

Supporting complex decision making by semantic technologies

Stefan Fenz

SBA Research and TU Wien, Floragasse 7, 1040 Wien, Austria
`stefan.fenz@tuwien.ac.at`

Abstract. Complex decisions require stakeholders to identify potential decision options and collaboratively select the optimal option. Identifying potential decision options and communicating them to stakeholders is a challenging task as it requires the translation of the decision option's technical dimension to a stakeholder-compliant language which describes the impact of the decision (e.g., financial, political). Existing knowledge-driven decision support methods generate decision options by automatically processing available data and knowledge. Ontology-based methods emerged as a sub-field in the medical domain and provide concrete instructions for given medical problems. However, the research field lacks an evaluated practical approach to support the full cycle from data and knowledge assessment to the actual decision making. This work advances the field by: (i) a problem-driven ontology engineering method which (a) supports creating the necessary ontology model for the given problem domain and (b) harmonizes relevant data and knowledge sources for automatically identifying decision options by reasoners, and (ii) an approach which translates technical decision options into a language that is understood by relevant stakeholders. Expert evaluations and real-world deployments in three different domains demonstrate the added value of this method.

Keywords: semantic technologies · ontology engineering · decision support

1 Introduction and problem statement

Making complex decisions is fundamental to business activities. In the context of this work, decisions are defined as complex if (i) they are unique, i.e., non-repeatable, (ii) they involve uncertainty, (iii) they require multiple decision makers to make the decision, (iv) the necessary domain knowledge is not present amongst all required decision makers, and (v) decision makers may not understand the full impact of their decision (cf. [7]). For example, the top management of a company has to decide about different IT security strategies, risk mitigation measures, or building refurbishment projects. In these scenarios, management may knowledgeable about maximum investment costs, business/political impact, and acceptable risks, but they may not know how an optimization in

one dimension affects the remaining decision dimensions. For instance, reducing the investment costs may result in lower energy efficiency in building refurbishment projects and spending less money on IT security may result in a higher IT security risk level. The challenge in complex decision making is to quantify these observations in such a way that a decision maker – who is not fully knowledgeable about the underlying technical issue – can still identify the options which maximize the cost/benefit ratio of the decision.

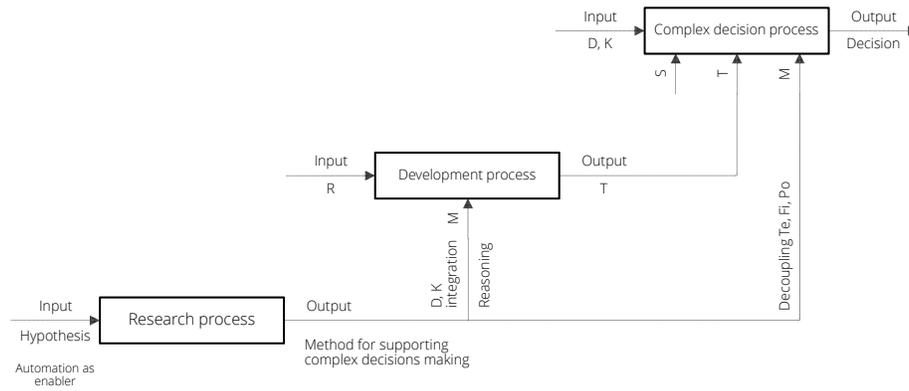


Fig. 1. Interaction of research, development and decision process

In general, a complex decision process (cf. Figure 1) requires data (D) and knowledge (K) as well as stakeholders (i.e., decision makers) (S) using tools (T) and methods (M) to process the input (D, K), with the final goal of reaching a decision. The necessary tools are the output of a development process which uses methods and requirements (R) to generate the decision support tool, i.e., the output. The research process of this work is driven by the hypothesis that only an increased degree of automation of the necessary data/knowledge integration and the subsequent automated reasoning can efficiently support the complex decision making process. Automation requires that (i) the knowledge necessary to identify decision options has to be available in a machine-readable way (e.g., in the form of ontologies), and (ii) automated reasoning engines have to automatically derive decision options based on this machine-readable knowledge body. The main challenges in achieving this goal is on the one hand to efficiently create the machine-readable knowledge body, and on the other hand to communicate the decision options identified by reasoning engines in a stakeholder-comprehensible way.

