

Verifiable Source Code Documentation in Controlled Natural Language

Tobias Kuhn¹ and Alexandre Bergel²

¹ Chair of Sociology, in particular of Modeling and Simulation, ETH Zurich
<http://www.tkuhn.ch> kuhntobias@gmail.com

² Pleiad Lab, Department of Computer Science (DCC), University of Chile
<http://bergel.eu>

Abstract

Writing documentation about software internals is rarely considered a rewarding activity. It is highly time-consuming and the resulting documentation is fragile when the software is continuously evolving in a multi-developer setting. Unfortunately, traditional programming environments poorly support the writing and maintenance of documentation. Consequences are severe as the lack of documentation on software structure negatively impacts the overall quality of the software product. We show that using a controlled natural language with a reasoner and a query engine is a viable technique for verifying the consistency and accuracy of documentation and source code. Using ACE, a state-of-the-art controlled natural language, we present positive results on the comprehensibility and the general feasibility of creating and verifying documentation. As a case study, we used automatic documentation verification to identify and fix severe flaws in the architecture of a non-trivial piece of software. Moreover, a user experiment shows that our language is faster and easier to learn and understand than other formal languages for software documentation.

1. Introduction

Software documentation is commonly understood as any description of software attributes, and documenting software is a critical factor to success [1]. Glancing at the website of the ten most popular SourceForge software programs¹ easily demonstrates that: all software comes with a fair amount of documentation, including screenshots, tutorials, and FAQs.

However, scrutinizing the documentation of these popular programs reveals that the offered documentation, with extensive installation procedures and prob-

¹VLC, eMule, Vuze, Ares Galaxy, Smart package, 7-Zip, Notepad++ Plugin Manager, PortableApps.com, FileZilla, and MinGW are the most downloaded software at the time this article is being written. See <http://sourceforge.net/>

lem troubleshooting instructions, is essentially focused on end users. None of the projects we studied present a clear description of their internal structure. The source code of eMule, for example, weighs more than 8 MB with 260 000 lines of code. The extensive website of eMule contains many documentation files², but none of them are targeted toward developers.

This situation is not an exception, especially when the documenting of the software internals is not a strong requirement [2], as is often the case with open source software. There are good reasons for this, including the lack of resources, the fragility of the documentation [3, 4], and the lack of immediate benefits. Keeping the internal documentation up to date with the source code is particularly difficult [5]. This article proposes an effective and convenient solution to write documentation that is interpreted as a formal model, is automatically checked, and is kept in sync with the source code.

We propose the use of controlled natural language [6, 7] to represent software structure and its relation to the outside world, which leads to documentation that is automatically verifiable. The idea is to let developers write statements in a subset of English, describing the design and the environment of software that is otherwise nowhere explicitly expressed in the source code (e.g., *Simon is a developer*, *Shapes is a component*, *Every subclass of MOShape belongs to Shapes*, and *Simon maintains every class that belongs to Shapes*).

Among the different kinds of controlled natural languages [7], our approach relies on the type of languages that provide a natural and intuitive representation for formal notations. The statements of such languages look like natural language (such as English), but are controlled on the lexical, syntactic, and semantic levels to enable automatic and unambiguous translation into formal logic notations. We propose to use such natural-looking statements to write documentation statements about the source code structure. Ideally, these documentation statements should be next to the source code, e.g. in code comments, and written, read, and evaluated in the IDE of the developer. In this way, changes in the source code or its documentation that would introduce an inconsistency can be automatically detected. Furthermore, the logical structure of the documentation is accessible to query engines, via questions written in the same controlled natural language.

We evaluate our approach using the language Attempto Controlled English (ACE) [8, 9] and the semantic wiki engine AceWiki [10]. To test the feasibility of our approach, we apply it to document Mondrian, an open-source visualization engine, and to improve its design. In addition, we conduct a user experiment to test how controlled natural language compares to other formal languages for source code documentation and querying.

In this article, we investigate the following two research questions:

- A – *Can controlled English be efficiently used to document software systems?*
- B – *Is controlled English easier to understand than other formal languages for*

²<http://www.emule-project.net/home/perl/help.cgi?l=1>

software documentation?

A growing amount of empirical evidence [11] shows that programmers spend a significant amount of time understanding source code and navigate between source code of software components. The generality of available software engineering tools and their poor performance are often pointed out as the culprit for engineering tasks being carried out in a suboptimal fashion [12, 13]. Informal, broken, or poorly conceived documentation seem to be at the root of many difficulties encountered in software engineering tasks [11, 14, 15]. The two research questions we have stated above explore a new direction: documentation should be written in restricted natural language and automatically verified against the actual source code.

We tackle the first research question by documenting a significant amount of source code in ACE using AceWiki. The second research question is investigated by conducting a user experiment with 20 participants.

We would like to stress that this article is *not* about providing yet another solution for querying source code. Effective solutions have been proposed already, using logic programming [16] or natural language [17]. Reverse-engineering and reengineering environments typically offer sophisticated way to query and extract information from source code. Our approach combines source code modeling techniques with general logic expressions for documentation, all accessed via a control natural language and accessible to formal reasoning. ACE has been successfully applied in areas such as the Semantic Web, however its application to software documentation is new and has not been considered as far as we are aware of.

We also would like to point out that our approach focuses on *internal* code documentation. Software documentation can describe a wide variety of aspects, including system requirements, architecture, design decisions, and implementation details. In this article, we focus on the low level source code structure where documentation statements talk about source code elements such as classes and methods, but express things that are not visible in the source code. This includes connections to the outside world, such as who is responsible for which classes and what external devices a particular piece of code depends on.

In order to answer the research questions introduced above, this article makes the following contributions:

- i – it presents a concrete approach to document source code with a controlled natural language;
- ii – it shows how source code modeling can be combined with controlled natural language and automated reasoning;
- iii – it evaluates the approach on a case study of documenting a medium-sized application;
- iv – it presents the results of a user experiment for the comparison of the understandability of different languages for source code documentation.

These contributions consist of applying a state-of-the-art approach on natural language processing to address the recurrent software engineering problem of writing understandable and verifiable software documentation.

This article is organized as follows: first, we motivate our approach and provide some background information (Section 2), before introducing the specifics of our approach (Section 3); then, we describe our implementation that is based on AceWiki (Section 4); subsequently, the case study on Mondrian is detailed (Section 5), the design of the user experiment and its results are explained (Section 6), a number of points are discussed (Section 7), before our work is put in a broader perspective (Section 8) and we draw some conclusions from this work (Section 9).

2. Motivation and Background

Even though the benefits of code documentation are well known [18, 19], many popular applications have poor or outdated documentation. Donald Knuth has for long advocated a radical change toward the way we build software programs [20]:

Let us change our traditional attitude to the construction of programs: Instead of imagining that our main task is to instruct a computer what to do, let us concentrate rather on explaining to human beings what we want a computer to do.

