

Assessment for Spreadsheets

Alexandru Toader & Michael Kohlhase
Jacobs University
Germany

Andrea Kohlhase
University of Applied Sciences Neu-Ulm
Germany

Abstract—Spreadsheets are tools to model personal, business and scientific data and to support decision-making based on such models. Moreover, spreadsheets are used for communicating data, models, and decisions, which makes the interpretability of spreadsheet intents an important consideration in the design of spreadsheet interaction.

In this paper we will focus on the problem of model assessment, i.e., the process of passing judgements on a situation based on modeled spreadsheet data, for instance, the development of profit over time represented in a controlling spreadsheet. We introduce a framework for specifying, documenting and personalizing assessments semi-formally and a set of plugins for the Semantic Alliance Architecture that can visualize assessments to assist the (spreadsheet) user in her decision-making.

I. INTRODUCTION: SPREADSHEETS AS DECISION-SUPPORT SYSTEMS

Spreadsheets are used to model quantitative real life scenarios and entire decision making processes have been built around them. In [1] they are even defined as “*software products that help users apply analytical and scientific methods to decision making*”. Here we take up this view and consider them special **decision support systems (DSSs)**. So far though, spreadsheet functionality does not directly support this quality, it rather assists indirectly by realizing calculations and enabling reports necessary for analyzing a situation and making decisions built on these. Thus, extending spreadsheet software to directly support its DSS quality would lower the hurdle and make DSS available to even more users, businesses, and institutions. We call systems that use spreadsheets for modeling and support the decision process **spreadsheet-based decision support systems (SSDSS)**, see [2], [3]. In a nutshell, an SSDSS assists users with two tasks: *i*) deciding whether a particular aspect of a model is favorable to the business goal, and *ii*) finding remedies if it is not. We call the first task **assessment** and we will concentrate on supporting it in SSDSS in this paper.

We can see assessment in spreadsheet documents as just another facet of the interpretation process of tabular data with respect to the situation modeled by the electronic ledger. Interpretation is a source of errors¹ that has been often neglected although it is of great importance in the case of long-lived spreadsheets. This is especially virulent in cases where the reader and the creator of a spreadsheet are different individuals, the reader is unfamiliar with the domain of the spreadsheet, and not all the knowledge of the problem domain

is transferred along with the document – a standard situation, when spreadsheets are used as a means of communication: An indication of this is given in a study that categorizes the help queries of users of a financial controlling spreadsheet in a large public/private research institution (see Figure 1). [6] reports that assessment-related questions make up 28% of the required explanation types.




Fig. 1. Types of Help Requests for a Financial Controlling System

This paper aims to reduce the occurrence of interpretation errors by enhancing the spreadsheet user’s assessment capabilities.

We build on the theoretical analysis in [6] and report on a concrete implementation by the first author in the scope of his B.Sc thesis [7] based on the “Semantic Alliance Framework” developed in our research group [8], [9]. This gives us a DSS that combines the ease and simplicity of modeling business data in spreadsheets with a knowledge-based assessment service.

In Section II we briefly review the concept “assessment”, in Section III the Semantic Alliance Framework, and in Section IV we present the assessment service itself. Section V evaluates the services via a use case, and Section VI concludes the paper and discusses future work.

II. ASSESSMENT IN SPREADSHEETS

Before devising an assessment service for spreadsheets, it is important to understand what assessment of spreadsheet data is and how it fits into the user’s workflow so that we can derive requirements for our solution. To ease reading we will not point out any longer that we are specifically concerned with assessment of spreadsheet data and in a spreadsheet context.

¹For general studies of spreadsheet errors see [4], [5].

Assessment consists of value and purpose judgments passed on situations modeled by the spreadsheet document or parts of it. Indeed, [6] classifies assessment queries in Figure 1 into two categories:

- *Assessment of Value* refers to the ability of a user to make judgments on the spreadsheet data after having interpreted the concrete values.
- *Assessment of Purpose* refers to the ability of a user to draw the proper conclusions or perform the correct actions based on the content of a spreadsheet document.

A. Running Example

To fortify our intuition about SSDSS consider the simple spreadsheet in Figure 2, which will serve as a running example for this paper. It is used to model the evolution of a small department of a multinational company over a number of years. It acts as a controlling tool on which all present and future decisions are based, thereby qualifying as a simple SSDSS.