As a result, this work contributes the following methods to support complex decision making via semantic technologies: (i) data/knowledge integration and a reasoning method for efficiently querying the knowledge body to identify decision options, and (ii) a method which separates the relevant dimensions (e.g., tech-

nical, financial, and political) in the decision process to provide the stakeholders with comprehensible decision options.

2 Theoretical background

The theoretical background of this work stems from the normative decision theory and the stakeholder theory. The comparison of options in terms of their cost/benefit categories is the core of the normative decision theory. A minimal amount of rationality is expected, and the decision maker usually prefers the option with the most valuable outcome (cf. [15]). While the normative decision theory is concerned with a single decision maker, the stakeholder theory [10] states that not only one party (e.g., shareholders), but numerous parties such as employees, political parties, financiers, trade associations, customers, etc. should be considered in organizational decisions. E.g., in a company context, the challenge for the management is to maximize the value for the stakeholders without negatively impacting business operations. A broad consensus, making sustainable growth and success more likely, can be supported by:

- Informing decision makers about the full range of decision options in a language that is understood by the decision makers
- Ensure that the decision options are technically and financially feasible
- Provide comprehensible information relevant for the decision option (cost, benefits)
- Support collaborative decision making

To efficiently meet these requirements in complex decision making scenarios (involving a large number of decision alternatives), a decision support system with the following characteristics is required:

- Knowledge relevant to the decision making is available in a machine-readable way to ensure processing in a timely manner
- Potential decision options on the basis of the knowledge body are identified by automated means
- Decision options are presented in a comprehensible way to the decision makers

In the following section we review the state of the art and derive the field's research gap.

3 State of the art and research gap

According to [13] existing decision support systems (DSS) can be classified as follows: *Model-driven DSS* access and manipulate finance, optimization or simulation models (examples: production planning management decision system [6] and production scheduling application [5]). *Data-driven DSS* access and manipulate time series of internal and external data (example: analytical airline

information management system [8]). *Communication-driven DSS* [3] support decision-relevant communication by using information and network technology as the main architectural component (examples: groupware, video conferencing and Wikis). *Document-driven DSS* support decision making through document retrieval and analysis (examples: full-text document search engines and document organization systems). *Knowledge-driven DSS* store the knowledge about a specific domain in a machine-readable form and are capable of automatically using this knowledge to support the decision maker in identifying solutions for a given problem. The roots of knowledge-driven DSS date back to 1965 (cf. [4]) and it has been steadily improved since.

Ontology-based DSS are a sub-field of knowledge-driven systems and use an ontology together with a reasoning engine to support the decision making. Ontologies are a formal way to define the structure of knowledge for a certain domain. Classes are used to represent and define concepts of the domain, and properties are used to establish relations between the classes. OWL¹ is currently the dominant language to encode ontologies.

Personal recommendation system (e.g., [9]) and clinical decision support systems (e.g., [1], [11]) are examples for ontology-based decision support systems. These systems use ontologies to encode the underlying knowledge body (e.g., symptoms and diseases) and reasoners to infer new knowledge (e.g., potential diseases based on observed symptoms).

The role of the ontology in these systems is critical because its capabilities (formally describing domain concepts, their relation/dependencies to each other, description logic statements which can be used to evaluate a certain state with a reasoner, etc.) are required to efficiently support the complex decision-making process by automatically finding technically feasible solutions for the decision problem. Ontologies enable us to define the domain once and reuse it in similar context with minimal additional costs. As OWL is used to encode the ontology, a wide range of editors and reasoners can be used and a vendor lock-in is prevented. Not using ontologies would require us to develop all these functionalities/tools from scratch and the reusability of the artifacts would suffer. The main strength of the reasoner is that it supports description logics. Intersection/union/negation of concepts, universal/existential restrictions, etc. are provided out of the box by OWL and compatible reasoners. These concepts can be combined to powerful description logic statements which go beyond basic matching of criteria. Examples are building renovation measures which have alternative requirement sets and each requirement set contains mandatory and optional components.

