



JACOBS
UNIVERSITY

Assessment Service for Spreadsheet Documents

by

Alexandru Valentin Toader

Bachelor Thesis in Computer Science

Dr. Andrea Kohlhase, Prof. Dr. Michael Kohlhase
Name and title of the supervisor

Date of Submission: May 11, 2014

Jacobs University — School of Engineering and Science

With my signature, I certify that this thesis has been written by me using only the indicated resources and materials. Where I have presented data and results, the data and results are complete, genuine, and have been obtained by me unless otherwise acknowledged; where my results derive from computer programs, these computer programs have been written by me unless otherwise acknowledged. I further confirm that this thesis has not been submitted, either in part or as a whole, for any other academic degree at this or another institution.

Signature

Place, Date

Acknowledgements

I am most grateful to my supervisors, Andrea and Michael Kohlhase, for the valuable comments, ideas and guidance to the writing and development of this research. Furthermore, special thanks to Mihnea Iancu for the great guidance and direct support throughout the whole guided research.

Abstract

Spreadsheets are used by millions of individuals worldwide for business and personal purposes. They work with simple concepts as ranges and formulas that allow users to start working immediately, with no prior knowledge, but, at the same time, they offer advanced data analysis and mathematical functions.

Most spreadsheet users are not professional programmers and the number of errors in their documents exceeds by far the allowed error-rate in other domains. As spreadsheets hold an increasingly important role in the decision making process, errors can influence the outcome of complex real-world processes. Out of all the type of errors, interpretation errors stand out as an important but under-researched source of fallacy.

This thesis will tackle the issue of interpretation errors in spreadsheet documents in the context of assessment by developing an *Assessment Service* that will play the role of a decision support system while enhancing the semantic transparency of the application. The service will be offered in the context of the Semantic Alliance Framework and will be based on the concept of theory graphs.

Contents

1	Introduction	1
2	State of the Art	4
3	Preliminaries	6
3.1	Semantic Transparency	6
3.2	Semantic Illustration	6
3.3	Invasive Design	8
3.4	Semantic Alliance Framework	8
3.5	MMT	9
3.5.1	The MMT System	9
3.5.2	Category Theory Pushouts	10
3.6	The Universal OpenMath Machine	12
4	The Value Assessment Service	13
4.1	Integration into the Semantic Alliance Framework	13
4.2	MMT Assessment Theories	16
5	Use Case	19
6	Future Work and Conclusions	21

1 Introduction

A spreadsheet is a large sheet of paper that organizes business data in tables. Spreadsheets have been used in accounting for hundreds of years, but the widespread use of electronic spreadsheets only happened after the personal computer boom in the 1980' and the introduction of the VisiCalc [1] software. Electronic spreadsheets are interactive documents that are heavily employed in administration, financial forecasting, education and science.

Spreadsheets are used to model real life scenarios and entire decision making processes have been built around them. In order to understand the role they play as a decision support tool and the benefits and dangers that their use brings, let us look at a concrete decision making process, which will serve as a running example for this thesis.

Running Example

Project Costs					
	Actual			Projected	
Year	2012	2013	2014	2015	2016
Revenues	486,000 €	620,500 €	793,000 €	845,100 €	1,153,000 €
Expense category	Expense				
Salaries	285,000 €	335,000 €	392,000 €	543,000 €	643,000 €
Utilities	25,300 €	64,500 €	89,000 €	80,000 €	95,000 €
Materials	123,000 €	153,000 €	175,000 €	210,000 €	250,000 €
Administration	32,500 €	45,200 €	56,200 €	64,200 €	72,000 €
Other	12,200 €	11,000 €	15,320 €	19,500 €	21,000 €
Total Expenses	478,000 €	608,700 €	727,520 €	916,700 €	1,081,000 €
Profit (Loss)	8,000 €	11,800 €	65,480 €	-71,600 €	72,000 €
%Profit (Loss)	1.65%	1.90%	8.26%	-8.47%	6.24%

Figure 1: Running Example: A Simple Project Management System Using LibreOffice

Alin C. was the manager of a small department of a multinational company. He used the spreadsheet in Figure 1 to model the division's financial evolution and to make managerial decisions based on that information. Sadly, Alin is not part of the company anymore, but the spreadsheet he created still functions as a starting point for all the decisions regarding the present and the future of his former division. Although the document Alin used is still available to his successor and former superior, the rules he used to judge the evolution of his department are not readily apparent.

John R. is a young manager that has just taken over Alin's former 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 1. As an ambitious manager, John wants to maximize resource consumption in order to maximize growth while at the same time maintain a profit margin. 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 S., John's boss, is in charge of managing multiple divisions and redistributing funds in order to maximize global efficiency and, at the same time, helping struggling projects survive long enough to reach profitability. He uses the same spreadsheet document, created by John's predecessor, to obtain an overview of the project's expenses status and decide on funds distribution. He has a general view on how a division should evolve, how much profit it should make and how fast the profits should grow over time. If a department's evolution does not respect his time-tested blueprint for success, he has to intervene in the division's management. In essence, both users, have to make decisions based on the information in the document in Figure 1, but neither of them has a full understanding of the data in the spreadsheet.

Spreadsheets as Decision Support Systems

As we can see from John's and Robert's work-flows, spreadsheets play an important role in their decision making process. The spreadsheet software gains the attributes of a decision support system. According to [2] decision support systems (DSSs) are *"software products that help users apply analytical and scientific methods to decision making"*. By using models and algorithms from multiple disciplines, DSS products can execute, interpret, visualize, and interactively analyze the modeled situations.

Decision support systems assist the user in deciding whether a particular aspect of a problem is favorable, or not, to the business goal he has. In essence, it all scales down to the simple question "Is this good for my business?". As spreadsheets are used to model real life scenarios, the question above translates into "Is this value or block good for me?".

The problem with building a standalone DSS is that it requires significant expertise in decision analysis, software development and user interface design. For these reasons, the broad use of DSSs has not occurred [2]. In contrast, a great number of professionals are using spreadsheet software in order to realize calculations, analyses and reports necessary for the decision making process. Decision support systems that are built using spreadsheet software can be called spreadsheet-based DSS [3]. Significant work has been done on the topic of developing spreadsheet-based decision support systems ([4], [3]), but they still require a complex development process and knowledge that often falls outside the training of the average spreadsheet user [4].

In essence, decision support systems aim to provide automatic assessment of modeled situations. As the situation is modeled by the spreadsheet document, the two managers' work can be simplified by offering them a service that provides assessment functionality for spreadsheet applications. Before devising such a service, it is important to understand what assessment is and how it fits into the user's work-flow. Assessment consists of value and purpose judgments passed on situations modeled by the spreadsheet document or parts of it. I will consider the spreadsheet in Figure 1 as a starting point for uncovering the different aspects of assessment. The following can be considered typical assessment statements:

1. "The revenues look good."
2. "The profits in the years 2012-2015 look good, but not the one in 2016".

3. "The values in cells B17 to E17 look good, but not the one in F17."
4. "The labor costs have increased faster than the revenues."
5. "The evolution of profits over time, doesn't look promising."

The first assessment refers to the fact that the values in the revenue row are positive, but this can only be ascertained by assessing the purpose of the respective semantic item. We can consider 2. and 3. as equivalent statements, but we must underline the fact that they do not refer just to the numbers in cells B17:F17, but the values in the context assigned by the author of the document, the size of the profit in the years 2012-2016. The statements 4. and 5. refer to the whole range containing profit values, to the evolution of a value over time. In this case, the behavior of a function is the one assessed, not the individual values.

In essence, assessment in spreadsheet documents is just another facet of the interpretation process of tabular data with respect to the situation modeled by the electronic ledger. A service that would provide assessment functionality would in fact enhance the user's interpretation capabilities and would have a great impact on reducing the number of errors in the spreadsheet document. The topic of spreadsheet errors has been the subject of numerous studies [5], [6], [7], [8], [9]. Interpretation is a source of errors that has been often neglected although it is of great importance in the case of long-lived spreadsheets. As in our example, the reader and the creator of a spreadsheet can be different individuals and not all the knowledge of the problem domain is transferred along with the document. The misinterpretation of data can lead to erroneous decisions and an exponential propagation of errors in newly created documents.