Project Costs					
	Actual			Projected	
Year	2012	2013	2014	2015	2016
Revenues	486,000 €	620,500 €	793,000 €	1,090,000 €	845,100 €
Expense category	Expense				
Salaries	285,000 €	335,000 €	392,000 €	643,000 €	543,000 €
Utilities	25,300 €	64,500 €	89,000 €	95,000 €	80,000 €
Materials	123,000 €	153,000 €	175,000 €	250,000 €	210,000 €
Administration	32,500 €	45,200 €	56,200 €	72,000 €	64,200 €
Other	12,200 €	11,000 €	15,320 €	21,000 €	19,500 €
Total Expenses	478,000 €	608,700 €	727,520 €	1,081,000 €	916,700 €
Profit (Loss)	8,000 €	11,800 €	65,480 €	9,000 €	-71,600 €
%Profit (Loss)	1.65%	1.90%	8.26%	0.83%	-8.47%

Fig. 2. Running Example: A Simple Controlling Tool for Project Management

Let us use Figure 2 to introduce an important concept: We call a rectangular fragment of a spreadsheet a **functional block**, if it corresponds to a function in the modeled situation. For instance, cells [B6:D6] model the revenues as a function on the years 2012-2014. Cells and functional blocks are the basic building blocks of spreadsheet models and, therefore, the natural domains of assessment.

B. Assessment is Intensional

The following can be considered typical assessment statements in the context of Figure 2:

- “Row [6] looks good.”
- “The revenues look good.”
- “I like this (points to cell [D17]) but that (points to cell [F17]) is a disaster.”
- “I like the profit/loss in 2014 but of course not the expected one in 2016.”

On the surface, the first statement refers to a row in the spreadsheet, but if we look closer, then we see that this cannot

really be the case, since if we shift the whole spreadsheet by one row, then we have to readjust the assessment. So it has to be about the intended meaning of row [6], i.e., the development of revenues over the years. Indeed we can paraphrase *i)* with *ii)* — another clue that the assessments are really about situations modeled by a functional block in the spreadsheet. But assessments are not restricted to functional blocks as statements *iii)* and *iv)* only refer to individual cells. Note again that the statements are not about the numbers 65,480 and -71,600 (numbers in themselves are not good or bad, they just are). Here the assessment seems to be intensional, i.e., about the intension “the profit in 2014 is 65,480€” rather than the extension. Another way to view this is that the latter two assessments are about the argument-value pairs {2014;65,480} resp. {2016;-71,600}.

We will make this view the basis of our treatment of assessment: We need to have a model of spreadsheets, which can capture the intensional level, so that we can assess it. In the next section (part III-A) we will present spreadsheets with a structured background ontology that captures the spreadsheet intention, and in IV-A we show how we can extend that for assessing spreadsheet values and functional properties.

C. Assessment is Personal

But there is another aspect of assessment we need to model: assessments depend on context and who assesses the situation. Indeed, example *v)* below shows a statement relating assessments by three stakeholders – arguably enabling the addressee of the statement to form her own assessment from that.

- “Upper Management will be happy about the leftover funds in [nn] that they can now use elsewhere, but the PI of the project will be angry that he got less work out of the project than expected. Not to mention the funding agency; they cannot be told of this at all, because it violates their subsistence policy.”

In particular, we cannot tie assessments rigidly to particular cells and functional blocks, but need a setup, where assessments are objects in their own right that can be “applied” to cells and functional blocks. We consider these “assessment objects” as reified background knowledge about values and functions of the model expressed in the spreadsheet.

D. Assessment is Computational

Finally, we note the obvious: An assessment is computed on the values in the respective cell or functional block – according to the specification in the “assessment object” we have stipulated above.

All of these requirements do not fit with the conventional approach to extending spreadsheets by macros in the scripting language of the application (e.g., .Net for MS Excel or Java for LibreOffice). Instead we make use of a semantic framework technology we have previously developed and which we will present now.

III. THE SEMANTIC ALLIANCE FRAMEWORK

The entry point for the Semantic Alliance Framework we use as a basis for our Assessment Service is the realization, that spreadsheets per se only represent the data and computation layers of a model, and leave the intended meaning and the background knowledge about the situation modeled in the spreadsheet implicit – usually hinted at in the “legend” cells – see [6] for details.