E.g., a heating system of a building can be renovated by replacing it by a heat pump or a gas condensing boiler. A heat pump ready building has at least, one floor heating system or wall heating system, and at least, one surface collector or deep drilling, and, a power grid connection which is greater than x kilowatts. A gas condensing boiler ready building needs at least, one chimney and at least, one access to the gas network with a capacity that is greater than the heating and warm water demand of the building. Such requirements can be

¹ W3C web ontology language. see <https://www.w3.org/OWL/> for further details

formally described by description logics and reasoners can be used to evaluate the statements in a specific context (such as a concrete building modelled inside the ontology).

Based on the requirements derived from the theoretical background of our work, we identified the following research gap to increase the practical value of ontology-based decision support systems:

- A problem-driven ontology engineering method which provides a design pattern for how to structure the ontology in order to maximize its value and expressiveness in a given decision support scenario.
- Stakeholder communication to present the output of the ontology-based DSS in a comprehensible form to enable the stakeholders to collaboratively identify the most suitable decision.

In the following sections we outline our research contribution to these challenges and show how we have evaluated our research results in practice.

4 Integration and reasoning

The main challenge of ontology-based decision support systems is to model the ontology in a way that it supports problem-solving reasoning. Existing ontology engineering methods support (i) the creation of general ontologies (e.g., [17], [12], [14]), (ii) specific user groups such as non-experts (e.g., [2]), (iii) specific scenarios such as ontology re-use (e.g., [16]), and (iv) specific organizational settings such as collaboration (e.g., [18]). While all of these approaches work in their intended fields, they do not give explicit guidance on how to build an ontology which is capable to assist in complex decision support scenarios.

We developed a problem-driven ontology engineering method which (i) uses the decision problem and available data/knowledge as design foundation, (ii) harmonizes knowledge sources with regard to the decision problem, and (iii) evaluates the maintainability of and contribution to problem solving in order to rate the quality of the produced ontology. The following paragraphs outline each step of the developed method:

4.1 State decision problem

The problem which should be addressed by the decision support system and the underlying ontology have to be clearly stated, since the ontology competency questions in Step 4 are derived from the stated problem. Examples: What energy efficiency measures have the best cost/benefit ratio in terms of costs and CO₂ emissions with regard to a specific building? What IT security measures have to be implemented in a given organization to reduce risks to an acceptable level and to not exceed allocated budgets?

4.2 Review relevant knowledge sources

To address the stated decision problem, it is necessary to review relevant knowledge sources such as guidelines, handbooks, laws, and process knowledge. The review process should strictly focus on understanding and extracting the knowledge which is necessary to solve the problem. This step is completed when the method of how to identify potential solutions to a given problem is fully understood. Examples: (i) identifying appropriate insulation products for a building refurbishment requires an understanding of the building heating demand calculation and all associated calculation parameters, (ii) identifying useful information security products requires an understanding which vulnerabilities are mitigated by these products.

4.3 Model-relevant problem parameters on an abstract layer

The knowledge source review in the previous step is used as input for the modeling process. The goal is to model only the most relevant problem parameters which are necessary to solve the stated problem on an abstract layer. Holistically model the domain or to include knowledge which goes beyond the purpose of solving the problem is a clear non-goal in this step. Figure 2 shows an example from the building energy efficiency domain. The building (refurbishment project) is characterized as a class which is connected to (i) existing building components (windows, roof, etc.) with specified thermal conductivity levels, (ii) access to energy networks, (iii) existing housing technology such as photovoltaics (PV), heat distribution systems, etc., and (iv) refurbishment candidate classes which allow an ontology reasoner to classify the building according to its specific refurbishment needs (see Step 4 for further details). Example: a building equipped with a very inefficient heating system would be classified as a heating system modernization candidate.

In general, the final model has to include:

- Ontology classes, properties, and individuals which can be used to describe the status quo in sufficient detail. If it is for example required to identify heating system replacement options, the model has to provide the necessary components to formally express the current heating system status and heating-relevant building parameters such as heat demand.
- Ontology classes, properties, and individuals to describe potential solutions to the problem. In the context of the heating system example, the model has to allow the integration of new heating systems including their technical and financial characteristics such as cost or efficiency parameters.
- Ontology classes which can be used to map potential solutions to the status quo. See Step 4 for further details.