One of the main causes of interpretation errors is the inherent lack of semantic transparency¹ of spreadsheet applications. In essence, this is a problem of knowledge management which can be tackled using knowledge management technologies, by recovering the content and semantics from the document and exploiting it for automation.

This thesis aims to reduce the occurrence of interpretation errors by enhancing the spreadsheet user's assessment capabilities. To that end, I will develop a decision support system that combines the ease and simplicity of spreadsheet use with an assessment service based on theory graphs in order to provide a tool for fast and efficient decision making. Typical decision support systems are made up of four components: data management, model management, knowledge management and user interface management [11]. By illustrating spreadsheet documents with a semi-formal background ontology², we can provide the model, data, and knowledge management components. The user interface component can be provided by the Semantic Alliance Framework [13].

The proposed service could assist decision-makers in problems involving risk management, the allocation of scarce resources, and the need to balance conflicting objectives. The *Assessment Service* will give the user access to the assessment knowledge used by the creator of the spreadsheet while enhancing the semantic transparency of the application and reducing the potential for error in the decision making process.

¹ A user interface is "semantically transparent, if it enables a user to access its semantic objects and their relations via the corresponding UI objects" [10].

²An ontology formally represents knowledge as a hierarchy of concepts within a domain. [12]

2 State of the Art

In this section, I will give an overview of the avant-garde of research performed on the topic of improving the interpretation capabilities of spreadsheet users and on the topic of creating spreadsheet based decision support systems. In order to understand why assessment knowledge is essential for a better understanding of spreadsheet documents I analyze the findings presented in *Spreadsheets with a Semantic Layer* by KOHLHASE and KOHLHASE [14].

The researchers performed an interview with a DFKI³ expert on the topic of the DFKI Controlling System with the goal of discovering what type of help a user assistance system must offer and where it should be offered. The study found that in order to fully understand the meaning of a knowledge item the user has to know and understand up to six of the following seven aspects:

- Definition (Conceptual)
- Purpose (Conceptual)
- Assessment of Purpose
- Assessment of Value
- Formula
- Provenance
- History

Figure 2 portrays the distribution of occurrences according to each type. We see that *Assessment of Purpose* and *Assessment of Value* explanations make up 28% of the required explanation types.

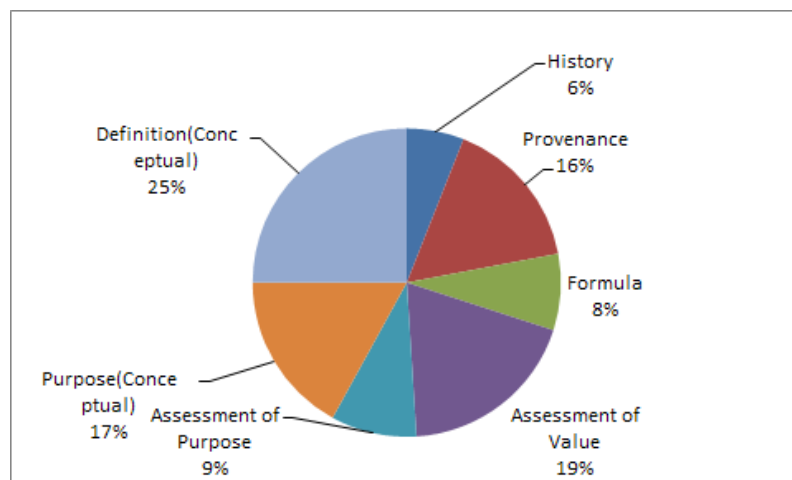


Figure 2: Types of explanations [14]

Assessment of Purpose refers to the ability of a user to draw the proper conclusions or perform the correct actions given the purpose of a knowledge item in a spreadsheet.

³German Research Center for Artificial Intelligence

Assessment of Value refers to the ability of a user to make judgments on the data in the spreadsheet after interpreting the concrete values.

We see that standalone spreadsheet applications are relatively semantically opaque and there is a need for external support in the form of a user assistance system in order to alleviate this problem. The challenge of creating a semantically transparent user assistance system, that offers a level of help similar to a human, has been tackled in the SACHS [15] and SiSsl⁴ [16] projects. In particular, the implemented SACHS system, together with the spreadsheet player, is able to give help for 33% of the explanations found in the help of a human expert [14]. These two ontology-based systems are well-suited for Definition and Purpose explanations, but there is no support whatsoever for *Assessment of Purpose* and *Assessment of Value* explanations and functionality.

A theoretical solution for an Assessment of Value service is proposed in *Spreadsheets with a Semantic Layer*. The background ontology that supports the *Definition lookup* service is extended by a set of assessment theories that judge the intended functions in the semantic blocks of the spreadsheet on their functional properties. The authors propose that the assessments are formalized as formulas which can be evaluated by the spreadsheet player.

The proposed solution has a few shortcomings that I have tried to overcome in my guided research. First of all, including the assessment theory into the definition of a concept is impractical and impure from an ontological point of view. The assessment of a semantic block does not only depend on the concept it is associated with, but also on the community of practice, the user's perspective and an array of other factors external to the ontology. Second, formalizing assessments as spreadsheet formulas bears a great implementation and portability cost i.e. spreadsheet formulas are not 100% portable between spreadsheet applications so performing the same evaluation in different environments might give inconsistent output.

The concepts presented in *Spreadsheets with a Semantic Layer* will be used as a starting point for the development of a value assessment service. The service will be offered as part of the Semantic Alliance Framework, thus becoming available for multiple spreadsheet applications and operating systems with minimal effort, and it will offer a system independent formalization and evaluation of assessments.

Significant work has also been performed in the domain of spreadsheet-based decision systems [4] [3], but the main focus has been on enabling spreadsheet creators to build better models with their data. All the solutions I have found so far require major changes to the structure of the workbook, changes, that left undocumented, will increase the opacity of the knowledge items, and will enhance the potential for error. To the extent of my knowledge, there is no method that allows a user to build or use a DSS around an existing spreadsheet document without modifying its structure.

The *Assessment Service* I will present in the following sections offers the functionality of a decision support system to spreadsheet creators and readers and enhances the spreadsheet application's semantic transparency.

⁴SiSsl is an instance of the Semantic Alliance Framework

3 Preliminaries

In this section, I will present an overview of the systems and frameworks that will be used in realizing the Assessment Service, as far as it is necessary to make this material self-contained. The service will be offered in the context of the Semantic Alliance Framework, but it will use multiple technology stacks and concepts to accomplish its goals.

3.1 Semantic Transparency

As stated before, one of the main goals of this thesis is to improve the semantic transparency of the spreadsheet applications that will use the Assessment Service. Before seeing how this service will improve the interaction potential of spreadsheet applications, we first need to understand what "semantic transparency" is.

In Linguistics the term "*semantic transparency*" is defined as "*a descriptive phrase that has been used [...] to describe endocentric compounds. Endocentric compound words are those whose whole meaning can be figured out by an analysis of its parts of morphemes*" [17]. In Artificial Intelligence, semantically transparent systems are described as "*ones that directly encode and exploit the kinds of information that a human agent might consciously access when trying to solve a problem*" [18]. From the Human-Computer Interaction standpoint, the perspective from which we tackle this issue, "*semantic transparency*" describes user interfaces that "*allow people to interpret them properly*" [19].

At this point, it is not completely clear how a user interface can be "interpret(ed) properly", so let us examine this issue further. A user interface is composed of meaning carrying objects that convey its intention, its purpose. These meaning carrying objects, which in the case of spreadsheets are cells, formulas, functional blocks, legends etc., are called **semantic objects** [19]. A user interface is **semantically transparent** if "it enables a user to access its semantic objects and their relations via the corresponding UI objects" [19].

The spreadsheet in Figure 1 seems transparent enough, but as soon as we ask:

- what is the definition of "projected revenues"
- should the profit be negative?
- should the expenses increase faster than the revenue over time?

we realize that we are missing implicit knowledge that is necessary to properly interpret the meaning of the semantic objects that compose spreadsheet.

By making this background knowledge explicit and using it through the Assessment Service and other semantic services, questions similar to the ones above can be automatically answered, thus exposing a higher level of semantic transparency to the user.