A. Semantic Illustration

To remedy this shortcoming, we give the spreadsheet creators the opportunity to document the background knowledge – not in a human-readable “manual”, but in a machine-actionable **structured background ontology**². The background ontology concepts themselves can be linked to the meaning-carrying fragments of the spreadsheets. Figure 3 shows the approach.

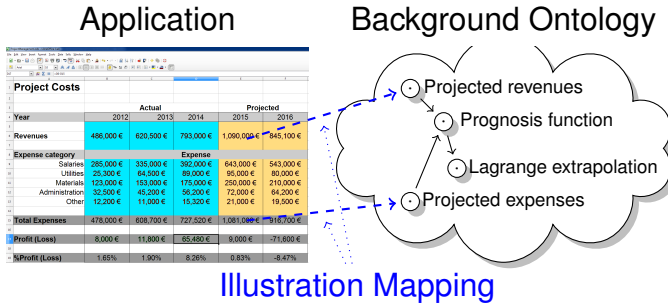


Fig. 3. The Semantic Illustration Architecture

In our system the background ontology is represented in the OMDoc [10] format, a markup format and data model for semi-formal, structured document collections. OMDoc provides a strong, logically sound module system based on structured theory graphs [11]. In contrast to other ontology modeling languages like OWL, the OMDoc format does not commit to a formal logical language. Therefore, it lacks a native concept of inference, but also does not force the author to fully formalize the ontology and to work around the expressivity limitations of the underlying logical system. Instead, OMDoc allows to locally formalize elements - and thus provide partial inference - with whatever formal system is most suitable.

After choosing a representation format for the ontology, the biggest challenge left is how to integrate it with the document and make it usable by the associated services. In [12] ERIKSSON suggested to combine documents and ontologies by “adding annotations to electronic-documents formats and including the ontologies in electronic documents” yielding semantic documents. This approach limits the reusability of the ontology for multiple documents that contain the same concepts and poses a problem for maintaining and enhancing

²We will use the term “ontology” in its broad meaning, where it stands for “a representation of the objects in a given situation, and their relations to each other”.

the information stored in the ontology. Enhancing the spreadsheet application with functionality to use a local ontology is prohibitively expensive. Therefore, an external ontology that can be used by multiple services and documents seems more sensible and is the approach taken by the Semantic Alliance Framework which will support the Assessment Service. This approach was materialized as the *Semantic Illustration architecture* in [13].

We will now give a brief overview of the Semantic Alliance Framework and explain how it supports semantic services like the Assessment Service.

B. Semantic Alliance Framework




Fig. 4. The Semantic Alliance Framework

The Semantic Alliance Framework is an architecture and software framework for semantic allies³. It combines an invasive approach that allows users to profit from semantic technology without leaving their accustomed workflows and tools with an application-independent way of extending applications with knowledge management technologies. The Semantic Alliance Framework provides the foundation for the development of the Assessment Service. As we will see later on, it offers functionality to tap into the user’s interactions with the document and to provide meaningful services in the appropriate context while giving the user the impression of (almost) perfect integration with the host application.

The Semantic Alliance Framework works by mashing up the GUIs of the knowledge management systems and applications themselves. The Semantic Alliance Framework has three components (see Figure 4):

- **Sally:** The main component of the mashup enabler, integrates the functionality of the application \mathcal{A} (e.g., spreadsheet application) and the semantic services \mathcal{S} (e.g., the Assessment Service) into a joint user interface and interaction model.
- **Alex:** The application \mathcal{A} is extended by a slim API (Application Programming Interface) “Alex” that reports and executes relevant user interactions within \mathcal{A} (e.g., cell clicks in spreadsheets) to and from **Sally**.

³Semantic allies are semantic systems that complement existing software applications with semantic services and interactions based on a background ontology.

- **Theo:** A screen-area manager that enables invasive design. Content supplied by \mathcal{S} is embedded as a pop-up into the GUI of \mathcal{A} .
- **Semantic Services:** are independent components which merge and process information from multiple sources including the machine-actionable ontology, the application \mathcal{A} and external sources. They return the result to the user through **Sally** (e.g., The Definition Service, which is offered through the Semantic Alliance Framework, when invoked returns the definition of a concept associated with a spreadsheet cell. For instance, when the user asks for the definition of cell [E15], he gets the definition of the concept *Projected Expenses*.)

