE-ORIENTED APPROACHES TO TECHNOLOGY DESIGN

The most prominent approach for considering human values during technology development is called *Value Sensitive Design* (VSD; Friedman & Nissenbaum 1996, Friedman et al. 2006, Friedman & Hendry, 2019). VSD was first conceptualized twenty-four years ago and has advanced ever since. In their recent book, Friedman and Hendry (2019) present unique methods that have been applied as part of the VSD's iterative tripartite methodology, i.e. for 1) conceptual, 2) empirical, and 3) technical investigations. The 17 methods cover value elicitation, value analysis, value source or stakeholder identification as well as various other purposes. *Value-based Engineering* (Spiekermann, 2016; Spiekermann & Winkler, 2020) has developed from the same motivations as VSD, but proposes a different methodology. In this approach, three different ethical theories (utilitarianism, virtue ethics and deontology) are applied in the value elicitation phase to identify values that are ethically salient for a specific product and context. It then conceptualizes these values and concretizes them in a technical analysis that is either iterative or risk-assessment based. Both value-oriented approaches share core ideas, such as the integration of direct and indirect stakeholders into the design and development process, the appreciation of values to go beyond functionality, the envisioning of long-term effects, and the consideration of context. At the same time, they also differ in some respects, for example, in how they understand values.

The concept of values is generally difficult to define. In psychology, values are considered to represent desirable behaviours, end states or transitional goals, and the source of a person's self-esteem (Pereira & Baranauskas, 2015; Schwartz, 1994; Verplanken & Holland, 2002). In this view, the relative importance of values depends on the person's culture, socioeconomic status, and practical context (Verplanken & Holland, 2002). The VSD community commonly refers to values as "what a person or group of people consider important in life" (Friedman, Kahn Jr., et al., 2006) with a focus on morality and ethics (Friedman & Hendry, 2019). Value-based Engineering, on the other hand, builds on the philosophical understanding of values developed by Material Ethics of Value (Hartmann, 1932; Scheler, 1913-1916/1973), which understands values as *ought-to-be* principles that should generally guide behaviour. Several scholars in different theoretical contexts have come up with lists of values to support an exemplary understanding of the concept of values (e.g. Friedman, Kahn, Borning, & Huldtgren, 2013; Winkler & Spiekermann, 2019). While value lists can provide a helpful resource for incorporating ethical considerations into technology design, but the cultural and subjective variety of values challenge the completeness of any such list. Also, a list does not provide a solution for the selection of the most relevant values that need to be taken into account for a certain technology and its specific context. Therefore, the acknowledgement of the context-specificity of values has formed a key characteristic of value-oriented projects, which usually start with the identification of values for a specific technology and its context.

While the different methodological and theoretical value frameworks might lead to the identification of similar values for a specific technology, they still influence how relevant values are selected. Thus, the underlying understanding of the concept of values as well as criteria for the elicitation and selection of values need to be made transparent. A transparent process also helps to avoid known challenges in qualitative research, including confirmation biases, culture biases and other cognitive biases (Kahneman, Slovic, & Tversky, 1982; Plous, 1993), which can endanger the validity, reliability and value consciousness of value-oriented projects. Especially since the roles of the designer, value prioritizer, interpreter, reporter and conflict solver are often subsumed in one person, keeping track of the value analysis process is important to avoid a power discrepancy among the actual affected stakeholders and those working on the data (Borning & Muller, 2012). This power discrepancy in combination with a lack of transparency could lead to resistance from managers, engineers, and designers involved in the development of the technological product and thus to a failure of the whole value-oriented project. This, in turn, requires an extremely careful treatment of value data and detailed documentation to ensure transparency.

Human values need to be discussed in detail to explore different contexts, interpretations and nuances (Steen & van de Poel, 2012). Building on rich and versatile value knowledge also helps in communicating values (Pommeranz, Detweiler, Wiggers, & Jonker, 2012) and most importantly in understanding their true meaning. For example, considering “freedom from harms” as a general definition of “security” does not provide enough information on what specifications or requirements a product needs to fulfil. Security could refer to the protection of private data using encryption as well as to the protection of a private estate using video surveillance. This necessary level of detail for (context) information poses another challenge for the representation and communication of values during the development of a technical product.