3.2 Semantic Illustration

In order to provide the Assessment Service and other semantic services, we need to be able to access and use the background knowledge of the spreadsheet document. As it is impossible to always ask the creator of the spreadsheet questions about the missing information, we try to store this implicit knowledge into a semi-formal domain ontology.

By using an ontology, the knowledge can be used for multiple purposes and by multiple services at the same time.

OMDoc [20] is a markup format and data model for **Open Mathematical Documents** that will be used to write the ontology. OMDoc is an XML-based, semantics-oriented format for representing semi-formal, structured document collections. OMDoc provides a strong, logically sound module system based on structured theories [21]. In contrast to other ontology modeling languages like OWL [22], the OMDoc format does not commit to a formal logical language, and therefore 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.

The implementation of the Semantic Alliance Framework that will support the Assessment Service uses an ontology represented in OMDoc and builds semantic services on top of it. OMDoc is an XML based format, therefore it is hard for humans to both author and read. In order to facilitate the development of complete and complex ontologies, sTEX [20] is used. sTEX is a collection of macro packages that allow semantic markup in T_EX and L^AT_EX documents. Using special tools [23], the semantic T_EX document is transformed into a machine understandable format like OMDoc. In the SiSsl [16] project, the ontology is written in sTEX and transformed into the OMDoc format.

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 semantic services. ERIKSSON [24] 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 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* [19].

In the Semantic Illustration architecture, semantic technology is made fruitful by "illustrating" existing software artifacts semantically via a mapping into a structured background ontology. This approach comes in contrast to the Semantic Web one where information resources are enhanced into semi-formal ontologies by annotating them with formal objects.

One of the main advantages of the Semantic Illustration approach is that it makes the ontology reusable and it avoids the high cost of reauthoring the ontology for each usage. Each document only needs to hold a mapping between the semantic objects and a link to an ontology term (e.g. URI). The same ontology can be shared across multiple documents of different types and services can use the background knowledge with disregard to application specific features or the document structure.

We will now look at the concept of Invasive Design [13], which together with Semantic Illustration, form the basis of the Semantic Alliance Framework.

3.3 Invasive Design

Invasive Design [13] is a new approach to software design that aims at obtaining a similar effect on the user as Invasive Technology [15] but with a significantly smaller implementation cost. From the user's perspective, a service "feels" embedded into an application if it occupies a screen area that overlaps with that of the application from which it was invoked. This perception is strengthened if the service information is displayed in the visual vicinity of the object on which it was invoked. This type of embedded services are called contextualized services.

In contrast to the Invasive Technology approach, where the framework components are actually embedded into the application, the Invasive Design approach relies on the perception of the user with regard to embedding. The document player is only invaded through a thin client that reports the user's actions to a physically independent framework.

In the next section, we will see how Semantic Illustration and Invasive Design work together in the Semantic Alliance Framework into which the Assessment Service will be integrated.

3.4 Semantic Alliance Framework

The Semantic Alliance Framework is an architecture and software framework for semantic allies⁵. It combines an invasive approach that allows users to profit from semantic technology without leaving their accustomed work flows 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.




Figure 3: The Sally Framework [13]

The Semantic Alliance Framework is based on two concepts: *Semantic Illustration* and *Invasive Design*. It works by mashing up the GUIs of the knowledge management systems and applications themselves. The Semantic Alliance Framework has three big components (see Figure 3):

- **Sally**: The main component of the mashup enabler, integrates the functionality of the application \mathcal{A} and the semantic services \mathcal{S} into a joint user interface and interaction model.
- **Alex**: The application \mathcal{A} is extended by a slim API⁶ "Alex" that reports and executes relevant user interactions within \mathcal{A} like cell clicks in spreadsheets to and from **Sally**.

⁵Semantic allies are semantic systems that complement existing software applications with semantic services and interactions based on a background ontology.

⁶Application Programming Interface

- **Theo:** A screen-area manager that enables invasive design. S supplied content is embedded as a "pop-up" into the GUI of A .

3.5 MMT

MMT [21] is a knowledge representation format that focuses on logic-independence and modularity. It is designed to be foundation independent, and applicable to all base languages based on theories. MMT plays a central role in the current implementation of the Assessment Service. It allows to formally represent assessment rules in a way that allows automatic evaluation and inference.

MMT concepts are distinguished into four levels: the document, module, symbol and object level. Documents act as open packages using logical identifier. They provide namespace management and every document may refer to every other document as long as referential cycles are not created. Documents contain **modules**, and MMT uses only two kinds of module declarations: **theories** and **views**.

Theories are defined by a set of typed **symbols** and a set of **axioms** describing the properties of the symbols. MMT uses the **Curry-Howard representation** to drop the distinction between symbols and axioms, and, as a result, MMT uses only theories and theory morphisms.

An MMT theory consists of symbol declarations. **Constants** represent declarations of the base language. A constant assignment provides a T-term for an S-constant and structure assignments provide a T-morphism for an S-structure. A **signature morphism** $\sigma : S \rightarrow T$ that translates all the symbols and theorems of S into objects and theorems in T is called a **theory morphism**. MMT represents theory morphisms as **structures** and **views**. An MMT **view** from S to T consists of a list of assignments from $c \rightarrow \omega$ where c is an S-constant and ω is a T-term. An MMT **structure** is a named import and it can be used to show that a theory inherits concepts from another theory. **Terms** form the MMT object level.

Theories and the relations between them form graphs where theories are nodes and theory morphisms are arrows. Theory graphs and the operations that can be performed on them via MMT plugins, will play a central role in the development of the Assessment Service by allowing the creation of derived theories and the evaluation of concrete views of existing theories.

Knowledge concepts represented in MMT documents can be uniquely identified via URIs as follows. Document identifiers doc are URIs without queries or fragments. Module identifiers, $doc?mod$, are formed from a document identifier and a module identifier separated by a "?" . Symbols are uniquely identified by a module identifier $doc?mod$ and a symbol identifier sym , separated by a "?" i.e. $doc?mod?sym$. MMT URIs will be used to enhance the semantic illustration by associating data items in the spreadsheet (cells, ranges) with theories that will be used for assessment.

3.5.1 The MMT System

The MMT System provides an API that uses the small number of primitives in MMT. This API is application-independent as it focuses on the data model of MMT in order to be

easily integrable with specific applications. The capabilities of the MMT System can be extended by writing small plugins that customize the behavior to one specific foundation.

The system provides a front-end in the shape of an HTTP server. This exposes all API functionality and knowledge management services via a RESTful [25] interface.

The Assessment Service will use the MMT System and its plugins via the RESTful interface to evaluate data contained in the spreadsheet in the context of corresponding MMT theories.

3.5.2 Category Theory Pushouts

The pushout mechanism offered by MMT is the central functionality that has allowed the creation of the Assessment Service. It allows the framing of arbitrary knowledge concepts and associated data in terms of established assessment theories.

Category theory is defined as a formalization of mathematics and its concepts as a *collection of objects and morphisms* [26].

A category \mathcal{C} consists of four kinds of data subject to three axioms [27]:

- **Data:**
 - **Objects**
 \mathcal{C} has objects $A, B, C \dots$
 - **Morphisms**
For each ordered pair of objects A, B in \mathcal{C} , there is a class of morphisms from A to B denoted by $Hom_{\mathcal{C}}(A, B)$.
 - **Composition**
For each ordered triple of objects A, B, C in \mathcal{C} , there is a law of composition: If $f : A \rightarrow B$ and $g : B \rightarrow C$, then the composite of f and g is a morphism $g \circ f : A \rightarrow C$.
 - **Identity**
For each object A there is a designated identity morphism on A , notated as 1_A , $1_A : A \rightarrow A$.
- **Axioms:**
 - **Unique typing**
 $Hom(A_1, B_1)$ and $Hom(A_2, B_2)$ are disjoint unless $A_1 = A_2, B_1 = B_2$.
 - **Associative Law**
 $h \circ (g \circ f) = (h \circ g) \circ f$
 - **Identity is a "neutral element"**
For the identity morphism $1_B : B \rightarrow B$ associated to each object B , two equations must hold for each pair of objects A and C and each pair of morphisms $f : A \rightarrow B, g : B \rightarrow C$: $1_B \circ f = f$ and $g \circ 1_B = g$.