Knuth argues that a software program should be conceived as a work of literature. Unfortunately, software engineers are far from being that disciplined. Several studies show that internal documentation is considered important but is rarely written [19, 21].

The need of better documentation is often perceived as one of the principal reasons for difficulties encountered in software development [11, 15, 22, 23]. Empirical studies [11] show that an average programmer spends a significant amount of time “finding initial points in the code that were relevant to the task”, finding “more information relevant to the task”, and understanding graphs of interacting code.

Recent work in software documentation [14, 15] has focused on relating source code to informal descriptions (e.g., API descriptions or StackOverflow posts). In this way, obsolete or erroneous pieces of documentation can be identified, and social networks can be analyzed to look for discussions associated to a particular part of the source code. Current approaches based on informal documentation mostly use simple text pattern matching, e.g., identifying typos in class names [15] and dividing camel-cased terms into compound terms [24]. In contrast, our two research questions we stated earlier (Section 1) focus on the use of a controlled natural language for software documentation, which allows us to carry out sophisticated analyses (in particular for queries and checks) that are not possible in a reliable way with uncontrolled natural language.

We identified three main challenges to overcome in order for a documentation system based on a control natural language to be practical, effective and reliable:

Validation. Documentation is only useful if it correctly reflects the implementation, otherwise it can be counterproductive. Documentation therefore needs

to be regularly validated.

Immediate feedback. The documentation of source code is usually fragile with respect to code changes. A small change can invalidate a considerable part of the documentation. Only with immediate feedback, divergences between documentation and source code can be quickly identified, which is essential to prevent desynchronization.

Minimal impact on software engineers. To motivate software developers to write sufficient and good documentation, the effort for writing documentation should be minimal and should not disrupt their work flow.

Numerous approaches have been proposed to address these challenges. Their impact, however, is rather limited so far.

Interface Definition Languages, for instance, are specification languages to describe software components (e.g., WebIDL³). Their application range is commonly limited to the very technical corner, such as remote procedure invocations. A description is meant to be read by a compiler, not by a human being. IDLs can be easily verified and can produce immediate feedback, but the extra work for software engineers is considerable. To document software in IDL, one has to first learn the IDL syntax, become familiar with the IDL tools, and understand how to detect and interpret a mismatch between documentation and source code.

Modern programming environments offer a large set of tools to navigate and query source code. Such queries are usually expressed in a proprietary domain specific language⁴ or in a kind of SQL-like language.⁵ Punctual queries on the source code can be used to verify structural properties, but they cannot easily be exported and stored for documentation purposes.

A fair amount of previous work exists on how to reduce the documentation effort that has to be provided by software engineers. For example, Sridhara *et al.* [25] generate comments for method arguments, and Buse *et al.* [26] automatically document exceptional events in natural language. These approaches only cover a rather narrow range of possible documentation statements. One can argue that the crucial role of documentation is to point out things that cannot be seen in the source code, and the above approaches are no help in this respect.

To deal with the challenges outlined above, we propose to use controlled natural language. Controlled natural languages are restricted versions of natural languages (e.g., English) that are easily processable by computers. Grammar and dictionaries are restricted in such a way that ambiguities are reduced or (as in our case) completely eliminated. Recently, controlled natural languages have started to be employed to query software source code [17].

³<http://www.w3.org/TR/WebIDL>

⁴e.g., <http://browsebyquery.sourceforge.net>

⁵e.g., <http://www.eclipse.org/modeling/emf/?project=query>

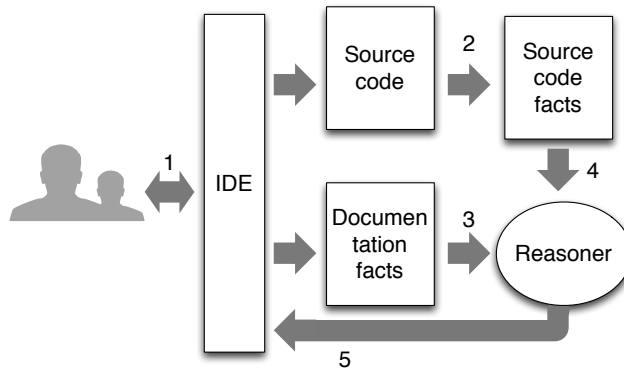


Figure 1: Workflow of our approach

In the subsequent sections, we present our approach on applying controlled natural languages to software documentation. With the help of a source code model, a controlled natural language, a proper editing environment, and a reasoning engine, software engineers can write documentation that can be validated, leads to immediate feedback, and imposes only minimal learning efforts on the software engineers.

3. General Approach

In what follows, we first present the general workflow of our approach (Section 3.1). Then, we describe how software source code is represented (Section 3.2), before we explain the main operations available to developers who have to document software while ensuring the produced documentation is synchronized with the source code (Sections 3.3 – 3.5).

3.1. Workflow

The different steps of our approach are shown in Figure 1:

- *Step 1*: Developers write documentation statements in controlled natural language while writing the source code, ideally using the same integrated development environment. Such documentation statements can, for example, be written in specifically tagged source code comments, which makes them easily extractable.
- *Step 2*: Facts about the code structure are extracted from the source code. This is based on a source code model that is specific to the given programming language. Section Section 3.2 explains the details of such source code models.

- *Steps 3 and 4*: The manually written documentation statements and the automatically generated source code statements are given to a logic reasoner.
- *Step 5*: The reasoner performs inference tasks like consistency checking and query answering, and present the results to the developer via the IDE. This is explained in the sections [3.4](#) and [3.5](#).

3.2. Source Code Model

In order to link the documentation to the source code, a source code model has to be defined. In our approach, this model is automatically extracted from the source code, and can later be used as a basis to manually write more general documentation statements. The generation of the source code model is just a prerequisite for the second part of our approach, but is not particularly interesting or novel on its own.

Examples of core concepts of such a source code model could be *code element*, *class*, *method*, *package*, or *interface*. Core relations could be *direct subclass of*, *defines*, *invokes*, or *instantiates*. We will use the language ACE for the examples to be shown below, but any other controlled natural language of sufficient expressiveness could be used. In any case, these concepts and relations should be defined together with morphological variants like *define* and *defined*, in order to be able to produce well-formed sentences. Restrictions on these core concepts and relations could be defined as follows:

Every class is a code element.
Every method is a code element.
No class is a method.
If X defines Y then X is a class.
If X defines Y then Y is a method.

When a word like *class* is just introduced, the reasoner does not know anything about the concept of a *class*. By adding statements like the ones above, the word *class* is formally linked to other words (by means of first-order logic), which captures (part of) the meaning of the word and enables the automatic calculation of inferences. The first example states that *class* is a subconcept of *code element*; the third one states that it does not overlap with the concept *method*; and the fourth connects it to the relation *defines*, saying that only classes can define something. In this way, we tell the reasoner about the domain of object-oriented source code (“we” means us as system developers, not end users).