Rich coding manuals have been developed to preserve context-specific information throughout the analyses (e.g. Friedman, Kahn, Hagman, Severson, & Gill, 2006; Hendry, Abokhodair, Kinsley, & Woelfer, 2017). Such a diligent data analysis procedure is essential, but a value-oriented design project that focuses on the integration of stakeholder perspectives ideally supports transparency throughout the whole design process. This is already time-intensive work for small scale projects and becomes more and more difficult to sustain in larger value-oriented projects. In large-scale projects, the amount of generated value content coming from numerous stakeholders and potentially numerous methods can be enormous. Additionally, the content needs to be updated and extended constantly due to the iterative nature of value-oriented design processes, through which new insights are continuously analysed and validated with stakeholders. These dynamics make it difficult to apply value-oriented design in large industry projects, where high-quality technology development can only be achieved through a transparent and traceable product design and development process. In summary, value-oriented approaches face several challenges. First, there are various ways to define values and related concepts, leading to potentially different selection criteria across individuals and teams. Second, the aim to represent different stakeholder perspectives puts a lot of responsibility onto those involved in the selection and analysis of original value data, which can lead to biases and problematic power discrepancies. Third, values are always bound to specific contexts, and this context-specificity needs to be preserved in every step of analysis. Fourth, the consideration of multiple stakeholder perspectives and the iterative nature of value-oriented approaches requires constant extension of value knowledge, leading to enormous and volatile datasets, which form the basis for representing and communicating values and associated value knowledge. In a recent case study (Spiekermann-Hoff, Winkler, & Bednar, 2019), we have encountered the challenges enlisted above. This inspired us to look for a solution, which we believe can be found in building ontologies and knowledge bases. After introducing the case study, we present a methodology that allows to define value concepts, fill them with rich datasets, and track changes throughout the design process.

3. A NEW WAY OF REPRESENTING AND COMMUNICATING VALUES

For a representation of gathered value knowledge that is independent from the underlying theoretical understanding of values and specific methods of e.g. value elicitation, we borrow techniques from the semantic web community. The semantic web is an extension of the current web that aims at converting unstructured and semi-structured information into information of which the underlying semantics are expressed in a formal machine-understandable way (W3C, 2014).

Within the vision of the semantic web, ontologies play a key role in providing formally defined terms for describing resources in an unambiguous manner. The term *ontology* describes a form of semantic knowledge representation (Ehrlinger & Wöß, 2016). An ontology is an explicit description of concepts (or *classes*) and their properties within an area of interest. In our case, the area of interest comprises human values for a specific technology context and their relations among each other as well as with the stakeholders. Ontologies can be used flexibly to define concepts and relations, but also allow the definition of constraints (e.g. a value being relevant only for the affected stakeholders is a constraint). Once an ontology is “filled” with individual instances (i.e. specific values for a specific technology context), a *knowledge base* is formed (Noy & McGuinness, 2001). The information contained in the knowledge base can be visually represented with graphs, which connect single concepts (or *nodes*).

Expressing the semantics of value knowledge is necessary in order to preserve the connections and avoid the shortcomings of descriptive data analysis. The *Resource Description Framework* (RDF) is a domain-independent data model that expresses information with a specific vocabulary, for instance as a RDF Schema (RDFS; Antoniou et al. 2012). The smallest entity in the RDF is an *RDF statement*, also referred to as *semantic triple* (W3C, 2014) as it consists of three elements: the subject, the object and the predicate. An RDF statement expresses a relation between the subject and the object and the predicate represents the nature of their relationship. The *property* denotes the relationship between the subject and the object (W3C, 2014). Representing knowledge with RDF statements supports the reuse and expansion of knowledge (Antoniou et al., 2012).