Given two morphisms $f : A \rightarrow B$ and $g : A \rightarrow C$ in a given category \mathcal{C} , a **pushout** is a triple (D, α, β) with $\beta \circ g = \alpha \circ f$ that satisfies the following universal property: for every

triple (Y, α', β') with $\beta' \circ g = \alpha' f$, there exists a unique morphism $\theta : D \rightarrow Y$ making the diagram commute [28].




Figure 4: Pushout [28]

The concept of pushouts is mapped into MMT [29]. Given a view $\sigma : \Sigma \rightarrow \Sigma'$ and a structure $\Sigma \leftrightarrow \Sigma, \Gamma$, the pushouts in MMT are defined by:

$$\begin{aligned} \text{dom}(\sigma(\Sigma, \Gamma)) &= \text{dom}(\Sigma') \cup \text{dom}(\Gamma) \\ \sigma(\Sigma, \Gamma)(c) &= \begin{cases} \Sigma'(c) & \text{if } c \in \text{dom}(\Sigma') \\ \sigma^{\Sigma, \Gamma}(E) & \text{if } c \in \text{dom}(\Gamma) \end{cases} \\ \sigma^{\Sigma, \Gamma} : c &\rightarrow \begin{cases} \sigma(c) & \text{if } c \in \text{dom}(\Sigma') \\ c & \text{if } c \in \text{dom}(\Gamma) \end{cases} \end{aligned}$$

Given morphisms $\sigma_1, \gamma_1 : \Sigma, \Gamma \rightarrow \Phi$ (where $\sigma : \Sigma \rightarrow \Phi$) and $\sigma_2 : \Sigma' \rightarrow \Phi$, the universal morphism $\sigma(\Sigma, \Gamma) \rightarrow \Phi$ is σ_2, γ_1 .




Figure 5: Pushout in MMT [29]

At the moment we have a very abstract description of what a pushout is, but let us make it a bit more concrete with an example. $\text{PosOrd}_>$ is the theory of positively ordered sets. \mathbb{N} , the set of natural numbers is a view of $\text{PosOrd}_>$. $\text{OrdList}_>$ represents positively ordered

lists. From the theories \mathbb{N} and $OrdList_{>}$ MMT generates the pushout theory $OrdList_{>}(\mathbb{N})$. In our generated theory we can use all the axioms from \mathbb{N} and $OrdList_{>}$ as we know they hold true.




Figure 6: Pushout over \mathbb{N} and $OrdList_{>}$

3.6 The Universal OpenMath Machine

MMT will be used to express the assessment theories, but these theories still need to be evaluated in the context of the spreadsheet document. The evaluation will be performed using the MMT system extension developed by Felix Mance for his bachelor thesis [30], which integrates the Universal OpenMath Machine built by Vladimir Zamdzhiev [31].

The Universal OpenMath Machine is a software system that can execute Scala [32] implementations of OpenMath [33] objects embedded in an OMDoc [34] document and simplify OpenMath expressions.

Felix’s work gave the MMT system the attributes of a declarative, deductive and computational system. Felix’s Practical OpenMath Machine, offers the power of a proof assistant system⁷ and a computer algebra system⁸ without any of their individual weaknesses. This system is composed from two subsystems. The first subsystem integrates the declarative and computational realms and the second subsystem integrates the declarative and deductive realms. By using MMT as the common declarative system of the two subsystems, the two components can be connected and become a complete, fully extensible declarative, deductive and computational system with support for arbitrary logics and programming languages.

The Universal OpenMath Machine will be used by the Assessment Service to compute the result of applying an assessment theory in a particular context. Without the Assessment Service, if John tries to decide if the profit is good or bad, he manually applies the assessment theory "A profit value is good if it is positive." to the value of the cell of interest i.e. he checks if the value of the cell is bigger than 0. But, with the help of the Universal OpenMath Machine, we can transfer the task of evaluating the assessment theory to the application.

In the following section I will explain how all these systems work together to realize the current implementation of the Assessment Service.

⁷"A software tool to assist with the development of formal proofs by human-machine collaboration" [35]

⁸"A software program that allows computation over mathematical expressions" [36]

4 The Value Assessment Service

In the first part of this section we will see how the Assessment Service is implemented, how it is integrated in the Semantic Alliance Framework and how it interacts with the MMT system. In the second part, we will look into what an "assessment theory" is and how it is represented and evaluated in MMT.

The Assessment Service will essentially support three operations from the user's point of view:

- Assess a cell
- Assess a functional block (a range of cells)
- Clear assessment results

4.1 Integration into the Semantic Alliance Framework

In order to understand how the Assessment Service works and how it is integrated into the Semantic Alliance Framework I developed each of the mentioned operations and broke down the actions of the individual components of the framework and the message flow between them, switching from the user's perspective to the system's perspective as often as it is necessary for a full comprehension.

We will use LibreOffice Calc [37] as our host spreadsheet application and the document illustrated in Figure 1 as our running example. For this service, the Alex add-in for LibreOffice Calc has been extended with functionality to highlight cell blocks and to lock the editing of the document.

The Assessment Service was created in Java [38] as a sub-module of the Sally component of the framework. The service registers itself with the framework at initialization and will be available to the user each time he tries to assess a semantic object that is linked to an assessment theory. A semantic object can be linked to an assessment theory by keeping the association in the Abstract Spreadsheet Model⁹(ASM) associated with the document.

Let us first look at the process of assessing a cell. For this purpose, we go back to our running example and see how John might use the Assessment Service. John will select the cell he wants to work with and right click it. At this point he has to select "Show Frames" from the context menu. A message containing the position of the selected range is sent to Sally. Sally looks through the list of registered services, selects all the services that apply and presents the user with a list of Frames(see Figure 7) in which these services can be used. The flow of messages between the different components can be seen in Figure 8.

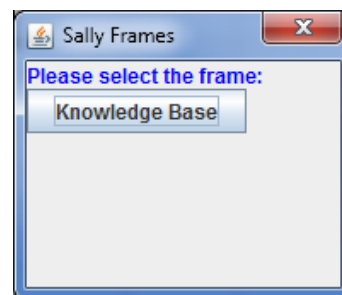


Figure 7: Sally Frames

⁹A software component that keeps information about the mapping between semantic objects and ontology concepts and assessment theories.

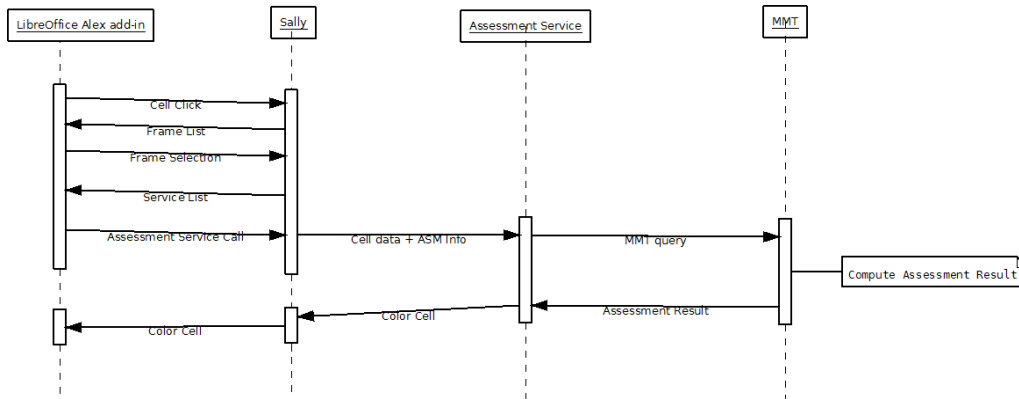


Figure 8: Value Assessment Process

The user will now select the "Knowledge Base" frame and then select the "Assessment Service" from the list of options (see Figure 9). The assessment module in Sally will now make use of information from the Abstract Spreadsheet Model (ASM) to decide what kind of functionality it can offer to the user. The user will then see a list of functionality offered by the Assessment Service for the selected cell.