Based on this model, we can automatically extract source code statements, for example:

EventHandler is a class.
EventHandler is a direct subclass of Handler.
EmergencyHandler is a class.
EmergencyHandler is a direct subclass of EventHandler.
EmergencyHandler-isActive is a method.
EmergencyHandler defines EmergencyHandler-isActive.

Again, it is important to note that users are not supposed to write such statements, but they are automatically extracted and ideally synchronized in real time as the source code changes. As we will see, these statements together with the source code model serve as a scaffold for user-defined documentation statements.

In addition to the statements shown above, we can extract statements about the invocation that occurs between methods:

EmergencyHandler-trigger invokes AlarmDesk-setAlarm.

AlarmDesk-setAlarm invokes Broadcaster-sendMessage.

Based on these core concepts and relations, we can define additional relations for convenience reasons, like *subclass of*, *direct superclass of*, *superclass of*, *method of*, *class of*, and *uses*:

If X defines Y then Y is a method of X.

If X is defined by something that belongs to Y then X is a method of Y.

If X invokes something that is defined by Y then X uses Y.

If X defines something that uses Y then X uses Y.

Altogether, this shows how the general structure of source code can be formally represented in controlled natural language, which is the basis for expressing actual documentation statements.

3.3. Verifying Documentation Written in ACE

We now come to the important and interesting part of our approach: user-defined documentation. The described source code model is the starting point, which can be extended by additional ontological entities (i.e. introducing new words) and definitions (i.e. writing new statements). Below, we give explain how individuals, concepts and relations can be defined and used for source code documentation in ACE.

Individuals correspond to logical constants and are the simplest form of ontological entities. They are typically represented in (controlled) natural language as proper names, for example *EventManager*, *Simon*, the *RMoD* group, or the *EventManager Tutorial*:

EventManager is a system.

Simon is a developer.

RMoD is a group.

The EventManager Tutorial is a user manual.

Classes and methods, as they are extracted from the source code, are also individuals and should therefore be represented as proper names.

Nouns are typically used to represent *concepts*, which correspond to unary predicates in logic. *Class*, *document* and *developer* are examples of such concepts. In order to have a properly defined knowledge base, the respective concepts should be organized in a hierarchy, called taxonomy. A clean taxonomy helps users and reasoners to interpret and use the given concepts in a coherent way. Taxonomies consist of subconcept relationships that can be expressed

by sentences of the form *every A is a B*. This is a simple example of such a taxonomy:

Every person is an entity.
Every developer is a person.
Every tester is a person.
Every group is an entity.
Every artifact is an entity.
Every document is an artifact.
Every user manual is a document.
Every code element is an artifact.

Indentation is used here to clarify the hierarchical structure. Typically, some — but not all — of the concepts on the same level are supposed to be disjoint. This can be defined with statements such as the following:

No person is a group.
No person is an artifact.
No group is an artifact.

Of course, this is not the only sensible way to define a taxonomy for source code documentation, and our approach does not depend on a particular one.

There are many ways in which *relations* between individuals, called binary predicates in logic, can be represented in (controlled) natural language, including transitive verbs (e.g. *maintain*, *write*), *of*-constructs (e.g. *member of*), and adjectives with prepositions (e.g. *related to*). These are some examples of how such relations can be used:

If X maintains something then X is a person.
If something is a member of X then X is a group.
Simon Denier is a member of RMoD.
Simon Denier maintains EmergencyHandler.
AlarmDesk is written by Simon Denier.
The EventManager Tutorial describes EventManager.

The general nature of logic allows users to write formal documentation about everything they consider relevant, from modules and bugs over developers and groups to tasks and business plans.

3.4. Queries

Putting together these different kinds of statements and feeding them all to a logic reasoner, we can automatically answer different kinds of sophisticated questions. Such questions can be written in the same controlled natural language as the other statements. As a start, we can query the source code structure and obtain the answers from the reasoner:

Which methods are invoked by more than 80 methods?
- Broadcaster-sendMessage
- EventMessage-getText
- ...

Which classes instantiate a subclass of Event?

- *AlarmDesk*
- *TemperatureSensor*
- ...

While there exist integrated development environments and other code analysis tools that are able to answer such questions, controlled natural language allows us to formulate them in a more intuitive way. More importantly, users can query the documentation in the same way they can query the code structure, as the following examples show:

Which members of RMoD maintain more than 20 classes?

- *Simon Denier*
- ...

Which classes are related to a class that is maintained by Simon Denier?

- *AlarmDesk*
- ...

Which document that describes EventManager is written by a member of RMoD?

- *the EventManager Tutorial*
- ...

More examples of such questions follow in the subsequent sections.

3.5. Consistency

Each new statement is checked by the reasoner — immediately after its creation — for consistency. If a certain statement is not consistent with the existing knowledge base, the user is informed and the statement is not included in the knowledge base. In this way, we can be sure that the knowledge base is always consistent, which is a prerequisite for any other reasoning task.

For example, let us consider a situation in which — among others — the following statements are present in our knowledge base:

- If X is a direct subclass of Y then X is a subclass of Y.*
- If X is a direct subclass of something that is a subclass of Y then X is a subclass of Y.*
- EventHandler is a direct subclass of Handler.*
- EmergencyHandler is a direct subclass of EventHandler.*
- No subclass of Handler is maintained by a member of Group-A.*
- Every member of Group-B is a member of Group-A.*
- Brian is a member of Group-B.*

The first two sentences are part of the source code model. The next two are automatically generated from the source code. The last three sentences are part of the user-defined documentation. In this situation, if a user intends to add the statement

EmergencyHandler is maintained by Brian.

then the reasoner is able to detect that it would introduce inconsistency and can inform the user that this statement cannot be added to the knowledge base. Deliberately, a non-trivial example has been chosen here that is well within the capabilities of state-of-the-art reasoners.

Source code changes are treated the same way: For every change in the source code or its documentation, it is checked whether this change would make the knowledge base inconsistent. If so, the user is informed and has to resolve the issue (by either modifying the source code or the documentation).

4. Implementation

We implemented our approach based on AceWiki [10, 27], an existing semantic wiki engine, in which developers can comfortably browse and edit software documentation written in the ACE language. The writing and improving of the documentation can be done in a collaborative way. AceWiki seems to be well-suited for the evaluation of the general feasibility and usefulness of our approach, but for the future we plan to implement it as an IDE plugin.

The goal of AceWiki is to provide a very intuitive and natural ontology interface, where the complex logic formalisms are hidden behind controlled English. In this way, AceWiki does not only provide a more natural interface, but also supports a higher degree of expressiveness than most existing semantic wikis. The main difference from other semantic wikis — like Semantic MediaWiki [28] or KiWi [29] — is the use of controlled English. Articles in AceWiki are written in ACE, a formal language that looks like natural English, but the wiki content is translated into logic in the background, which allows for automatic reasoning. Recently, AceWiki has been extended to support multilinguality [30].