3.1. Case study: A telemedicine communication system

Our case study is based on an idea for a telemedicine communication system, which was analysed in accordance with Value-based Engineering (for a detailed description of the case study, see Spiekermann-Hoff, Winkler, & Bednar, 2019). This online communication system connects patients to a general practitioner (GP) who first records patients’ medical history and symptoms and then refers them to a specialized doctor who was recommended by other doctors. Several values can immediately be identified with the envisioned beneficial effects of this IT product, e.g. “health”. However, the underlying recommendation system and the telecommunication system also raise ethical issues, especially as they are used in a medical context. Consider, for example, the underlying motivation of recommendations among doctors (who might know each other) and the fact that a digital platform does not allow physical interaction, which could influence the GP’s decision and undermine values such as “fairness” and “accuracy”. Because of these important ethical implications, the envisioned system fits perfectly as a case study for a value-oriented project.

The value elicitation was conducted by 35 students (age: *M* = 24.56, *SD* = 2.61, 38.2% female, 14 different nationalities) that formed teams of two. 13 participants were female (38.2%) and 21 male (61.8%; 1 missing value). All participants were students at the Vienna University of Economics and Business. Value-related data was gathered following the Value-based Engineering approach (Spiekermann, 2016; Spiekermann & Winkler, 2020), which deploys three ethical theories to elicit values: consequentialism, virtue ethics, and deontology. Once participants had identified relevant values, they were asked to come up with design ideas to further foster beneficial value effects and to avoid negative effects.

3.2. Building an ontology for value representation

For building an ontology for Value-based Engineering, we followed steps and guidance provided by Noy and McGuinness (2001). Figure 1 shows a visualisation of the resulting ontology with the most important terms and their relations.

Figure 1. Exemplified ontology for Value-based Engineering.

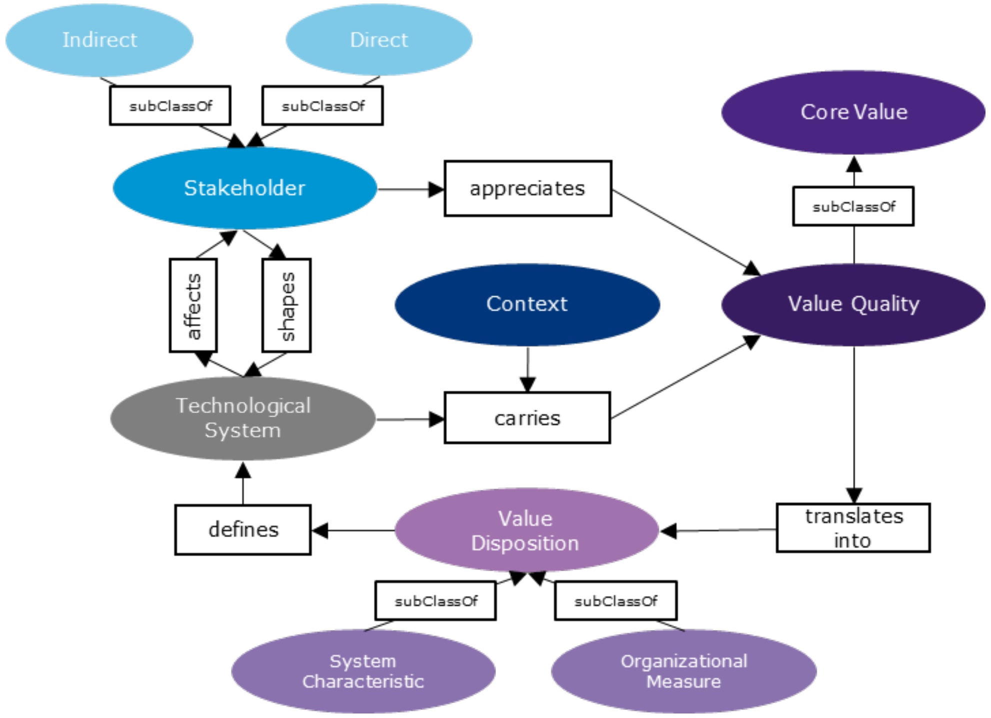

As a first step, we determined that the main domain of this ontology is the representation of value knowledge for Value-based Engineering. This scope definition has several consequences for deciding on appropriate terms (step 2) and defining classes and class hierarchies (step 3).