Note that the Semantic Alliance Framework relies on the Semantic Illustration architecture – the machine-actionable background ontology is the basis for the semantic services and the illustration mapping (the green dashed arrows in Figure 3) show which parts apply at a given point in the spreadsheet. But from the point of view of the Semantic Alliance Framework, the background ontology is encapsulated into the semantic (web) services reducing the knowledge integration problem to a mashup of (web) services (this is what the Sally component does).

IV. THE VALUE ASSESSMENT SERVICE

In this section we will give a brief overview of how the Assessment Service is integrated into the Semantic Alliance Framework, describe the interaction between the user and the service, and discuss about the representation and computation of assessments with theory graphs. An Assessment Service differs from – say – a Definition Lookup Service which only relays and aggregates parts of the background ontology, in that it has to *compute* a verdict about possibly changing spreadsheet data based on the background ontology and feed it back to the user on the fly in an intuitive way.

A. Representing/Computing Assessments in Theory Graphs

The Semantic Alliance Framework organizes the background ontology into **theories** – collections of concept definitions together with statements of properties of and relations between the objects denoted by these concepts. Theories are connected by **views**, that is, meaning-preserving mappings that allow to interpret objects from the source theory as objects of the target theory. The **theory-graph** paradigm in OMDoc/MMT [10], [11] gives us a powerful, modular/object-oriented representation framework for formal and informal mathematical knowledge. Crucially, OMDoc/MMT allows for **parametric theories** and embedded computation in the MMT web service [14]: Representing assessment knowledge in (collections of) theories that can be connected to cells and functional blocks in spreadsheets meets the requirement of “object-like” assessment knowledge from subsection II-C. The object-oriented/parametric nature allows us to make general theories of “assessment patterns” and specialize them in the

illustration mapping as the Sally component that operationalizes it also contains an “abstract spreadsheet model” that can store concrete parameters.

Let us make this concrete with an example from Figure 2. Say we want to assess the projected profits with the intuition that profits are good as long as they are positive. As we want to assess a value in a functional block, we use a general assessment theory that takes a function f and an argument p as parameters. For the assessment of the value in cell [F17] we instantiate f to the profit function π and p to the year 2016. As π is the target concept in the background ontology for the functional block [B17:F17] under the illustration mapping and 2016 the argument of the functional block [B17:F17]. These instantiations were already in effect in the abstract spreadsheet model. For the particular assessment in our example, we make use of the “positive-balance-is-good” theory and bind it to the functional block [B17:F17]. This theory can be constructed by simply instantiating the general theory of “monetary quantities”, which inherits from “real numbers” (from elementary math). We obtain a comparison operator for the general theory of “partially ordered sets” (also from “general math”) and instantiate the base of the comparison (a parameter in our assessment model) to the real number zero. Given this illustration binding, the MMT API can instantiate all parameters in the formulae in the various theories and compute⁴ the assessment value. Here, the computation amounts to $\pi(2016) = -71.600 \text{ €}$ which is negative, leading to a negative assessment.

Note that we only had to provide the “positive-balance-is-good” theory⁵ and extend the illustration mapping to get the assessment for the full functional block – all other theories were so general that we can assume them to pre-exist in some Semantic Alliance ontology.

B. Integration into the Semantic Alliance Framework

The Assessment Service is just one of the many semantic services [8] offered through the Semantic Alliance Framework. For the purpose of this example, we will use LibreOffice Calc [15] as our host spreadsheet application and the document illustrated in Figure 2 as our running example.

The service registers itself with the framework at initialization and is available to the user each time he tries to assess a semantic object that is linked to an assessment rule.

When the Assessment Service is invoked on a cell range, it collects information from the spreadsheet application regarding the selected object, merges it with knowledge from the ontology, and it runs the assessment for the concrete values in the document as described above.

⁴In the declarative OMDoc/MMT framework nothing restricts the computation to two-value “good/bad” assessments: Instead of a predicate, we can specify functions into any assessment scale without changing the framework.

⁵Arguably, this theory is rather general as well. In our (admittedly limited) experience assessment is governed by a tractable number of theories that correspond to a limited number of assessment stereotypes.