Figure 2 shows a screenshot of the AceWiki interface. It resembles existing wiki interfaces like the one of Wikipedia. The essential difference is that the content of the wiki articles is not free text, but has to comply with the restrictions of the ACE language.

In the background, the wiki articles are translated into OWL (Web Ontology Language) [31] and given to a reasoner. The reasoner is used to ensure the consistency of the knowledge base and to answer questions. For the case study to be presented, the OWL reasoner FaCT++ [32] has been used, but others can be used within AceWiki.

To assist users writing sentences in ACE, AceWiki provides a special editor that is able to predict the possible next words or phrases for partial sentences. In this way, users can create well-formed sentences step by step without the need to learn the restrictions of the language in advance. Figure 3 shows a screenshot of this editor. Apart from function words like *every* or *not* that are predefined and cannot be changed, users can add and edit lexical entries for proper names, nouns, *of*-constructs, transitive verbs, and adjectives with prepositions.



Figure 2: A screenshot of AceWiki

5. Case Study

As a case study to evaluate our research question *A*, AceWiki has been employed to document Mondrian,⁶ an agile visualization engine written in Pharo,⁷ a Smalltalk-inspired programming language. This section introduces Mondrian, describes our documentation effort on it, and shows how this fits into its overall development process.

When Mondrian was released in 2006, its design was clear and neat [33]. The growing number of users, however, had triggered an explosion of new requirements that were not originally planned. Satisfying these user requests had been given priority, leaving the documentation behind. After just a few years, the discrepancy between documentation and implementation was considerable. Mondrian’s architecture had suffered from the multitude of fixes and feature additions.

Source Code Model. We first had to define a source code model for the Smalltalk-style programming constructs of Pharo, which was very straightforward. Based on this model, we extracted from the source code the following four types of statements:

MORectangleShape is a class.
MORectangleShape is a direct subclass of MOShape.

⁶<http://www.moosetechnology.org/tools/mondrian>

⁷<http://www.pharo-project.org>

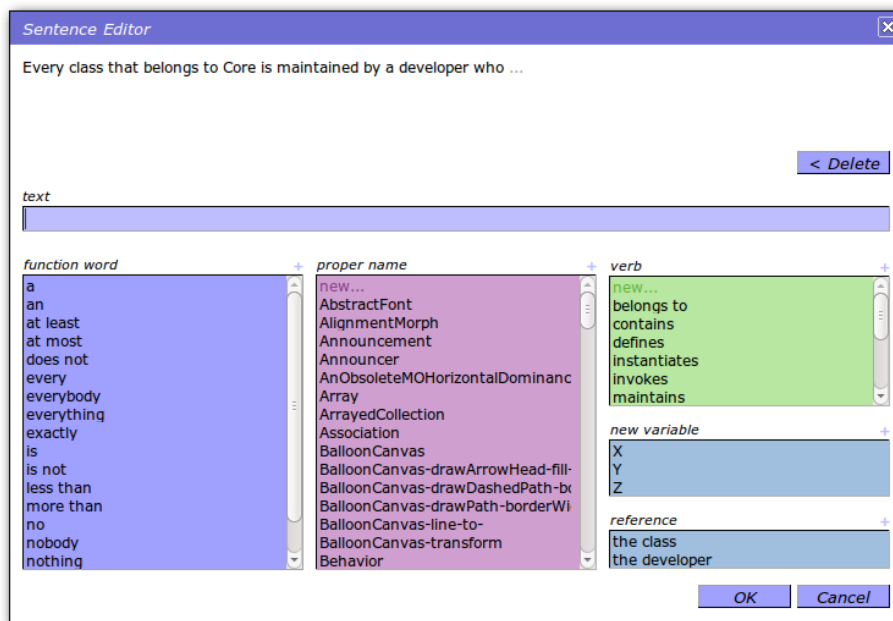


Figure 3: The editor of AceWiki

MOViewRendererTest-testCachedShape invokes MOShape-isCached.
MOViewRendererTest-testCachedShape instantiates MORectangleShape.

Next, we developed an AceWiki exporter for Pharo and the Moose software analysis platform, which we made available under the MIT license.⁸ The source code structure is extracted in a simple exchange format and then loaded into AceWiki.⁹ In total, 23 341 ACE statements were extracted from the Mondrian source code, using the FAMIX metamodel.¹⁰

Documenting Mondrian. Mondrian is a visualization engine that offers a rich domain specific language to define graph-based rendering. Each element of a graph (i.e., a node or an edge) has a shape that defines its visual aspects. Nodes may be ordered using a layout. An interactive scripting engine is offered that is called Easel. The interface with the graphical engine of the host platform is called Morph.

Mondrian is an excellent candidate for a documentation effort: We have access to its domain knowledge and its source code; it is a piece of legacy code (older than 6 years); it went through several evolution steps; and it is a

⁸<http://www.squeaksource.com/AceWikiExporter.html>

⁹<http://bergel.eu/download/mondrian.acewikidata>

¹⁰<http://www.moosetechnology.org/docs/famix>




Figure 4: Hypothetical architecture of Mondrian

critical part of the Moose platform. Mondrian comes with a reasonable set of documents, including book chapters,¹¹ tutorials,¹² research papers [33, 34] and mailing list archives. However, these documents are essentially for end users and do not describe the software’s internals.

Figure 4 shows the general architecture of Mondrian, as declared by its developers. There are nine components with `Core` as the central one.

Defining Components. In order to represent the Mondrian architecture in terms of software components, we first need to define some new words in Ace-Wiki. The terms *component* and *belongs to* were defined to assign classes to components:

Everything belongs to at most 1 component.
If X belongs to Y and uses Z then Y uses Z.

Basically, what we do here is to manually extend the model that we introduced above, which contains the concepts *class* and *method* but not *component*. This flexibility to add new concepts and relations allows for expressing things that were not anticipated when the system was designed. This is exactly what source code documentation should mainly do: explain things that cannot be expressed in the given programming language, or only in an implicit way.

We can subsequently define the components that compose Mondrian. First, the individuals need to be defined:

Core is a component.
Shapes is a component.
Tests is a component.
 ...

Then, each class can be assigned to a component, using the *belongs to* relation introduced above:

Every subclass of `MOGraphElement` belongs to `Core`.

¹¹<http://www.themoosebook.org/>

¹²<http://www.moosetechnology.org/tools/mondrian>

MOGraphElement belongs to Core.
MOShape belongs to Core.

In this way, class hierarchies and individual classes are assigned to the respective components.

The definition of some relations like *uses* can now be extended to take the new relation *belongs to* into account (and similarly for *class of* and *method of*):

If X uses something that belongs to Y then X uses Y.

We can now ask different kinds of questions such as

What belongs to Core?
Which component uses Core and uses Examples?

and state architectural properties like:

No component is used by Core.