In step 2, we accumulated appropriate terminology from the Value-based Engineering literature (Spiekermann-Hoff et al., 2019; Spiekermann, 2016; Spiekermann & Winkler, 2020), literature on material value-ethics (Hartmann, 1932; Scheler, 1913-1916/1973) and value lists (Winkler & Spiekermann, 2019). This resulted in numerous important terms, including *core value, value quality, indirect stakeholder, direct stakeholder, affects, appreciates,* or *system characteristics*.

For the third step, we mainly used a top-down approach, by first defining the most general concepts (or *classes*) and then further defining them with sub-concepts (or *sub-classes*). These classes need to be able to describe subjects and objects in an unambiguous manner (Ye et al., 2015), which is a challenge as many terms from step 2 are inherently ambiguous. For instance, according to material value-ethics, a single value can be at the same time a *core value* and a *value quality*. The value "security", for example, can be a value on its own, but also a *value quality* of the value "privacy". We solved this by allowing the attribution of both classes for one value, i.e. "security" can be defined as a value quality and as a value. In our ontology we used *core value*, *stakeholder*, and *value disposition* as main classes and *value quality*, *indirect stakeholder, direct stakeholder, system characteristic,* and *organizational measure* as subclasses.

The fourth step in ontology development is the definition of class relations and class properties. Class relations are described in the Value-based Engineering literature with terms such as *affects, shapes, defines, carries, appreciates* and *translates into*. Class properties in a design case could be, for instance, the *degree of availability*, *social background*, *mean age*, *gender*, or *degree of importance* for a certain stakeholder. Such information is not included in the illustrated example, but can easily be added. The same goes for the type definition of a property, which forms the fifth step. For example, *mean age* would be defined as a *number* here.

The last step, creating instances by a) choosing a class, b) creating an individual instance of that class and b) filling in the type definition, was achieved during coding. Building the underlying ontology helps to make assumptions about the data explicit and can guide the qualitative analysis of the original value data. Ontology development is necessarily an iterative process, which means that a more detailed differentiation of classes, their relation and properties can be achieved when the ontology is filled with specific instances.

3.3. Formulating instances and RDF statements

We prepared all value-related data for the formulation of machine-readable RDF statements or triples, which fill the structure of the ontology, i.e. the defined classes, relations, and properties, with specific instances. Figure 2 visualises the value "health" and related instances for the predefined ontology.

For formulating RDF statements, the data resulting from the value elicitation phase had to be divided into the defined classes (e.g. *core values* or *stakeholders*). This process equals the coding of qualitative data in any research project, as we summarized all value-related ideas, checked their logical structure and named them adequately. To produce RDF statements, we put the coded data into a *subject* – *predicate* (relation) – *object* structure. First, we developed coding rules and formulated examples, as it is standard when developing a coding manual. As a second step, a subset of the original dataset was coded and transformed into RDF statements by two independent coders. We found an inter-coder agreement of 80%, including triples that were completely identical and triples that represented the same entities with a slightly differing wording (example: coder 1: "Leaked doctor information", coder 2: "Information of doctor is leaked"). Afterwards, the coding manual was improved and the whole dataset was coded by coder 1. Coder 2 analysed the resulting semantic triples and checked their logical structure.

Figure 2. Exemplified instances related to “health” for the telemedicine case study.