4.4 Create description logic statements to validate the model

The goal of this step is to create description logic statements which can be used to validate the model with competency questions that are suited to solve the

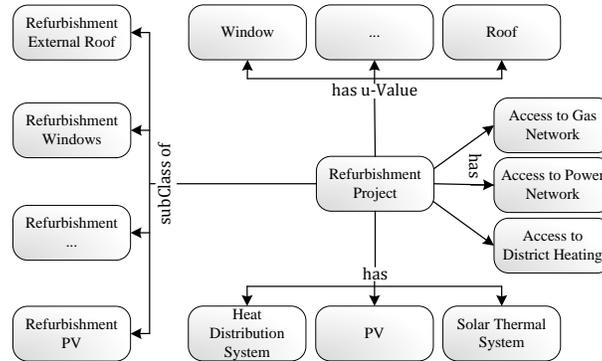


Fig. 2. Step 3 - output example

problem stated in Step 1. Figure 3 shows an example of statements to analyse whether the heating system of a building can be replaced with a modern gas-driven one, therefore checking if the building has (i) a chimney to get the exhaust out of the building, (ii) access to the gas network, and (iii) a heat distribution system which would work with a gas-driven system (floor/wall heating system or radiators). We used the Protege ontology editor to create the description logic statements and a reasoner to validate the model created in Step 3. The validation is successful if the reasoner classifies the building into the correct refurbishment candidate categories.



Fig. 3. Step 4 - output example (Translation: Raumheizung - space heating, Warmwasser - hot water, verwendet Energietraeger - uses energy carrier, Gas - natural gas, Sanierungsprojekt - renovation project, Kamin - chimney, Gasnetzzugang - natural gas network access, hat Abgabesystem - has delivery system, Flaechenheizung - panel heating, Fussbodenheizung - underfloor heating, Heizkoerper - radiators)

In general, the following pattern is applied to create the description logic statements based on the ontology classes, properties and individuals defined in Step 3:

- Define the technical requirements which have to be met in order to map a status quo element to a potential solution. Example: classifying a building as a candidate for a modern gas-driven heating system requires that the building is already equipped with a chimney, has access to the gas network, and has a compatible heat distribution system.
- Express these requirements by description logic statements as anonymous classes which are defined as equivalent to the mapping class. See Figure 3 for an example.
- Test the statements by running the reasoner and checking if the classification was done as intended. Extend and modify the statements until the reasoner produces correct results.

4.5 Harmonize and integrate relevant knowledge sources based on the model

Based on the abstract model created in the previous steps, concrete and relevant knowledge necessary for solving the problem is harmonized and integrated. Examples: integration of concrete building refurbishment products including feature and price data, compatibility information such as heating technology and heating distribution systems, and legal requirements regarding minimal thermal conductivity values of insulation products.

4.6 Validate the model

In Step 6 the ontology contains all necessary components (domain knowledge, domain model, environment data, solution data, and description logic statements) to assist, in combination with a reasoner, with solving the stated decision problem. The validation of the ontology reasoner output is done by modeling concrete environment data (e.g., data of a concrete building we have to refurbish) in the ontology. After running the reasoner, the output (see Figure 4 for an example) is validated by experts with regard to its correctness and usability in further decision support operations.

5 Decoupling

The output of the reasoner defines potential solutions for the stated problem. For instance, as shown in Figure 4, the energy efficiency of the building can be improved by putting insulation on the outer walls, switching to a modern gas heating system or putting photovoltaics on the roof. All measures suggested by the reasoner comply with the building and the requirements which have been modeled in the ontology (e.g., legal or technical requirements).

Every suggested measure can be part of the final solution which has to be identified by the decision makers based on their preferences. The following process ensures that the correct solution data is presented in a comprehensible way:

1. Extract cost and benefit data for each measure from the ontology



Fig. 4. Step 6 - output example

2. Identify feasible solution sets by creating all combinations of the identified measures (e.g., solution set X would be to put photovoltaics on the roof and insulate the outer walls, solution set Y would be solution set X plus replacing the heating system)
3. Visualize each solution set with data which is relevant to the decision maker (costs, break even, etc.)
4. Enable the decision maker to sort and filter the solution sets

Figure 5 shows the user interface of a building refurbishment decision support system which was built based on the developed methods. The user specifies the maximum investment costs, relevant goals such as minimum CO₂ emission reduction and renewable energy share. By clicking on the optimization button the system conducts the aforementioned steps and presents a list of solution sets which can be filtered and sorted. By selecting a specific solution set the user learns about its specific measures. The main benefit of this method is that all presented solution sets are completely compatible to the given technical and financial requirements (only feasible measures are suggested by the reasoner based on the knowledge and data modeled in the ontology).