C. User Interaction

The Assessment Service addresses both assessment needs: to be able to assess a value and the purpose. The former refers to an interaction on cell level, the latter requires an interaction concerning a functional block. If the service visualizes the assessment outcome, then the original layout should be regathered with a simple click. From the user’s point of view, therefore, the Assessment Service has to support three operations:

- Assess a value,
- Assess a domain (a group of cells), and
- Clear assessment results (i.e., revive the original spreadsheet state).

In order to assess the value of cell [F17], the user right-clicks on the cell and selects the Assessment Service “Assess value of F17” from the list of services offered by the framework for that particular object (see Figure 5). At this point, the service mashes-up information about the selected object from the spreadsheet and the knowledge base, and returns the result of the associated assessment theory evaluated for the object.

The result of an assessment is shown to the user by highlighting the assessed cell or cell ranges in bright green for a positive result, and in bright red for a negative result. At this point, LibreOffice locks down the sheets containing the colored cells so that no modification is possible. We can see the result of assessing the projected profit for the year 2016 in Figure 5.

Projected	
2015	2016
1,090,000 €	845,100 €
643,000 €	543,000 €
95,000 €	80,000 €
250,000 €	210,000 €
72,000 €	64,200 €
21,000 €	19,500 €
1,081,000 €	916,700 €
9,000 €	-71,600 €

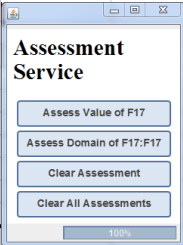


Fig. 5. Assessed Projected Profit

The process of assessing an entire functional block is almost identical. After, for example, the range [B17:D17] was selected and the Assessment Service “Assess Domain of B17:D17” was invoked, the system checks if the selected domain represents one homogeneous functional block. If there are cells in the block that are associated with a different concept, i.e., a different intended function, we cannot assess the block. If this sanity check passes, the process continues as in the case of the assessment of value scenario.

After performing an assessment the spreadsheet cannot be edited by the user. This locking is enforced in order to offer the user the possibility of returning the spreadsheet document to a pre-assessment state. In particular, this enables the “Clear Assessment” and “Clear All Assessments” options

of the Assessment Service (see Figure 5). After clearing all the results, the user can edit the document as before the Assessment Service was invoked.

Our Assessment Service provides a simple, intuitive interface that hides the formal theory structure and the evaluation of arbitrarily complex assessment rules. The theory graph that is the basis of the service can be easily extended with new assessment rules with a minimal implementation cost while incrementally increasing the value of the entire system.

V. USE CASES

In this section we use our running example to understand in what ways the Assessment Service will be of value in a real life scenario simple as it may be. While this is no substitute to a usability evaluation, it already gives us an indication of the potential of the new technology.

A strong point of our theory-graph-based approach is that it can handle different assessments – even adversarial ones – simply by associating different theories and views; composition and computation are supplied uniformly by the OMDoc/MMT framework. We expose this by assuming two users of the same spreadsheet document with distinct junior and senior perspectives and responsibilities, John and Robert. Neither of them is the author of the document, but both rely on it for assessing the past, present and future evolution of the company department modeled by the spreadsheet.

John is a young manager who has just taken over the department. He is in charge of managing the project’s team, schedule, and budget. On his first day on the job, John received the spreadsheet document in Figure 2. His challenge lies in understanding the data contained in the spreadsheet document and determining if there are changes needed to the division’s budget structure and business plan. **Robert**, John’s boss, is in charge of managing multiple divisions. He uses the same spreadsheet document, created by John’s predecessor, to obtain an overview of the project’s expenses state and to decide if an intervention in the management of the department is necessary.

So, both users have to be able to grasp very quickly financial risks based on the information summarized in the spreadsheet to make decisions and act fast, but neither of them has a full understanding of the spreadsheet.