This last sentence, however, triggered an inconsistency in our case. According to Figure 4, this sentence is supposed to be true, and the second question should give *Easel* as the only answer. However, as it is often the case when the documentation is not synchronized with the application source code, the exactitude of what is written on a white board is no more than a mere illusion. Asking the question

Which component is used by Core?

gives us the answers *Easel*, *Utils*, *Layouts*, *Shapes* and *Events*, which explains why the sentence above triggered an inconsistency. Some dependencies apparently violate the modular design: The *Core* component should not depend on other components.

One of the reasons why *Core* depends on *Layout* can be found in the `applyLayout` method defined in *MORoot*. *MORoot* is a subclass of *MOGraphElement*, therefore belonging to *Core*. This method is used to order the elements of an object root (it behaves as a composite). *Layouts* can be applied to an instance of any subclass of *MOAbstractLayout*. The source code of `applyLayout` is:

```
MORoot>> applyLayout
"Layout the receiver and all its children."
self resetMetricCachesResursively.
self do: [ :each | each applyLayout ].
self layout applyOn: self.
self nodesDo: [:each |
    self bounds corner:
        (self bounds corner
         max: (each extent + each origin)) ]
```

The reason for the odd dependency between *Core* and *Layouts* stems from the method call indicated in **bold** in the code shown above. The method `applyLayout` sends the message `applyOn:` to the object returned by `self layout`, an accessor to the *layout* instance variable defined in *MOGraphElement*. The method `applyOn:` is

defined in the class `MOAbstractLayout`. An instance of `MORoot` therefore directly and explicitly invokes a method defined in the `Layouts` component. Whereas having such a method call at runtime is not a problem per se, but explicitly having this call in the source code goes against the stated modularization structure.

This problem was subsequently fixed and included in version 543 of Mondrian. The class `MOAbstractLayout` had been moved to the `Core` component and renamed as `MOLayout`, to follow unique naming patterns. In this new version of Mondrian, `Layouts` depends on `Core`, but not vice versa.

As another example, the method `announce:` of `MOGraphElement` is responsible for the dependency of `Core` on `Events`. This method is used to emit an event (called “announcement” in the Pharo terminology), which is dispatched by an announcer. The definition of this method is as follows:

```
MOGraphElement>> announce: anAnnouncement
"public method"
announcer ifNil: [ ^self ].
anAnnouncement element: self.
announcer announce: anAnnouncement
```

The message `element: self` is sent to the object referenced by the variable `anAnnouncement`. The type of `anAnnouncement` is not explicit, but the method `element:` is defined in the class `MOEvent`, which belongs to the component `Events`. In addition, the method `announce:` is defined in the class `MOAnnouncer`, which belongs to `Events` too. The dependency of `Core` on `Events` stems from these two method invocations.

With our approach, improper component dependencies are easily identified by formulating adequate questions:

*Which class of Events is used by a class of Core?
Which method of Core uses a class of Events?*

The classes `MOAnnouncer` and `MOEvent` are among the results of the first question. `MOGraphElement-announce-` is in the answer of the second question. This problem was fixed by moving the method `announce:` to the definition of the `Events` component. This is consistent with the philosophy of Pharo and Smalltalk to assign methods to different modules rather than their associated classes. This mechanism is commonly designed as “class extension” [35].

About ten anomalies were identified in Mondrian’s architecture in the form of improper component dependencies. Versions 511 to 544 of Mondrian fixed all of them, making its architecture significantly cleaner.¹³

Code and documentation synchronization. With our approach, developers are notified when some parts of the documentation are invalidated by a certain change in the source code. This is illustrated by the documentation effort on Mondrian that began with version 511. The definition of components was the

¹³All the different versions and their associated comments are available online on <http://www.squeaksource.com/Mondrian.html>

first action we realized. The `Shapes` component, for example, was simply defined as all subclasses of `MOShape`:

Every subclass of `MOShape` belongs to `Shapes`.

At that time, the two questions

*Which method of `Shapes` uses `Events`?
Which method of `Shapes` uses `Layouts`?*

returned an empty answer, which is what one would expect according to the initial design of Mondrian (Figure 4). With the following two statements, we had forbidden unwanted dependencies:

*No method of `Shapes` uses `Layouts`.
No method of `Shapes` uses `Events`.*

However, changes were made over time, and version 525 introduced the method `display:on:` to the class `MOChildrenShape`:

```
MOChildrenShape>> display: anElement on: aCanvas
"The element to display can only have
 a formshape as a shape"

| builder bounds |
...
builder := anElement shape builder.
bounds := builder boundsOf: self.
...
```

This method belongs to `Shapes` but calls `boundsOf:`, a method that returns the geometrical shape and location of graphical elements. `boundsOf:` is defined by the class `MOFormsBuilder`, which belongs to `Layouts`. Therefore, a dependency of `Shapes` on `Layouts` is introduced, which is inconsistent with the documentation statement *No method of `Shapes` uses `Layouts`*. Programmers were notified by the reasoner about this problem, and version 544 of Mondrian is now consistent again.

What we have gained. Before our effort, Mondrian lacked documentation about its internal structure and contained a number of unknown structural errors. Now, its architecture and its components are described by entirely formal statements that are at the same time as easy to read and understand as natural English. Our documentation effort has been realized without interrupting the development of Mondrian. We made a direct, positive impact on the developers, seeing some engineering tasks and discussions on mailing lists in response to our discoveries.

6. Experiment

The case study presented above indicates that our approach is feasible, but it does not tell us whether controlled natural languages are better or worse than

other formal languages. For that reason, we performed a user experiment to find out how controlled natural language compares to other formal languages when used for source code documentation. The goal of this experiment is to assess how easy or difficult it is to understand such documentation statements in different formal languages.

This section investigates research question *B*. In contrast to question *A*, we are *not* using the Mondrian example to investigate this issue, which might require some explanation. For relying on Mondrian in such an experiment, we would have to let participants get familiar with Mondrian or we would have to only select participants that already know Mondrian sufficiently well. In the first case, preparation time would have to be much longer (probably several hours) than the time participants are usually available for. In the second case, it would be very hard to control for the previous experiences of the participants with the given exemplary code base, and it would be difficult to disentangle the effects of our approach with other types of documentation the users might have been in contact with before. Also, it would be virtually impossible to get enough participants with the required expertise. For these reasons, we decided to base the experiment on neutral and fictitious scenarios, presented as schematic diagrams.