By decoupling the technical from the financial/political dimension we enable stakeholders to focus on the decision parameters they understand. The developed semantic decision support methods ensure that potential solutions are identified in an automated way and fully compatible with the actual problem environment (e.g., the concrete building stock).

The implemented data visualization enables stakeholders to explore the solution space in an interactive way and to find the most suited solution by collaboratively evaluating the consequences in relevant cost/benefit categories. As such, the developed method supports the stakeholder theory (see Section 2).

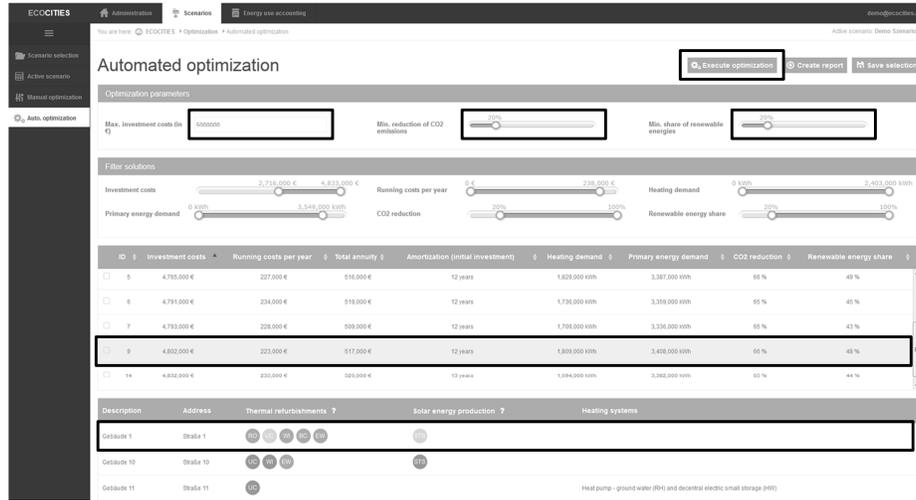


Fig. 5. User interface

6 Evaluation

The developed methods have been evaluated in four European small- and medium-sized companies (three different productive application fields) and two governmental institutions.

6.1 Single building refurbishment

Semergy.net is a decision support system for single building refurbishment which addresses both home owners and professionals. The user models the building and sets maximum investment costs; based on that the system calculates numerous ways of how to improve the building energy efficiency at certain investment resp. running energy costs and payback periods (see Figure 6).

Within an extensive validation phase, experts checked the system output regarding technical and financial feasibility and confirmed that the system provides refurbishment suggestions which are compliant with their respective expectations. 38 of these expert tests were conducted. Currently 590 users are registered on semergy.net and use it for their personal and professional energy efficiency calculations. Semergy.net reduces the time for identifying the most suitable energy efficiency strategy for a given building by up to 81 percent compared to traditional energy performance certificate (EPC) tools. This substantial time reduction is based on the high automation level w.r.t. the identification of appropriate energy efficiency measures. Table 1 shows the evaluation results.

Please note that traditional EPC tools do not provide an automated extensive search for refurbishment options. The user has to manually adjust the building parameters (e.g., outer wall insulation) to see how they affect the output in terms of energy efficiency. Costs of the energy efficiency measures are