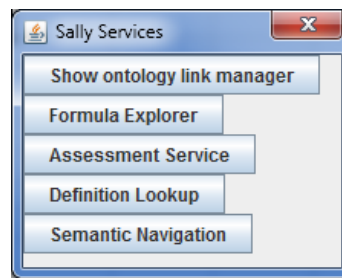


Figure 9: Knowledge Base Services

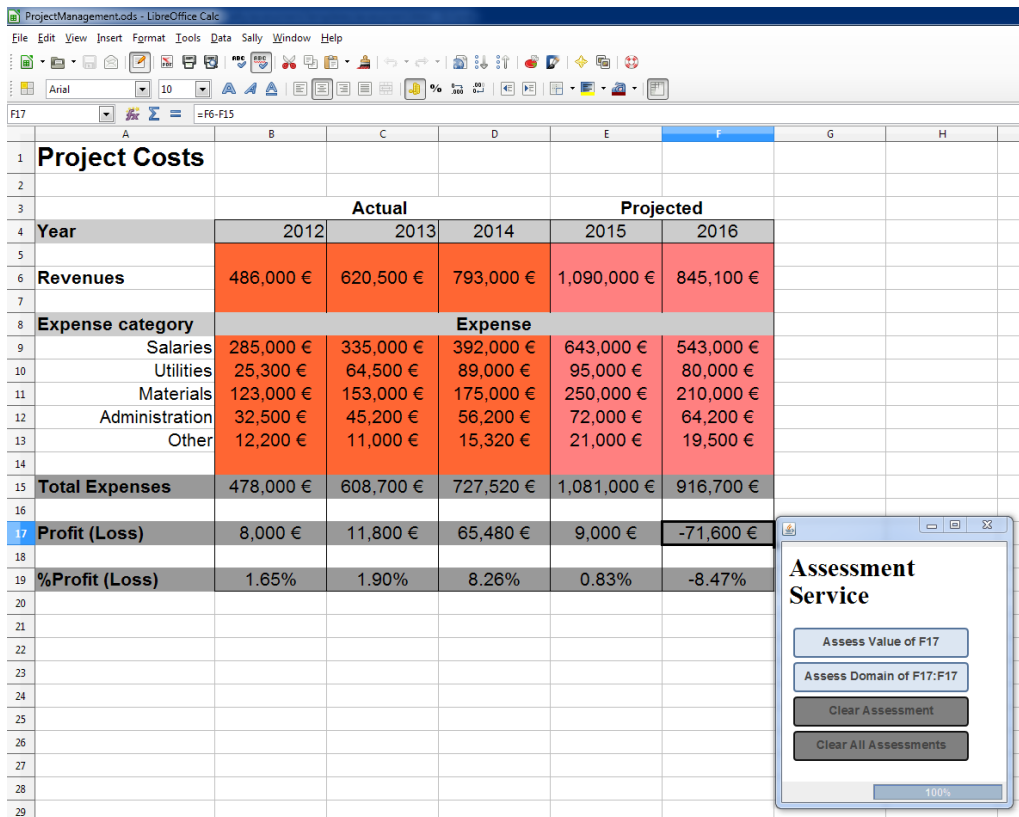


Figure 10: Assessment Service

In order to assess the value in the selected cell, the user will click "Assess Value of " (see Figure 10). Before performing the assessment, the services has to determine the cell's intended function. This is realized by retrieving the required information from the concept associated with the cell. As the association is kept in the ASM as a URI, this operation requires only accessing the ontology using the respective URI. The ASM also holds information about the assessment theories associated with each cell. The next step is to use the concept and the assessment theory associated with this cell to assess the value of interest.

The service will now send a POST [39] request to the MMT system that is running in a separate process to assess the value of the selected cell. This POST is addressed to the URIs of the assessment theory and it contains information about the concept associated with the cell and the value of the cell. We will take a closer look at how an assessment theory looks and how it is evaluated in the next section.

The MMT system will perform the assessment and return a boolean value, indicating whether the evaluation was positive or not, to the Assessment Service. Depending on the result, the Assessment Service will send a request to the Alex add-in in LibreOffice to color the cell of interest. The result of an assessment is shown to the user by highlighting the assessed cells or ranges in bright green for a positive result, and in bright red for a negative result. Although the choice of colors might pose a problem to users that suffer from some forms of colorblindness, it is sufficient for a research prototype. 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 11.

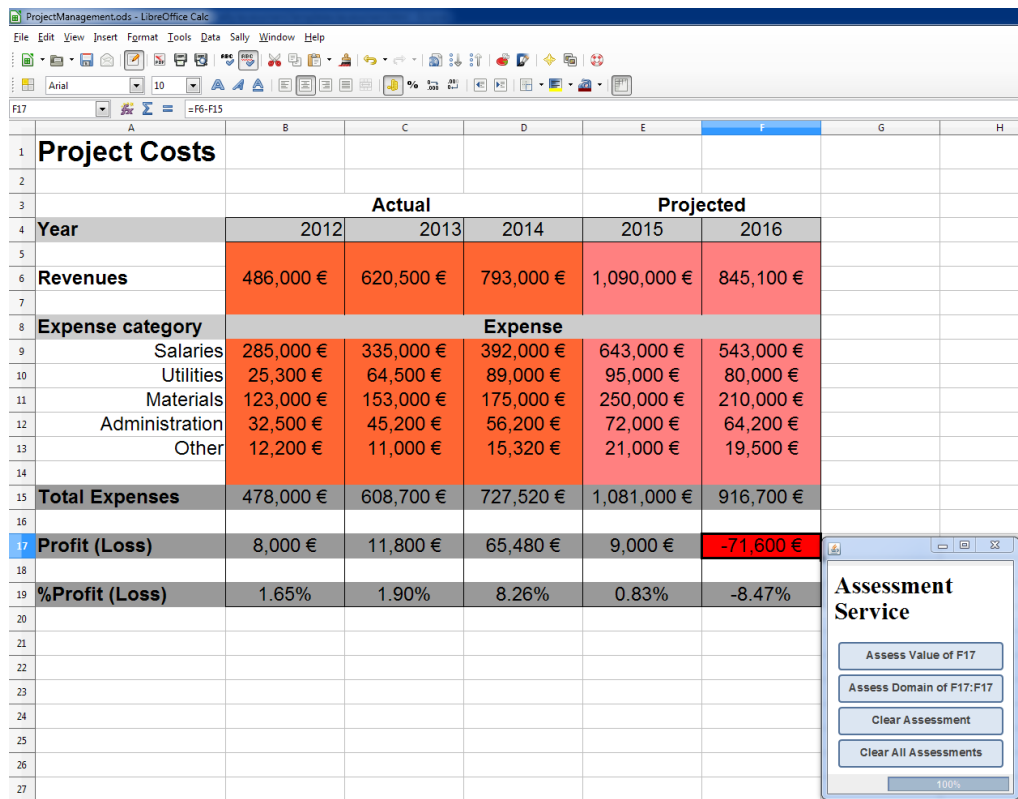


Figure 11: Assessed Projected Profit

The process of assessing an entire functional block(a domain) is almost identical. After the range is selected and the Assessment Service is 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 it. After this sanity check, the process continues as in the case of the single cell assessment.

After performing an assessment, the spreadsheet cannot be edited by the user. This locking is performed in order to offer the user the possibility of returning the spreadsheet document to a pre-assessment state. The option to clear the result of the last assessment or all the assessment results is offered to the user by the Assessment Service. After clearing all the results, the user can edit the document without any restrictions.

As we have seen, the main functionality of the service is concentrated in the MMT system, so we will now take a closer look at its role.

4.2 MMT Assessment Theories

The MMT system sits at the core of the Assessment Service. It allows to formalize the assessment process as a set of structured theories that can be viewed in the specific context of the spreadsheet.

Let us return to our running example to see how assessment knowledge is formalized and what an "assessment theory" represents. The first spreadsheet item that John wants to evaluate is the one of "Profits". Alin, our former domain expert, considered that "A profit value is good if it is positive.". This is a simple and, one might say, obvious evaluation, but it showcases the concepts behind assessment. These guidelines for evaluating an item are laid out by the domain expert which is, in most cases, also the creator of the document. Unfortunately, the system cannot compute if the values in the workbook respect this informal assessment rule. It requires a formal representation of this assessment that it can understand and process. The premise of the entire Assessment Service is that the domain expert can easily formalize the knowledge he uses to evaluate the spreadsheet as MMT theories. We call these formal representation of evaluation rules "assessment theories".