Total Expenses	478,000 €	608,700 €	727,520 €	1,081,000 €	916,700 €
Profit (Loss)	8,000 €	11,800 €	65,480 €	9,000 €	-71,600 €
%Profit (Loss)	1.65%	1.90%	8.26%	0.83%	-8.47%

Fig. 6. Assessed Actual Profit Block

John wants to maximize resource consumption in order to maximize growth and the division’s potential market value, while maintaining a profit margin. The electronic ledger offers the Assessment Service maintained before by John’s predecessor. To get a quick impression on how the department has been doing up to current year, he selects the cell range [B17:D17] and the “Assess Domain of B17:D17” service. In turn, the range is colored in bright green (see Figure 6) reflecting the

fact that the profit has increased over time, which is reassuring to John. He could have chosen the service “Assess Value of B17:D17” showing the same result, this time caused by all values in the range being positive. Turning his attention to the Projected Profit running the Assessment Service “Assess value of E17” on the cell representing the Projected Profit for the year 2016, he directly perceives a problem by the subsequent red coloring of cell [E17] in Figure 5. This requires John to immediately take a closer look at the business plan of the division to determine whether the decisions taken by his predecessor were sound.

Robert starts by assessing the Actual Profit expressed in percentages of the Revenues. He considers that a division is reaching its target if the profit grows faster than the Dow Jones Industrial Average and he assigned the according assessment theory some time ago. Now, when he selects the entire Actual Profit range and chooses the “Assess Domain of B17:D17” service, the block’s color turns red and he has to take according counter measures.

Robert and John are both worried that the division is not reaching its full potential and they now know that something has to be done about that, but they are grateful that the Assessment Service has helped them in taking a fast decision on what problem area requires their interest.

VI. CONCLUSIONS & FUTURE WORK

We have presented a knowledge-based approach to assessment in spreadsheets and reported on our experiences of implementing an assessment service in the Semantic Alliance Framework. Given a small set of carefully constructed “assessment theories” (constructed by OMDoc/MMT specialists) and a larger set of theories that specify the background knowledge about the domain modeled by the spreadsheet (created by domain specialists), individual assessment views can be set by linking to specific theories via the illustration mapping (set up by each user). From these the MMT web service can compute assessments and visualize them to the spreadsheet user. Extended by assessment knowledge, spreadsheet-based models can in turn act as decision support systems.

Future Services

In Section V we hid the underlying multifaceted question “*Why is this assessment positive/negative?*”. On the one hand, it is concerned with how the assessment theory is defined, but on the other hand, it is concerned with the factors that determined the assessment to have a positive or negative result. The following two envisioned services will attempt to tackle both issues.

Assessment Explanation: To explain an assessment result the user has to understand the assessment theory behind it. As the computation in the MMT web service is inference-based, we can generate explanations of every step of the computation and the assessment. In fact, it should be relatively simple to generate assessment explanations specific to the particular cell or functional block, since the MMT API has access to all parameter instances.

Assessment of Dependencies: Recall the two tasks of a DSS mentioned in the introduction. With the assessment service presented in this paper, we have only addressed the first one. Finding remedies based on the assessments still remains unaddressed.




Fig. 7. Assessment of Dependencies

It is conceivable that the information in the assessment can be used for pinpointing causes of negative assessments and thus problems the remedies might attack.

Consider for instance the cell [F17] in the running example, which represents the Projected Profit for the year 2016. The Projected Profit is defined as “*Projected Revenues - Projected Total Expenses*”. It makes sense to trace the assessments through the dependency tree in an attempt to gather more insight about the assessment’s result.

In Figure 7 we can see the definition graph for the negatively assessed cell [F17] for the Projected Profit in year 2016. The node associated with Projected Revenue for the year 2016 is colored in red, meaning that it was also negatively assessed, so the problem might be with the revenues.

The JavaScript InfoVis Toolkit [16] provides all the necessary functionality to visualize the dependency graph and the Assessment Service already provides functionality for assessing the individual cells and functional blocks. At this point, the only missing functionality to implement this service is a Formula Parser that would transform spreadsheet formulae into a tree structure, which is left for further research.

ACKNOWLEDGMENT

We gratefully acknowledge fruitful discussions with Mihnea Iancu: his extensions of the MMT API made the modular computation of assessments possible. This work has partially been supported by the German Research Council under Grants KO 2428/11-1 and KO 2428/13-1.

REFERENCES

- [1] H. K. Bhargava, S. Sridhar, and C. Herrick, “Beyond spreadsheets: tools for building decision support systems,” *IEEE Computer*, vol. 32, no. 3, pp. 31–39, Jan. 3, 2002, [Online]. Available: <http://dblp.uni-trier.de/db/journals/computer/computer32.html#BhargavaSH99>.