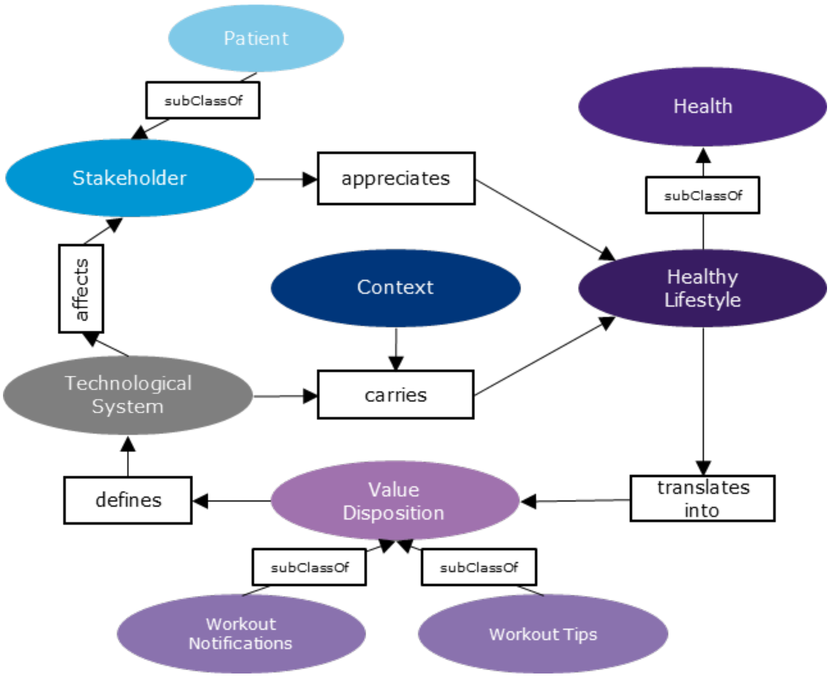

The resulting RDF statements can be combined into a value knowledge base, which represents the original value-related data following the structure of the predefined ontology. This knowledge base maintains the relations suggested in the original data, but can be formulated in a machine-readable way, providing the basis for further design steps. Triples can also be visualized as connected *graphs*, consisting of nodes (representing subjects and objects) and arcs (representing the predicate; W3C 2014). Additionally, node (or class) *properties* can be displayed visually. The open-source software developed for network exploration and manipulation “Gephi” (Bastian, Heymann, & Jacomy, 2009) allows an initial visualization and exploration of these graphs and includes and an adequate spatial visualization through its included “ForceAtlas2” layout algorithm (Jacomy, Venturini, Heymann, & Bastian, 2014). The visual representation is especially powerful, as it makes the most frequent concepts and relations immediately apparent through the relative size of the nodes and the thickness of the arches. An interactive platform that facilitates such visualisations would be an especially powerful contribution to an effective communication among teams and stakeholders.

4. DISCUSSION

The benefits of ontology-building and the development of knowledge bases come from the representation of data in machine-readable form, which supports a structured representation of concepts and their relation. Even for large-scale value-oriented projects, this methodology can offer a common framework for creating high-quality value knowledge. It also facilitates the communication of such knowledge among stakeholders, teams, and across different approaches. The methodology inspired by techniques from the semantic web community offers ways to deal with several challenges in value-oriented design projects.

First, different ways of defining key concepts such as values and stakeholders can lead to different selection criteria across individuals, teams, and methods. Building ontologies helps to share information, analyse and reuse knowledge, and to make assumptions explicit (Noy & McGuinness, 2001). A common knowledge base supports the sharing of gained value knowledge in specific projects (as in our telemedicine case study) while making it possible for every approach to maintain and express its own theoretical foundation (e.g. defining value qualities related to core values). Thus, it allows the representation and communication of values independent from the underlying theoretical background and specific methods used. This supports that value-oriented projects build upon and learn from previous projects more easily. Ontologies can also help to make relations between key concepts explicit (e.g. value qualities are carried by the technological system) and thus support the coding of original value-related data (e.g. by defining value qualities and core values). From an engineering perspective, formally defined terms and classes can make the fuzzy concept of values more tangible, which might encourage a wider adoption of value-oriented approaches.

Second, the aim to represent different stakeholder perspectives puts a lot of responsibility onto those involved in the selection and analysis of original value data. The presented methodology supports transparency of this process by providing detailed information (e.g. by indicating the stakeholders for whom a value quality is important). Representing value knowledge in a formal way also makes it easier to track changes, e.g. when introducing a category in the coding process, and does so in a machine-readable form, which can be queried, visually represented, and explored by stakeholders. Enabling transparency and the exploration of existing value knowledge by any stakeholder can help to decrease power discrepancies (Borning & Muller, 2012) and potential biases (Kahneman et al., 1982; Plous, 1993). Still, researcher values should always be made explicit and considered during the design process (Steen & van de Poel, 2012).