As we have seen in the previous subsection, the user can assess an individual value or he can assess the behavior of a function over a domain. I will first look into the process of formalizing the theory "A profit value is good if it is positive." and then I will give a brief overview of the theory graph for the assessment "A profit should increase over time.".

The theory graph in Figure 12 is the one used to assess if a profit value is positive or not. It is composed of a number of generic theories which are framed according to the concept we are evaluating. The following theories represent generic assessment components:

- **Comparison** contains two symbols: **base** and **comp**. **base** represents a point of reference with respect to which we would like to compare another value. **comp** represents a generic comparator that given two values, returns a boolean value expressing if the values respect its rule.
- **Valuation Function** includes the **Comparison** theory and contains three symbols: **domain, params** and **f**. **f** represents a function which we want to assess in a particular point, **params** represents the coordinates of the function value we are interested

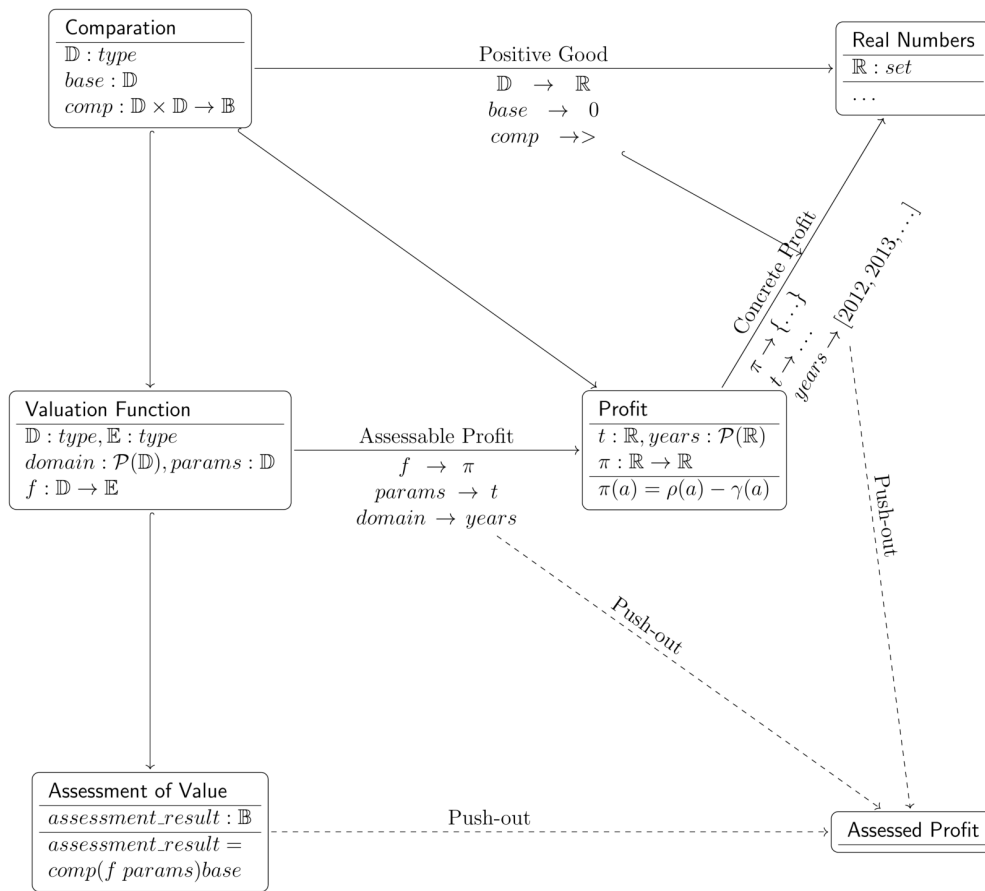


Figure 12: Assessment of Value Theory Graph

in and **domain** represents the domain over which we want to evaluate the function. For the purpose of assessing a value, the **domain** symbol can be ignored.

- **Assessment of Value** includes the **Valuation Function** theory and contains one symbol: **assessment_result**. **assessment_result** is the result of comparing the value of a function in a certain point with a base reference.
- **Real Numbers** represents the set of real numbers and it will allow us to view our theories from a more concrete perspective.

Profit represents the concept of profit. It contains three symbols π, t and **years** which represent the profit function, the value of the year in which we want to compute the profit and a set of years representing its domain. **Profit** can be substituted at runtime with any other concept for which we want to assess a value. This setup allows us to evaluate arbitrary functions and their values as long as we transfer all the necessary information from the spreadsheet to the MMT system.

Having the "theory framework" set into place, the next step is to particularize it for our context. We realize this by generating the following views:

- **Positive Good** is a view from **Comparison** to **Reals** and assigns to **base** the value 0 and to **comp** the greater than operator($>$). This will allow us to check if a value is

greater than 0.

- **Assessable Profit** is a view from **Valuation Function** to **Profit**. The **f** symbol is mapped to π , the **params** symbol is mapped to **t** and the **domain** symbol is mapped to **years**. This view makes the profit assessable.
- **Concrete Profit** is a view from **Profit** to **Real Numbers**. This is the point where the profit function and the year of interest are instantiated with concrete values. This is the theory that is closest to our representation of profits in the document.

At this point, we have integrated our concrete values into the theory graph and all we have to do is perform the assessment. Basically, we are interested in the value of **assessment_result** in the frame of **Assessable Profit** and **Concrete Profit**. The theory **Assessed Profit** is generated by MMT over the two views using the push-out mechanism. The Universal OpenMath Machine will simplify the value of the pushed-out **assessment_result** and the result will be returned to the Assessment Service component in Sally.

We now turn our attention to the process of assessing a function over a domain. Our informal assessment rule says "A profit should increase over time.". The formalization of this rule is represented as a theory graph in Figure 13. As you might have noticed, the theory graph has a similar structure to the one used for assessing an individual value so I will only highlight the most important differences.