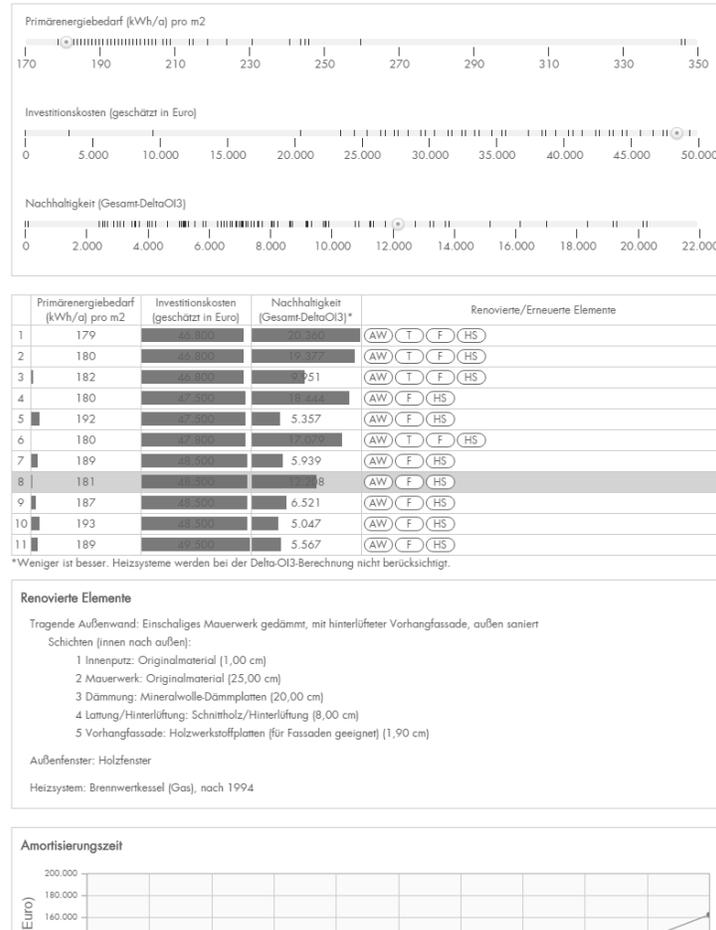


Fig. 6. semergy.net - decision support user interface

	Semergy.net	EPC tools
Entering building data	9min	10min
Identifying solutions	2min	75min (limited to 15 solutions)
Decision making	6min	9min
Total	17min	94min

Table 1. Semergy.net evaluation results - times on average across 38 test runs

Construction type	Candidates
Outer walls	268
Inner walls (load bearing)	77
Inner walls (non-load bearing)	77
Earth-facing walls	30
Earth-facing floor	216
Subceiling	6
Basement ceiling	250
Roof	467
Door	2
Window	5
Roof top window	5
Heating system	6
Total combinations	$2.16 * 10^{18}$

Table 2. Semergy.net - candidates per construction type on the example of a two story single family home with heated basement

also not provided by these tools and the compatibility of measure combinations has to be judged manually by the user. Because of this high manual effort at traditional tools, we required experts to identify only 15 refurbishment strategy options and measured the corresponding execution time. The evaluation has shown that the semantic knowledge base of semergy.net enabled the reasoner to identify all potential measure combinations and to prepare the data (investment costs and energy efficiency) for a comprehensible visualization of the decision options. Compared to traditional methods, the developed method provides not only significant time savings but also a broader range of decision options and a comprehensible visualization of the decision dimensions.

6.2 Multiple building and energy network refurbishment

Ecocities.at (see the UI in Figure 5) is a decision support system for identifying energy efficiency measures in large building groups. Compared to semergy.net, the system operates in the context of multiple buildings (e.g., 30), potential synergies/dependencies among these buildings, and energy networks within a building group. The solution set of ecocities.at (concrete energy efficiency measures on each building of the building group and global impact data such as costs and CO₂ emission reduction) was validated together with experts and pilot customers. The validation phase showed that a lot of feasible energy efficiency strategies were overlooked in non-automated considerations of the problem. The calculated solution space was much bigger and allowed the users to collaboratively identify the most suitable solution. Ecocities.at is the first product of its kind on the market. In comparison to traditional methods – which combine manual and tool work – ecocities.at reduces the time required for identifying appropriate energy efficiency strategies by 74 percent. Table 3 shows the evaluation results which are average times based on 7 test runs including 10 to 32 buildings.

	Ecocities.at	EPC tools and Excel
Entering building data	370min	330min
Identifying solutions	47min	1350min (limited to 150 solutions)
Decision making	15min	15min
Total (10 buildings)	432min	1695min

Table 3. Ecocities.at evaluation results - times on average across 7 test runs

Construction type	Candidates
Outer walls	4
Earth-facing walls	4
Uppermost ceiling / roof	4
Window	4
Heating system	71
PV and solar thermal	10
Total combinations	181.760