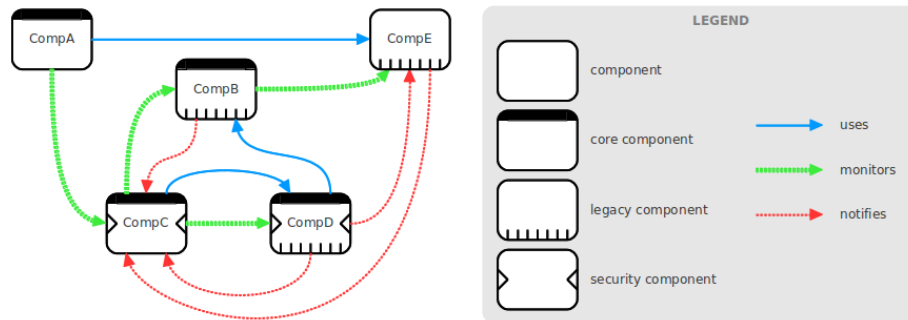
Besides ACE, we chose the SOUL syntax (as shown on the official SOUL website¹⁴) and Prolog syntax as the languages to be tested against each other. The latter two languages have both been applied to query and document source code structure [36, 37, 16, 38], and are the state of the art in this respect. Supplementary material of the experiment is available online.¹⁵

6.1. Experiment Design

The experiment design is based on a modified version of the ontograph framework [39]. The basic idea is to use simple diagrams as a neutral basis for comparison. Diagrams have the advantage that they are very different from textual languages, which minimizes the possible bias towards one of the languages to be tested. One of the four diagrams used in the experiment is shown in Figure 5. These diagrams show several components of different types and with different relations between them in a very simple graphical notation. They represent a very simple abstract model of a software system, which we can use to evaluate and compare the three languages we have chosen. For each of these diagrams, twenty statements had been written in each of the three languages in a way that each statement has semantically equivalent statements in both other languages. Some of these statements are true with respect to the diagram, others are false. The experiment is designed in a way that each participant is tested on all three languages. The order in which the languages are presented is shuffled in a balanced way.

¹⁴<http://soft.vub.ac.be/SOUL/>

¹⁵<http://bit.ly/AceWikiExperimentMaterial>



Which of the statements are true and which are false?

- | true | false | don't know | |
|-----------------------|-----------------------|-----------------------|--------------------------------------------------------------------------------------|
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | Every security-component that is used by a component notifies the component. |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | Every component that is notified by CompD is a legacy-component. |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | Every component that uses a legacy-component is a core-component. |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | Every security-component that does not notify CompE is monitored by CompA. |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | Every component that is notified by a legacy-component is a security-component. |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | ?x isCoreComponent if ?x isLegacyComponent |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | ?x notifies: ?y if ?x isSecurityComponent, ?y isComponent, ?y uses: ?x |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | ?x isCoreComponent if ?x isComponent, ?y isLegacyComponent, ?x uses: ?y |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | CompD notifies: ?x if ?x isComponent, or(?x monitors: CompB, CompA uses: ?x) |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | ?x isCoreComponent if ?x isComponent, or(?x isLegacyComponent, ?x monitors: CompC) |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | ?x notifies: CompC if ?x isLegacyComponent |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | is_legacy_component(x) :- is_component(x), notifies(compD,x). |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | is_core_component(x) :- is_component(x), \+ is_legacy_component(x). |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | notifies(compD,x) :- is_component(x), (monitors(x,compB) ; uses(compA,x)). |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | is_core_component(x) :- is_component(x), is_legacy_component(y), uses(x,y). |
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | notifies(x,compC) :- is_legacy_component(x). |

Figure 5: This figure is a blending of three screenshots (one for each of the three tested languages) of the testing phases of the experiment. Participants had to classify the statements at the bottom with respect to the diagram shown at the top.

During a learning phase of at most ten minutes, a short explanation of the language is shown to the participants together with a diagram and five true statements (in the respective language) marked as “true” and five false statements marked as “false”. After the learning phase, there is a testing phase of up to six minutes, during which the participants are presented another diagram and ten statements in the respective language. The objective of the participants is to classify each of the statements as “true” or “false” (they can also choose “don’t know”). Figure 5 shows how the screen looks during the testing phase. These learning and testing phases are repeated for each of the languages with different, but similar diagrams for the testing phases (these diagrams for the testing phases can be transformed into each other by a simple reordering and replacement of names, concepts and relation types). Thus, each of the twenty participants has to go through three testing phases consisting of ten statements each, which leads to overall 600 data points in the form of correct or incorrect classifications. After the three learning and testing phases, the participants have to fill out a short questionnaire, asking about their background and their experiences during the experiment.

Twenty experienced graduates and advanced undergraduates in the area of software engineering were recruited as participants for the experiment. They were on average, 27 years old, and all but one of them were male. Most of them (70%) had never used before a formal language similar to ACE; 60% and 45% had never used a language like SOUL or Prolog, respectively.

6.2. Experiment Results

The design of the experiment allows us to evaluate the differences of the three languages along three dimensions: the classification score, the time needed, and the preference of the participants derived from the questionnaire. Figure 6 shows the results.

The most important dimension is the classification score, i.e. the number of correctly classified statements out of ten. Since it is just a binary choice, a score of 5 can be achieved by mere guessing. “Don’t know” answers and missing answers (when the time limit ran out) are counted as halfway correct (0.5). With ACE, the participants managed to correctly classify an average of 9 out of 10 statements. With SOUL and Prolog, the average score was only 6.1 and 7.1, respectively. The dominance of ACE shows high statistical significance (see below). Thus, the participants understood ACE much better than SOUL or Prolog.

Looking at the time dimension, we see that ACE performed again clearly and significantly better than the other two languages. To learn ACE and to finish the classification task, the participants needed on average about 9 minutes out of 16 (a maximum of 10 for the learning phase plus 6 for the testing phase). In contrast, they needed about 12 minutes in the case of SOUL and Prolog.

The third dimension is derived from the questionnaire, which contains a question “How easy or hard to understand did you find the statements of language X?” for each of the languages. The possible answers are “very easy to understand” (value 3), “easy to understand” (2), “hard to understand” (1), and



Figure 6: The results of the experiment concerning classification score, time, and questionnaire score for ACE, SOUL and Prolog.

Table 1: The p -values for the differences between the three languages, using Wilcoxon signed rank tests.

	ACE~SOUL	ACE~Prolog	SOUL~Prolog
classification	0.0000076	0.000092	0.091
time	0.00012	0.000038	0.62
questionnaire	0.00011	0.00069	0.30

“very hard to understand” (0). ACE was perceived on average as being between “easy” and “very easy” to understand with a tendency to the latter. Prolog, on the other hand, was in the middle of “hard” and “easy”, and SOUL even a bit lower. Thus, ACE was perceived as much more understandable than the other languages.

To check the observed differences between the languages for statistical significance, we use the Wilcoxon signed rank test. This is a non-parametric test to check whether or not we can conclude from observed differences within a paired sample that the difference exists in reality. This test method is robust in the sense that it does not require the statistical population to be normally distributed. Table 1 shows the obtained p -values. These p -values represent the probability that the observed differences result by chance from a population where no differences are present. The lower these values, the more confident we can be that a difference exists in reality. All differences involving ACE are statistically significant on a 95% confidence level (and would even be significant on a 99.9% level). The differences between SOUL and Prolog are not significant.