- [2] M. M. Şeref and W. L. Ahuja Ravindra K. Winston, "Spreadsheet-based decision support systems," English, in *Handbook on Decision Support Systems 1*, Springer, 2008, pp. 277–298, DOI: 10.1007/978-3-540-48713-5_14.
- [3] V. Toader and V. R. Rus. (2014). Spreadsheet-based decision support systems, [Online]. Available: <http://steconomiceuradea.ro/anale/volume/2008/v4-management-marketing/275.pdf> (visited on 01/26/2014).
- [4] T. SH Teo and M. Tan, "Quantitative and qualitative errors in spreadsheet development," in *Proceedings of the 30th Hawaii International Conference on System Sciences: Information System Track-Organizational Systems and Technology - Volume 3*, ser. HICSS '97, Washington, DC, USA: IEEE Computer Society, 1997, pp. 149–, [Online]. Available: <http://dl.acm.org/citation.cfm?id=938435.938681>.
- [5] K. Rajalingham, D. R. Chadwick, and B. Knight, "Classification of spreadsheet errors," *CoRR*, vol. abs/0805.4224, 2008.
- [6] A. Kohlhase and M. Kohlhase, "Spreadsheets with a semantic layer," *Electronic Communications of the EASST: Specification, Transformation, Navigation – Special Issue dedicated to Bernd Krieg-Brückner on the Occasion of his 60th Birthday*, vol. 62, T. Mossakowski, M. Roggenbach, and L. Schröder, Eds., pp. 1–20, 2013, [Online]. Available: <http://journal.ub.tu-berlin.de/easst/article/view/870>.
- [7] A. V. Toader, *Assessment service for spreadsheet documents*, Bachelor's Thesis, 2014.
- [8] C. David, C. Jucovschi, A. Kohlhase, and M. Kohlhase, "Semantic Alliance: a framework for semantic allies," in *Intelligent Computer Mathematics*, J. Jeuring, J. A. Campbell, J. Carette, G. Dos Reis, P. Sojka, M. Wenzel, and V. Sorge, Eds., ser. LNAI, Berlin and Heidelberg: Springer Verlag, 2012, pp. 49–64, [Online]. Available: <http://kwarc.info/kohlhase/papers/mkm12-SALLY.pdf>.
- [9] C. Jucovschi, "Towards an interaction-based integration of mkm services into end-user applications," in *Intelligent Computer Mathematics 2014*, S. Watt, J. Davenport, A. Sexton, P. Sojka, and J. Urban, Eds., ser. Lecture Notes in Computer Science, Springer, 2014, pp. 344–356, [Online]. Available: <http://arxiv.org/abs/1403.5058>.
- [10] M. Kohlhase, *OMDOC – An open markup format for mathematical documents [Version 1.2]*, ser. LNAI 4180. Springer Verlag, Aug. 2006, [Online]. Available: <http://omdoc.org/pubs/omdoc1.2.pdf>.
- [11] F. Rabe and M. Kohlhase, "A scalable module system," *Information & Computation*, vol. 0, no. 230, pp. 1–54, 2013, [Online]. Available: <http://kwarc.info/frabe/Research/mmt.pdf>.
- [12] H. Eriksson, "The semantic-document approach to combining documents and ontologies," *International Journal of Human-Computer Studies*, vol. 65, no. 7, pp. 624–639, 2007, Knowledge representation with ontologies: Present challenges - Future possibilities, DOI: <http://dx.doi.org/10.1016/j.ijhcs.2007.03.008>.
- [13] A. E. Kohlhase and M. Kohlhase, "Semantic transparency in user assistance systems," in *Proceedings of the 27th ACM International Conference on Design of Communication*, ser. SIGDOC '09, Bloomington, Indiana, USA: ACM, 2009, pp. 89–96, DOI: 10.1145/1621995.1622013.
- [14] F. Rabe, "The MMT API: A Generic MKM System," in *Intelligent Computer Mathematics*, J. Carette, D. Aspinall, C. Lange, P. Sojka, and W. Windsteiger, Eds., ser. Lecture Notes in Computer Science, Springer, 2013, pp. 339–343, DOI: 10.1007/978-3-642-39320-4.
- [15] *Libreoffice*, <http://www.libreoffice.org/>, [Accessed on 21-April-2014].
- [16] *Javascript infovis toolkit*, <http://philogb.github.io/jit/>, [Accessed on 21-April-2014].