Table 4. Ecocities.at - candidates per construction type

Please note that because of the large number of possibilities a full identification of refurbishment options across the entire building group is not feasible by manual means. Therefore, we measured the time necessary to manually calculate the effect of a single refurbishment measure and extrapolated it to all potential measures. The manual calculation task includes the decision which refurbishment measure to implement (e.g., putting 20cm wall insulation on building X and replacing the heating systems at building Y) and checking how this affects the energy efficiency of the entire building group. Conducting this single task by traditional EPC tools and Microsoft Excel requires on average 9 minutes. However, for each building ecocities.at considers 218.700 refurbishment strategies based on the following refurbishment options: (i) three different qualities for earth-facing floor, outer wall and uppermost ceiling insulation, (ii) three different window qualities, (iii) three different photovoltaic systems, (iv) six different solar thermal systems, (v) ten different heating systems, and (vi) 15 different hot water production systems. By manual means this would result in a time effort of around 1 year for one building only. In reality an expert would not calculate all possible combinations, but would choose feasible combinations based on her/his experience (around 15 combinations per building). Compared to traditional methods, the developed method identifies a much larger space of refurbishment options and significantly reduces the time which is necessary to conduct the entire planning process.

6.3 IT security risk management

AURUM is the first ontology-driven IT security risk management product. It supports organizations in identifying the optimal information security measures in terms of costs, effectiveness, and compliance to standards. It is designed to (i) minimize the necessary interaction between user and system, and (ii) provide

decision makers with an intuitive solution that can be used without in-depth information security knowledge. The integration and reasoning components of AURUM make sure that only technically and financially feasible security measures are suggested to the decision maker. The selection of the final security measure strategy is based on its investment, running costs, and organization-wide risk level after implementation (see Figure 7). 18 test runs were conducted with experts and end-users, one productive installation is currently deployed at an governmental institution in Europe. The evaluation results have shown that AURUM provides no time savings but a deeper and broader range of comprehensible security measure strategies compared to traditional risk and compliance management tools.

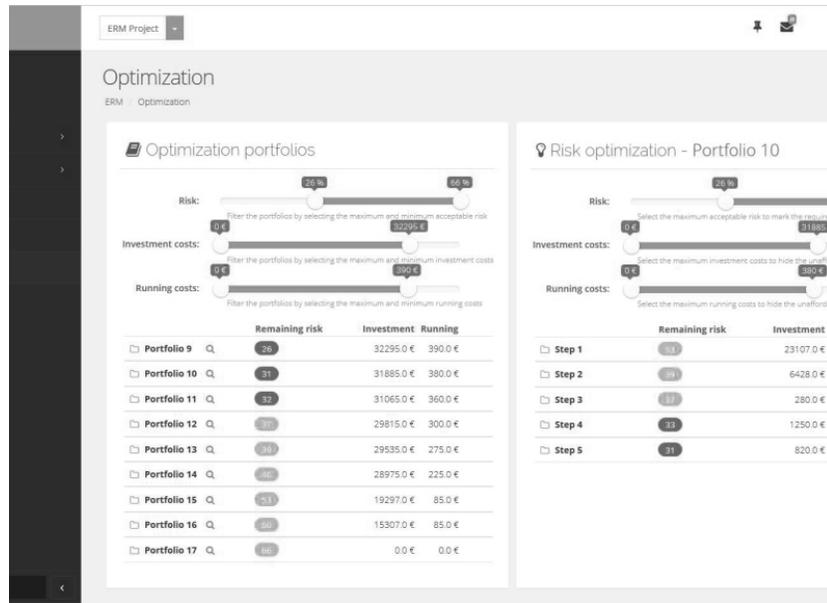


Fig. 7. AURUM

7 Discussion and further research

The purpose of this research is the development of methods to enable sustainable decisions by stakeholder inclusion. Especially in complex decision scenarios, stakeholders need decision options to be presented in a language they understand.

In order to process detailed domain knowledge and create comprehensible decision options in an automated way, we have developed an ontology engineering method that supports researchers and practitioners to efficiently build the

required ontologies independently of the application field. The decoupling approach translates detailed technical knowledge into financial/political dimensions and therefore enables decision makers to identify adequate options. The added value of the research results was assessed in an extensive evaluation phase, including several real-world deployments. The evaluation showed that the method supports the full cycle from data/knowledge assessment to the actual decision making, independent of the application field. Compared to traditional methods, the developed method provides (i) a broader range of technically feasible decision options, (ii) substantial time savings at identifying the decision options, and (iii) the possibility to visualize the decision options in a comprehensible way to relevant stakeholders. While users valued the time savings and comprehensible presentation of decision options, some users criticized the broad range of decision options. While all the options were compliant to the underlying rule sets and correct from a technical point of view, they sometimes deviate from common solutions. E.g., in the building industry there are best practices for constructions, i.e., how building materials are combined. In the some cases the building renovation decision support system produced options that included a technically correct, but uncommon, combination of building materials. This limitation will be addressed in further research.