In addition to the three dimensions explained above, the questionnaire contains the question “If you were a software developer who has to document a piece of software, which formal language would you prefer to use?”. 65% of the participants chose ACE, 20% chose SOUL and 15% chose Prolog. Thus, they did not only consider ACE to be more understandable, most of them would also prefer to use ACE for *writing* source code documentation.

In summary, ACE was understood much better, required less time to learn, and was preferred by the participants. In all explored dimensions, ACE clearly outperformed the other languages.

7. Discussion

7.1. Applicability to Other Programming Languages

Our results are based on the experiments we have carried out using the Pharo programming language, but our general approach is not restricted to a particular language. In our source code model, we have considered the following language elements: packages, classes, methods, invocations between methods and class instantiation. These elements are sufficient to capture the general architecture and dependencies of Pharo programs, but they also apply to many other (object-oriented) programming languages.

It is relatively easy to adopt the approach to other programming languages (but this would typically be a task of the tool developer and not of the end user). For example, to add the concept of *interfaces*, as e.g. Java supports it, the source code model could be extended by the following statements:

Every interface is a code element.
No interface is a class.
No interface is a method.
If X implements Y then X is a class.
If X implements Y then Y is an interface.

From this, the ACE parser can automatically generate the formal model, which can be mapped to the source code structure as provided by tools such as the Moose software analysis platform¹⁶, which offers a wide range of language parsers, including Pharo, VisualWorks, Java, C, Cobol, Delphi, and C++. Additionally, we use the FAMIX meta-model, which provides a language-independent model of programming language elements [40]. It includes concepts such as a class, method, package, and invocation, and is general enough to cover most structural and behavioral elements commonly found in programming languages.

Once these mappings are in place for a particular language such as Java, we can represent source code in ACE:

Runnable is an interface.
Thread is a class.
Thread implements Runnable.

These facts introduce the elements necessary to later attach documentation statements and to provide a mapping to the given FAMIX model.

7.2. Keeping queries simple

During our documentation effort, we had to query the source code of Mondrian multiple times in order to pinpoint anomalies. Our experience shows that we had to face a number of complex queries, making heavy use of subordinate clauses introduced with the conjunction *that*.

For example, to identify the causes for the dependencies between **Core**, **Layout** and **Events** we first formulated these questions:

Which class that belongs to Layout is used by a class that belongs to Core?
Which method that is defined by a class that belongs to Core uses a class that belongs to Layout?
Which method that is defined by a class that belongs to Core uses MOPortEvent?

Although these sentences are grammatically correct from a linguistic point of view, they are difficult to read. To reduce the complexity of these questions, we defined the relations *method of* and *class of*, which have already been mentioned above. In the case of *method of*, we first had the following facts to express the relation between a method and its class:

¹⁶<http://moosetechnology.org>

If X is a method of Y then X is a method.
If X defines Y then Y is a method of X.

We later extended the relation to cover the *belongs to* relation:

If X is defined by something that belongs to Y then X is a method of Y.

Thus, apart from connecting methods and classes, the relation *method of* can now also relate methods and components. This simple statement greatly simplifies the formulation of method belonging. The relation *class of* has been defined in a similar way. The three questions above can now be formulated in a much simpler way:

Which class of Layout uses a class of Core?
Which method of Core uses a class of Layout?
Which method of Core uses MOPortEvent?

7.3. Kind of documentation

In order to illustrate that various kinds of documentation statements can be expressed with our approach, we take the ten design entity attributes as defined by the IEEE standard 1016-1998 [41]: Identification, Type, Purpose, Function, Subordinates, Dependencies, Interface, Resources, Processing and Data. Below, we exemplify how each of them can be expressed, at least partially, in ACE:

- Unique identification is done by assigning a name, e.g., “Pleiad group”, to an entity by adding it to the AceWiki lexicon.
- Types like subsystem, module, process, algorithm or data store are defined as nouns in the AceWiki lexicon and entities can be assigned to these types by simple “*is a*”-sentences:

Core is a module.
The Medical Database is a data store.

- Even though the purpose of a certain entity cannot be fully defined in ACE (it is a complex issue to formally represent purposes), it is very easy to give names to certain purpose aspects and to link them to the entities:

A purpose of ResultsCache is performance improvement.
A purpose of ResultsCache is system stability.

- The function of an entity like input, transformations, and output can be documented:

*The Attempto Parsing Engine processes Attempto Controlled English
and returns the ACE DRS format.*
The Medical Database stores every patient record.

- Subordinate entities can be expressed in AceWiki with relations like *contains* or *part of*:

*MySystem contains exactly 10 modules.
Every module is a part of MySystem.
Core contains the Data Base Manager.*

- Dependencies can be expressed as follows:

*ModuleA depends on ModuleB.
Every subclass of QueryWindow depends on a subclass of DataBase.*

- Much information about interfaces, e.g., which methods to call from the outside and which parameters to use, is already present in the source code in many cases. However, AceWiki allows for additional documentation:

*Every method that invokes DataStore-acquireLock invokes
DataStore-releaseLock.
Every method that invokes a test method belongs-to the Test module.*

- Dependencies on external resources can be defined as follows:

*The Data Base Manager requires MySQL.
The Receipt Generator requires a printer.*

- The processing attribute is about how an entity achieves its function, e.g., which algorithms it uses:

ListSorter uses the Quicksort algorithm.

- The data attribute should contain information about the internal data structures:

ElementDirectory uses a hash table.

Thus, ACE and AceWiki allow for the documentation of all ten of these design entity attributes, at least in principle and to a certain degree. In some cases like for the purpose attribute, it is not possible to give a full description in ACE, because it is difficult to represent such things in logic. But, as the examples show, even in these cases, it is possible to name the involved concepts and relate them to the affected individuals.

7.4. Performance

To measure the performance of AceWiki for larger code bases, we micro-benchmarked our approach on seven applications, including ArgoUML,¹⁷ a popular large software for case studies. The amount of facts to describe these applications ranges from around 1 400 to more than 170 000 (ArgoUML). Mondrian leads to around 23 000 facts.

The results are presented in Figure 7. For each of the applications, it shows the time needed (i) to load and parse the source code, (ii) to add a fact to the knowledge database, (iii) to check the consistency of the knowledge database, and (iv) to detect an inconsistency.