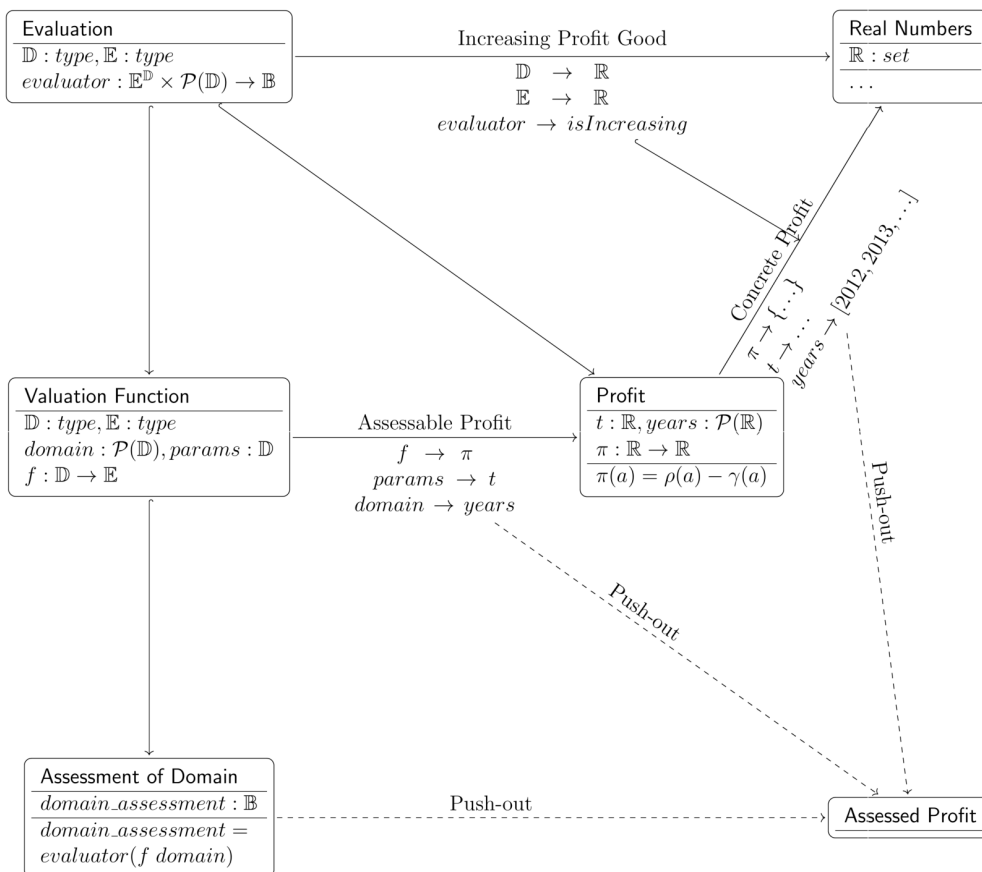


Figure 13: Assessment of Domain Theory Graph

The **Comparison** theory is replaced with the **Evaluation** theory. This theory contains the symbol **evaluator** which, given a function and a list of input values representing a sub-domain, will return a boolean value indicating whether the function has the expected behavior over the given sub-domain. The last theory of importance is **Assessment of Domain**. This theory contains the symbol **domain_assessment** which is the result of evaluating a function over a domain. As we can see from the theory graph in Figure 13, assessing a function over a domain is very similar to assessing a single value. In fact, the two theory graphs can be merged, but they have been kept separate for the benefit of simplicity.

The Assessment Service I have developed 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.

5 Use Case

In this section we take a closer look at our running example and see how the Assessment Service can be used in a real life scenario. We have two users of the same spreadsheet document, John and Robert, that have a slightly different view on the ideal outcome of the management process.

John's Assessments

John wants to maximize resource consumption in order to maximize growth and the division's potential market value, while maintaining a profit margin. As John is on his first day on the job, he has to rely on the data left behind by his predecessor. As we have already discussed, it would be extremely difficult for our young manager to get a full picture of the department's status just by analyzing the data contained in the spreadsheet document. In this fortunate case, the electronic ledger is accompanied by a background ontology and a set of assessment theories developed by its creator, John's predecessor.

The first step our manager takes is to analyze the current and expected profits. After selecting the "Actual Profit" functional block, the Assessment Service gives him two options, to assess the individual cells or to assess the functional block as a whole, the entire domain. He decides to first assess the individual cells. After choosing the option "Assess Value. . .", the entire "Actual Profit" block is colored in bright green, as we can see in Figure 14. This result is reassuring to John, but he still wonders why this is. Answering "Why the assessment is positive?" is left to a future service. After inspecting the assessment theory, John suspicion is confirmed: An actual profit is positively assessed if its value is positive. Now, John assesses the entire functional block "Actual Profit", and again, the whole block is colored in green, reflecting the fact that the profit has increased over time.

John turns his attention to the Projected Profit functional block and the assessment of the individual cells returns a negative result. Projected Profit should also be positive, but for some reason, the expected profit for the year 2016 is negative. This result forces John to

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%

Figure 14: Assessment Service

look deeper into the business plan of the division and determine if the decisions taken by his predecessor were sound.

Using the Assessment Service has enabled John to enhance his decision making process and to direct his attention into the right direction. He can now take a closer look at the critical points of the business plan and decide if a restructuring is necessary.

Robert's Assessments

Now, we look at assessing data in our document from Robert's point of view. Robert, John's boss, is in charge of managing multiple divisions and redistributing funds in order to maximize global efficiency and growth. Robert main source of information about John's division is the same spreadsheet document we have used until now.

Although both Robert and John want to make profit, they have slightly different views on how this should be achieved. For instance, Robert has a different idea on how a profit should grow, when it should be allowed to transform into a loss and how it should evolve over time. This different opinion materializes as a different association between cell ranges and assessment theories. This switch in perspectives can be easily implemented by allowing users to choose and edit the mapping between cells and functional blocks and assessment theories and by storing the mapping in a private extension of the ASM.

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 [40]. He selects the entire Actual Profit block and chooses Assess Functional Block. Unfortunately, the whole block is colored in red. This makes Robert worry that the results of this division might have a negative effect on the company's public appearance and stock value.

When assessing the Projected Profit block, Robert discovers another problem. Again, the projected profit for the year 2016 raises a red flag. He decides to further investigate the issue and talk with his new subordinate in order to create a viable future for this division. If all the assessments would have been positive i.e. the blocks would have been colored green, Robert would have been assured that the department John is managing is faring well and he could have focused his attention somewhere else.

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.

6 Future Work and Conclusions

Implementation Improvements

The assessment theories we have used so far have been fairly basic, but, in real life, a decision support system would have to support more complex assessments that factor in data from multiple sources, inside and outside the spreadsheet application.

MMT has the power to represent arbitrary languages based on theories, so it can be used to express powerful and complex assessment theories. At the moment, the main difficulty in creating high-level assessment theories is providing the framework for their evaluation. The current implementation of the Universal OpenMath Machine supports only a limited number of primitive operations on basic data types. In order to offer a framework that supports fast and efficient development of assessment theories, the Universal OpenMath Machine would have to be extended to support a larger number of primitive data types and operations.

An important challenge for the usability of the system at this point is associating assessment theories with the functional blocks and individual cells. The current implementation forces the user to create the Abstract Spreadsheet Model using its Java interface. This seriously impairs the usability of the framework in a production environment. A solution to this problem would be to extend the Abstract Spreadsheet Model Editor¹⁰ with functionality to associate assessment theories with individual cells and functional blocks. As an assessment theory is represented by a URI, this extension would require only a few developer hours but it would considerably improve the work flow of the creator of the spreadsheet document.

¹⁰A graphical user interface for editing the Abstract Spreadsheet Model

Future Services

In the presented use case example, one question was repeatedly raised during the assessment process: *Why is this assessment positive/negative?*. This question is multifaceted. 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

The first aspect of understanding the result of an assessment is understanding the assessment theory behind it. The current implementation stores a formal representation of the assessment theory in the MMT language which is also used in the evaluation process. As the MMT language is constantly evolving and has a highly formal character, it is unfeasible to ask of the average spreadsheet reader to fully understand it.

In essence, we need to combine the formal and informal aspects of the assessment theory and present them to the user in a simple, context sensitive interface. One approach would be to pair each assessment theory with an informal definition stored in the ontology. When the user assesses an object, he can also choose to see the informal definition of the assessment theory. Although this approach can be implemented with the current technology, it would lead to a system where it would be very expensive to develop and maintain assessment theories and definitions.

A better approach would merge the formal and informal representations of the assessment theory in a single entity. The work of MIHNEA IANCU on Flexiformal Mathematics [41] will allow documents with a variable degree of formality on which different semantic services can be used. By leveraging computations, adaptable explanations can be offered depending on the result of the assessment.

Assessment of Dependencies

In this section, the interest falls on determining the causes of an assessment result as we extend on the work done by KOHLHASE and KOHLHASE in *Spreadsheets with a Semantic Layer*.

Let us consider 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.

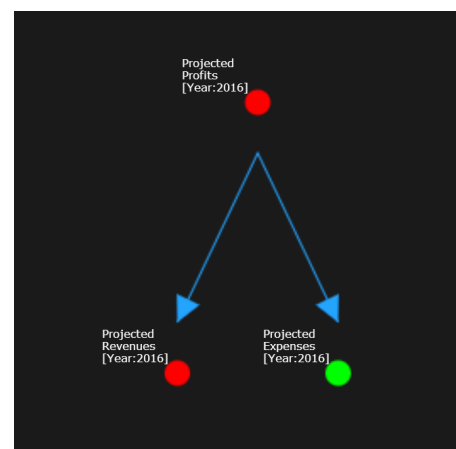


Figure 15: Assessment of Dependencies

In Figure 15 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 [42] provides all the necessary functionality to visualize the dependence graph and the Assessment Service already provides functionality for assessing the individual cells and functional blocks. At this point in time, the only functionality that is required to implement this service is a Formula Parser that would transform spreadsheet formulas into a tree structure. The development of the presented functionality is left for further research.