Further limitations are a missing concept for integrating large data sets into the ontology and a solid approach for maintaining the ontology in a collaborative way. In further research we will work on these limitations and are planning to apply this method in the field of farming decision support where large data sets such as historic weather information play an important role in decision making. We will have to research on methods of how to aggregate this highly granular data to a level processable by ontology-based decision support systems. Furthermore, knowledge sources for decision making are becoming increasingly dynamic – data is added to the knowledge body in ever shorter time periods. For instance, software vulnerability information is updated several times a day. We will look into collaborative ontology editing methods and adapt them with the goal to enable and encourage people to contribute to the ontology maintenance process.

Acknowledgements

This research was funded by the FFG – Austrian Research Promotion Agency (COMET K1 center SBA Research).

References

1. Bouamrane, M.M., Rector, A., Hurrell, M.: Using owl ontologies for adaptive patient information modelling and preoperative clinical decision support. *Knowledge and information systems* **29**(2), 405–418 (2011)
2. De Nicola, A., Missikoff, M.: A lightweight methodology for rapid ontology engineering. *Communications of the ACM* **59**(3), 79–86 (2016)

3. Engelbart, D.C.: Augmenting human intellect: a conceptual framework (1962). PACKER, Randall and JORDAN, Ken. Multimedia. From Wagner to Virtual Reality. New York: WW Norton & Company pp. 64–90 (2001)
4. Feigenbaum, E.A., Buchanan, B.G.: Dendral and meta-dendral: Roots of knowledge systems and expert system applications. *Artificial Intelligence* **59**(1-2), 233–240 (1993)
5. Ferguson, R.L., Jones, C.H.: A computer aided decision system. *Management Science* **15**(10), B–550 (1969)
6. Gorry, G.A., Scott Morton, M.S.: A framework for management information systems (1971)
7. Keane, M., Thorp, S.: Chapter 11 - complex decision making: The roles of cognitive limitations, cognitive decline, and aging. *Handbook of the Economics of Population Aging*, vol. 1, pp. 661 – 709. North-Holland (2016). <https://doi.org/https://doi.org/10.1016/bs.hespa.2016.09.001>
8. Klaas, R.L.: A dss for airline management. *ACM SIGMIS Database* **8**(3), 3–8 (1977)
9. Liang, T.P., Yang, Y.F., Chen, D.N., Ku, Y.C.: A semantic-expansion approach to personalized knowledge recommendation. *Decision Support Systems* **45**(3), 401–412 (2008)
10. Mitroff, I.I.: Stakeholders of the organizational mind. Jossey-Bass Inc Pub (1983)
11. Musen, M.A., Middleton, B., Greenes, R.A.: Clinical decision-support systems. In: *Biomedical informatics*, pp. 643–674. Springer (2014)
12. Noy, N.F., McGuinness, D.L., et al.: *Ontology development 101: A guide to creating your first ontology* (2001)
13. Power, D.J.: Decision support systems: a historical overview. *Handbook on Decision Support Systems* 1 pp. 121–140 (2008)
14. Spyns, P., Tang, Y., Meersman, R.: An ontology engineering methodology for dogma. *Applied Ontology* **3**(1-2), 13–39 (2008)
15. Steele, K., Stefánsson, H.O.: Decision theory. In: Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016 edn. (2016)
16. Suarez-Figueroa, M.C., Gomez-Perez, A., Fernandez-Lopez, M.: The neon methodology framework: a scenario-based methodology for ontology development. *Applied ontology* **10**(2), 107–145 (2015)
17. Sure, Y., Staab, S., Studer, R.: Ontology engineering methodology. In: *Handbook on ontologies*, pp. 135–152. Springer (2009)
18. Tempich, C., Simperl, E., Luczak, M., Studer, R., Pinto, H.S.: Argumentation-based ontology engineering. *IEEE Intelligent Systems* **22**(6), 10 (2007)