Third, human values are bound to specific contexts, and need to be discussed in detail to explore different contexts, interpretations and nuances (Steen & van de Poel, 2012). This context-specificity needs to be maintained in every step of analysis. The methodology we propose produces coded data that is semantically coherent, that is, represents the underlying logical structure in the form of semantic triples (e.g. when expressing that *patients – appreciate – high quality medical service*). An initially specified set of rules for the formulation of semantic triples supports the completeness of value data and secures important information such as the context of meaning.

Fourth, the consideration of multiple stakeholder perspectives, the iterative nature of value-oriented approaches, and the context-specificity of values leads to enormous datasets that need to be updated constantly. This is especially challenging for large-scale projects. Drawing definitions from a sound theoretical background and following a pre-defined set of rules already form the basics of good research practices in social sciences. But with increasing complexity of a dataset, it becomes more and more difficult to keep track of changes. The machine-readable form of RDF statements supports the digital handling of data and could thus provide a solution here. Building ontologies and knowledge bases could be especially beneficial to the handling of date in large-scale value-oriented projects and increase the recognition of value-oriented approaches outside of academia.

This paper wants to offer an inspirational starting point for making value knowledge explicit, transparent, and accessible to all stakeholders. As this paper presents only first experiences in utilizing semantic web techniques for value-oriented data analysis, the results we present for the case study are only rudimentary. We also acknowledge that methods that produce less structured data might be more difficult to translate into semantic triples. Still, we hope to inspire future value-oriented projects to build more elaborate ontologies and thus to improve the sharing of knowledge among stakeholders, teams, and different approaches. Future work in this area could also explore new ways of value knowledge discovery, e.g. by including query functions, filter mechanisms, interactive visualisation or developing value knowledge patterns, counting triples, nodes, and relations (Presutti et al., 2011).

5. CONCLUSION

In this paper, we show that ontology-building and the development of a knowledge base, methods from the semantic web community, facilitate the representation and communication of values independent from the underlying theoretical framework and the specific method used to acquire value data. At the same time, the proposed methodology supports essential quality criteria for value-oriented projects. The formal representation of value knowledge in form of semantic triples ensures a high level of detail while respecting the context-specificity of any information. The underlying ontology helps to represent key concepts and their relations and supports the transparency of data analysis, including the initial coding of the value-related data. Furthermore, the resulting knowledge base can be shared with stakeholders and researchers, supporting the joint evolution of value-oriented approaches and technology. The case study of a telemedicine communication system shows the steps that a value-oriented project would need to run through to develop an ontology, which can be extended it into a knowledge base to be shared among stakeholders, teams, and across different approaches. Future value-oriented projects could apply this methodology to jointly build an extensive value knowledge base and experiment with more advanced applications such as data query and interactive visualisations.

ACKNOWLEDGEMENTS

We want to thank our PhD supervisor Sarah Spiekermann for supporting us in pursuing this project. Also, Oliver Burmeister, Batya Friedman, David Hendry and the international VSD community have provided valuable feedback and comments, which we highly appreciate.

REFERENCES

Antoniou, G., Groth, P., Harmelen, F. van, & Hoekstra, R. (2012). *A semantic web primer* (3rd ed.). Cambridge, MA: MIT Press.

Barley, S. R., Meyerson, D. E., & Grodal, S. (2011). E-mail as a source and symbol of stress. *Organization Science*, *22*(4), 887–906. https://doi.org/10.1287/orsc.1100.0573

Bastian, M., Heymann, S., & Jacomy, M. (2009). Gephi: An open source software for exploring and manipulating networks. In *Third international AAAI conference on weblogs and social media*.