Conclusion

My bachelor thesis builds on the theoretical work conducted by KOHLHASE and KOHLHASE in *Spreadsheets with a Semantic Layer* in the area of assessment knowledge in spreadsheets.

The developed Assessment Service, based on the concept of theory graphs, that given only minimal user input in the form of associations of assessment theories and spreadsheet items, is able to deduce and perform value judgements on the concrete workbook data. Assessment theories can be easily created by domain experts with minimal effort while incrementally increasing the value of the system for end users.

This service acts as a decision support system and at the same time enhances the spreadsheet application's semantic transparency. I believe that, integrated in the Semantic Alliance Framework, the Assessment Service is an invaluable tool for spreadsheet users.

References

- [1] Visicalc. <http://en.wikipedia.org/wiki/VisiCalc>. [Accessed on 9-May-2014].
- [2] Hemant K. Bhargava, Suresh Sridhar, and Craig Herrick. Beyond spreadsheets: Tools for building decision support systems. *IEEE Computer*, 32(3):31–39, 1999.
- [3] Valentin Toader and Veronica Rozalia Rus. Spreadsheet-based decision support systems. <http://steconomiceuoradea.ro/anale/volume/2008/v4-management-marketing/275.pdf>, 2014. [Online; accessed 26-January-2014].
- [4] Michelle M.H.eref and Wayne L. Ahuja, Ravindra K. Winston. Spreadsheet-based decision support systems. In *Handbook on Decision Support Systems 1*, International Handbooks Information System, pages 277–298. Springer Berlin Heidelberg, 2008.
- [5] Polly S. Brown and John D. Gould. An experimental study of people creating spreadsheets. *ACM Trans. Inf. Syst.*, 5(3):258–272, July 1987.
- [6] Thompson SH Teo and Margaret 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*, HICSS '97, pages 149–, Washington, DC, USA, 1997. IEEE Computer Society.
- [7] Kamalasen Rajalingham, David R. Chadwick, and Brian Knight. Classification of spreadsheet errors. *CoRR*, abs/0805.4224, 2008.
- [8] Raymond R. Panko. Spreadsheet errors: What we know. what we think we can do. *CoRR*, abs/0802.3457, 2008.
- [9] Stephen G. Powell, Kenneth R. Baker, and Barry Lawson. A critical review of the literature on spreadsheet errors. *Decision Support Systems*, 46(1):128 – 138, 2008.
- [10] Andrea Kohlhase and Michael Kohlhase. Semantic transparency in user assistance systems. In Brad Mehlenbacher, Aristidis Protopsaltis, Ashley Williams, and Shaun Slattery, editors, *SIGDOC*, pages 89–96. ACM, 2009.
- [11] Components of Decision Support Systems. <http://dssystem.blogspot.de/2010/01/components-of-decision-support-systems.html>, 2014. [Online; accessed 31-March-2014].
- [12] Thomas R. Gruber. A translation approach to portable ontology specifications. *Knowl. Acquis.*, 5(2):199–220, June 1993.
- [13] Catalin David, Constantin Jucovschi, Andrea Kohlhase, and Michael Kohlhase. Semantic alliance: A framework for semantic allies. In Johan Jeuring, John A. Campbell, Jacques Carette, Gabriel Dos Reis, Petr Sojka, Makarius Wenzel, and Volker Sorge, editors, *AISC/MKM/Calulemus*, volume 7362 of *Lecture Notes in Computer Science*, pages 49–64. Springer, 2012.
- [14] Andrea Kohlhase and Michael Kohlhase. Spreadsheets with a semantic layer. *Electronic Communications of the EASST:Specification, Transformation, Navigation Special Issue dedicated to Bernd Krieg-Bruckner on the Occasion of his 60th Birthday*.

- [15] Andrea Kohlhase and Michael Kohlhase. Compensating the semantic bias of spreadsheets. In Joachim Baumeister and Martin Atzmler, editors, *LWA*, volume 448 of *Technical Report*, pages 49–58. Department of Computer Science, University of Würzburg, Germany, 2008.
- [16] Software engineering for spreadsheet interaction. <http://trac.kwarc.info/sissi>. [Accessed on 21-April-2014].
- [17] Wikipedia. Transparency (linguistic). [http://en.wikipedia.org/wiki/Transparency_\(linguistic\)](http://en.wikipedia.org/wiki/Transparency_(linguistic)), April 2014. [Accessed on 19-April-2014].
- [18] Andy Clark. *Mindware: An Introduction to the Philosophy of Cognitive Science*. New York: Oxford University Press, 2001.
- [19] Andrea E. Kohlhase and Michael Kohlhase. Semantic transparency in user assistance systems. In *Proceedings of the 27th ACM International Conference on Design of Communication*, SIGDOC '09, pages 89–96, New York, NY, USA, 2009. ACM.
- [20] Michael Kohlhase. OMDoc – *An open markup format for mathematical documents [Version 1.2]*. Number 4180 in LNAI. Springer Verlag, 2006.
- [21] Florian Rabe and Michael Kohlhase. A scalable module system. *CoRR*, abs/1105.0548, 2011.
- [22] Deborah L. McGuinness and Frank van Harmelen. Owl web ontology language overview. w3c recommendation. <http://www.w3.org/TR/2004/REC-owl-features-20040210/>, February 2004. [Accessed on 21-April-2014].
- [23] KWARC. Localmh - a cross-repository administration tool for the local authoring in mathhub.info. <http://gl.mathhub.info/MathHub/localmh>. [Accessed on 21-April-2014].
- [24] Henrik Eriksson. The semantic-document approach to combining documents and ontologies. *International Journal of Human-Computer Studies*, 65(7):624 – 639, 2007. Knowledge representation with ontologies: Present challenges - Future possibilities.
- [25] Representational state transfer. http://en.wikipedia.org/wiki/Representational_state_transfer. [Accessed on 21-April-2014].
- [26] Category theory. http://en.wikipedia.org/wiki/Category_theory. [Accessed on 9-May-2014].
- [27] Category theory/categories. http://en.wikibooks.org/wiki/Category_Theory/Categories. [Accessed on 9-May-2014].
- [28] Joonwon Choi. *Pushouts, Pullbacks and Their Properties*. Bachelor's thesis, Mathematics, Seoul National University, Seoul, KR, May 2013.
- [29] Florian Rabe. How to identify, translate, and combine logics?
- [30] Felix Gabriel Manche. *A Practical OpenMath Machine*. Bachelor's thesis, Computer Science, Jacobs University Bremen, Bremen, DE, May 2013.
- [31] Vladimir Zamdzhiev. *Universal OpenMath Machine*. Bachelor's thesis, Computer Science, Jacobs University Bremen, Bremen, DE, May 2011.

- [32] Scala Programming Language. <http://www.scala-lang.org/>, 2014. [Online; accessed 26-January-2014].
- [33] OpenMath. <http://www.openmath.org/>, 2014. [Online; accessed 26-January-2014].
- [34] OMDoc. <https://trac.omdoc.org/OMDoc>, 2014. [Online; accessed 26-January-2014].
- [35] Proof assistant. http://en.wikipedia.org/wiki/Proof_assistant. [Accessed on 21-April-2014].
- [36] Computer algebra system. http://en.wikipedia.org/wiki/Computer_algebra_system. [Accessed on 21-April-2014].
- [37] Libreoffice. <http://www.libreoffice.org/>. [Accessed on 21-April-2014].
- [38] Java programming language. [http://en.wikipedia.org/wiki/Java_\(programming_language\)](http://en.wikipedia.org/wiki/Java_(programming_language)). [Accessed on 21-April-2014].
- [39] Post(http). [http://en.wikipedia.org/wiki/POST_\(HTTP\)](http://en.wikipedia.org/wiki/POST_(HTTP)). [Accessed on 21-April-2014].
- [40] Dow jones industrial average. http://en.wikipedia.org/wiki/Dow_Jones_Industrial_Average. [Accessed on 21-April-2014].
- [41] Mihnea Iancu. *Towards Flexiformal Mathematics* PhD Research Proposal, Computer Science, Jacobs University, Bremen, DE, February 2013.
- [42] Javascript infovis toolkit. <http://philogb.github.io/jit/>. [Accessed on 21-April-2014].