Bednar, K., Spiekermann, S., & Langheinrich, M. (2019). Engineering Privacy by Design: Are engineers ready to live up to the challenge? *The Information Society: An International Journal*, *35*(3), 122–142. https://doi.org/10.1080/01972243.2019.1583296

Borning, A., & Muller, M. (2012). Next steps for value sensitive design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'12)* (pp. 1125–1134). ACM. https://doi.org/10.1145/2207676.2208560

Cadwalladr, C. (2017, May). The great British Brexit robbery: how our democracy was hijacked. *The Guardian*. Retrieved from https://www.theguardian.com/technology/2017/may/07/the-great-british-brexit-robbery-hijacked-democracy

Chung, L., & do Prado Leite, J. C. S. (2009). On non-functional requirements in software engineering. In *Conceptual modeling: Foundations and applications* (pp. 363–379). Berlin, Heidelberg: Springer.

Detweiler, C., & Harbers, M. (2014). Value stories: Putting human values into requirements engineering. *CEUR Workshop Proceedings*, *1138*, 2–11.

Ehrlinger, L., & Wöß, W. (2016). Towards a definition of knowledge graphs. In *SEMANTiCS 2016*. https://doi.org/10.1007/978-1-349-19066-9_2

Epstein, R., & Robertson, R. E. (2015). The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *Proceedings of the National Academy of Sciences (PNAS)*, *112*(33), E4512–E4521. https://doi.org/10.1073/pnas.1419828112

Friedman, B., & Hendry, D. G. (2019). *Value Sensitive Design: Shaping technology with moral imagination*. Cambridge, MA: MIT Press.

Friedman, B., Kahn Jr., P. H., & Borning, A. (2006). Value sensitive design and information systems. In P. Zhang & D. Galletta (Eds.), *Human-computer interaction and management information systems: Foundations* (pp. 348–372). Armonk, NY: M.E.Sharpe.

Friedman, B., Kahn, P. H., Borning, A., & Huldtgren, A. (2013). Value sensitive design and information systems. In *Early engagement and new technologies: Opening up the laboratory* (pp. 55–95). Springer.

Friedman, B., Kahn, P. H., Hagman, J., Severson, R. L., & Gill, B. (2006). The watcher and the watched: Social judgments about privacy in a public place. *Human-Computer Interaction*, *21*(2), 235–272. https://doi.org/10.1207/s15327051hci2102_3

Friedman, B., & Nissenbaum, H. (1996). Bias in computer systems. *ACM Transactions on Information Systems*, *14*(3), 330–347. https://doi.org/10.1145/230538.230561

Hartmann, N. (1932). *Ethics*. London: George Allen & Unwin.

Hendry, D. G., Abokhodair, N., Kinsley, R. P., & Woelfer, J. P. (2017). Homeless young people, jobs, and a future vision: Community members' perceptions of the Job Co-op. *ACM International Conference Proceeding Series*, *Part F1285*, 22–31. https://doi.org/10.1145/3083671.3083680

Isomursu, M., Ervasti, M., Kinnula, M., & Isomursu, P. (2011). Understanding human values in adopting new technology: A case study and methodological discussion. *International Journal of Human-Computer Studies*, *69*, 183–200. https://doi.org/10.1016/j.ijhcs.2010.12.001

Jacomy, M., Venturini, T., Heymann, S., & Bastian, M. (2014). ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software. *PLoS ONE*, *9*(6), e98679.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge University Press.

Miller, J., Friedman, B., Jancke, G., & Gill, B. (2007). Value tensions in design: The value sensitive design, development, and appropriation of a corporation's groupware system. *Proceedings* *of the 2007 International ACM Conference on Supporting Group Work (GROUP '07)*, 281–290. https://doi.org/10.1145/1316624.1316668

Noy, N. F., & McGuinness, D. L. (2001). Ontology development 101: A guide to creating your first ontology. Retrieved from http://www.ksl.stanford.edu/people/dlm/papers/ontology101/ontology101-noy-mcguinness.html

O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. New York: Crown.

Pariser, E. (2011). *Filter bubble: What the Internet is hiding from you*. New York: The Penguin Press.

Pereira, R., & Baranauskas, M. C. C. (2015). A value-oriented and culturally informed approach to the design of interactive systems. *International Journal of Human Computer Studies*, *80*, 66–82. https://doi.org/10.1016/j.ijhcs.2015.04.001

Plous, S. (1993). *The psychology of judgment and decision making*. New York: McGraw-Hill.

Pommeranz, A., Detweiler, C., Wiggers, P., & Jonker, C. (2012). Elicitation of situated values: Need for tools to help stakeholders and designers to reflect and communicate. *Ethics and Information Technology*, *14*, 285–303. https://doi.org/10.1007/s10676-011-9282-6

Presutti, V., Aroyo, L., Adamou, A., Schopman, B., Gangemi, A., & Schreiber, G. (2011). Extracting core knowledge from Linked Data. In *Proceedings of the Second International Conference on Consuming Linked Data* (Vol. 782, pp. 37–48). CEUR-WS.

Scheler, M. (1973). *Formalism in ethics and non-formal ethics of values: A new attempt toward the foundation of an ethical personalism [1913-1916]*. (M. S. Frings & R. L. Funk, Eds.). Evanston, Ill: Northwestern University Press. https://doi.org/10.2307/2707101

Schwartz, S. H. (1994). Are there universal aspects in the structure and contents of human values? *Journal of Social Issues*, *50*(4), 19–45. https://doi.org/10.1111/j.1540-4560.1994.tb01196.x

Spiekermann-Hoff, S., Winkler, T., & Bednar, K. (2019). *A telemedicine case study for the early phases of value based engineering* (Working Paper Series/Institute for IS & Society No. 001_vs 1). *Working Paper Series/Institute for IS & Society*. Vienna: WU Vienna University of Economics and Business. Retrieved from https://epub.wu.ac.at/7119/

Spiekermann, S. (2016). *Ethical IT innovation: A value-based system design approach*. Boca Raton: CRC Press.

Spiekermann, S., Korunovska, J., & Langheinrich, M. (2018). Inside the organization: Why privacy and security engineering is a challenge for engineers. *Proceedings of the IEEE*, *107*(3), 600–615. https://doi.org/10.1109/JPROC.2018.2866769

Spiekermann, S., & Winkler, T. (2020). Value-based Engineering for Ethics by Design. *ArXiv*, *2004.13676*. Retrieved from https://arxiv.org/abs/2004.13676

Steen, M., & van de Poel, I. (2012). Making values explicit during the design process. *IEEE Technology and Society Magazine*, *31*(4), 63–72. https://doi.org/10.1109/MTS.2012.2225671

van den Hoven, J. (2017). Ethics for the digital age: Where are the moral specs? In *Informatics in the future: Proceedings of the 11th European Computer Science Summit (ECSS 2015), Vienna,* *October 2015* (pp. 65–67). Cham: Springer Nature. https://doi.org/10.1007/978-3-319-55735-9

Verbeek, P.-P. (2008). Morality in design: Design ethics and the morality of technological artifacts. In S. A. M. Pieter E. Vermaas, Peter Kroes, Andrew Light (Ed.), *Philosophy and design: From engineering to architecture* (pp. 91–103). Springer Science + Business Media. https://doi.org/10.1007/978-1-4020-6591-0_7

Verplanken, B., & Holland, R. W. (2002). Motivated decision making: Effects of activation and self-centrality of values on choices and behavior. *Journal of Personality and Social Psychology*, *82*(3), 434–447. https://doi.org/10.1037/0022-3514.82.3.434

W3C. (2014). RDF 1.1 Primer. Retrieved from https://www.w3.org/TR/rdf11-primer/#section-triple

Winkler, T., & Spiekermann, S. (2019). Human values as the basis for sustainable information system design. *IEEE Technology and Society Magazine*, *38*(3), 34–43. https://doi.org/10.1109/MTS.2019.2930268

Ye, J., Dasiopoulou, S., Stevenson, G., Meditskos, G., Kontopoulos, E., Kompatsiaris, I., & Dobson, S. (2015). Semantic web technologies in pervasive computing: A survey and research roadmap. *Pervasive and Mobile Computing*, *23*, 1–25. https://doi.org/10.1016/j.pmcj.2014.12.009