¹⁷<http://argouml.tigris.org>

Software	#classes	#methods	#facts	load facts	add fact	check for consist.	detect inconsist.
Adore	12	163	1 442	14.12	0.019	0.02	0.07
Gofer	32	234	2 276	20.05	0.15	0.46	0.06
Autotest	56	200	1 701	16.22	0.014	0.09	0.24
Mondrian	148	1 902	23 342	25.56	0.55	0.55	4.88
Glamour	157	1 286	17 798	159.35	0.94	0.44	5.26
Moose	357	4 218	75 021	2774.17	1.98	0.76	6.40
ArgoUML	2 526	20 338	173 982	7835.25	2.40	2.50	14.76

Figure 7: Micro benchmarks of AceWiki (the unit of time for the last four columns is seconds)

The figures show that the OWL reasoner FaCT++ behaves relatively well when adding facts and checking for consistency. Adding a fact to a database with 1 442 facts takes 0.019 seconds. For the largest application, it takes 2.40 seconds. Loading the knowledge base is much slower. It takes about 14 seconds for 1 441 facts, 25 seconds for Mondrian and more than 7 800 seconds for ArgoUML. This is because the reasoner requires the complete ontology to be loaded into memory in order to perform reasoning tasks. However, there is certainly a lot of room for improvement in this respect, e.g., by adopting techniques used by highly scalable triple stores or by excluding some of the most complex OWL features, which typically lead to a drastic drop in time consumption.

AceWiki did not encounter problems loading a larger knowledge base consisting of more than 23 000 statements. We were not able to precisely measure the memory consumption, but we experienced a consumption of 817 MB of memory to construct and load the Mondrian database. Considering the resources available in modern computers, this does not seem to be a critical issue. This benchmark was carried out on an Ubuntu Linux 2.6.32-24 with 2GB of RAM and a 2.4 GHz Intel Core 2 Duo.

7.5. Usability

The Mondrian case study described above indicates that AceWiki is suitable for the representation of software documentation. The results of the experiment show that ACE statements are easy to understand for software engineers and that they prefer ACE over other formal languages to document source code. One question that is not answered by the presented studies, however, is whether developers are *able to write* sensible ACE statements with AceWiki. This aspect has been researched in previous work.

Two usability experiments for AceWiki have been performed with altogether 26 participants [27]. The results showed that AceWiki and its editor are easy to learn and use. The participants — without being trained to use AceWiki or ACE — managed to add many statements to AceWiki in a short period of time. On average, each user created a correct and sensible statement every 5–6 minutes, *while* defining a new word every 5–7 minutes. 78%–81% of the statements were correct and sensible, and 61%–70% of them were complex in the sense that they contained negations, implications, disjunctions, or number restrictions. The

participants were mostly students, with no particular background in software engineering. Since software engineers are used to working with complex software tools, it is reasonable to assume that they perform at least as well as average students. Another study confirmed that writing ACE sentences with the editor used in AceWiki is easier and faster than writing in other formal languages [42].

Another experiment with 64 participants evaluated the understandability of ACE, independently of a particular tool [43]. The results showed that ACE is not only easier to understand than a comparable logic language but also needs less time to be learned and is preferred by users. Under a strict time limit, the participants managed to correctly classify more than 90% of given ACE statements (which semantically covered a broad subset of OWL) as true or false with respect to certain situations. The participants were students or graduates with no higher education in computer science or logic. Again, we believe that software developers can be expected to be at least as skilled in logical thinking as average students, and should therefore be able to correctly interpret ACE statements with a high accuracy.

8. Related Work

The use of logic to describe and check software architecture has been explored in depth [44, 45]. In this context, a number of different tools have been implemented, e.g., Reflexion Models [46], ArchJava [47], Lattix Inc’s dependency manager [48], and Intentional Views [49]. However, in contrast to the approach presented here, the resulting formal models of the architecture cannot be read and queried in a natural way and, as far as we are aware of, they have not been used in practice to *document* software. Users are supposed to learn how to read and write statements in some sort of formal logic. This could be a major hindrance for the broad adoption of such systems, especially if the model of the architecture has to be verified by non-computer-scientists.

The use of wikis for software documentation is quite common. There even exists a wiki engine — Dokuwiki — that is specifically dedicated to documentation [23]. However, the documentation cannot be checked for consistency or be used to answer specific questions since wikis lack semantic capabilities.

Würsch et al. [17] present an approach to use controlled natural language in the context of software engineering that is in some respects very similar to our approach. Their system allows developers to ask questions about their source code in a controlled language. The difference to our approach is that only questions are supported. There is no possibility to augment the underlying model by annotations or documentation statements.

Kaplanski [50] also explored the use of controlled English to model object-oriented systems, which is very similar to our approach in many respects. However, he does not provide evaluation results or a running prototype at this point.

Kimmig *et al.* [51] propose an approach for querying source code using a natural language. Users write queries in a simple question that follows the pattern question word - verb - noun - verb (e.g., “Where is balance read?” with “balance” being an instance variable). Our approach differs on two points: (i) With

our approach, new concepts and relations are easy to define. Their approach is fairly limited in that respect. (ii) They apply a sophisticated procedure of cleaning and tokenizing the natural language queries, before formalizing them. With our approach, documentation is written in a precise and unambiguous language in the first place.

Buse and Weimer [22] synthesize succinct human-readable documentation from software modifications. They employ an approach based on code summarization and symbolic execution, summarizing the runtime conditions necessary for the control flow to reach a modified statement. It is unclear, however, whether such techniques are useful for more than helping to write version log messages.

9. Conclusion and Future Work

To conclude, we can come back to the three challenges we identified in Section 2 for a documentation system based on ACE to be effective and reliable, and we can review these challenges against our approach:

Validation. With our approach, the consistency of the documentation can be automatically checked by an OWL reasoner. The documentation is incrementally validated when new facts are added or removed.

Immediate feedback. After the initial loading of the whole code source and documentation data, our approach considers only incremental modifications. As we have shown, the cost of adding a new fact is only a few seconds or less. We can reasonably assume that the code production rate of a software developer remains below the processing time of our approach. Our system is therefore able to quickly provide feedback for any given modification.

Minimal impact on software engineers. In our approach, the documentation is written in ACE, a language that can be learned quickly. Our experiment showed that ACE is easy to understand and preferred by developers over other languages. Ideally, the documentation statements could be embedded in the source code as comments and supported by the IDE. This minimizes the impact on software engineers.

The importance of writing documentation about how a software operates does not need to be demonstrated. However, this activity is rarely a priority when not made an explicit requirement. In this article we showed that a reliable solution may be achieved with a controlled natural language interpreter and facilities for checking consistency and answering questions. We showed that controlled natural language is easier to understand than other languages, needs less time to be learned, and is preferred by the users. Our approach validates documentation against the actual source code, provides immediate feedback, and leads to only a small amount of additional work for software engineers.

Agile methods promote the use of unit tests as an executable kind of documentation. In the same spirit, we consider our work as a small step towards the vision of a fully automated code documentation system.

For the future, we plan to propose a concrete format to include ACE statements in source code comments, and to provide an IDE plugin. The relevant features of AceWiki should be included in this plugin and a reasoner should be integrated. In this way, the IDE can help developers to write ACE sentences, notify them about problems after changes in the source code or the documentation, and answer questions about the code structure and its documented environment. In addition, we plan to include an explanation module, which some OWL reasoners provide, in order to show to the developers *why* a particular question gave the respective answers or *why* a statement triggered inconsistency. This system should be designed in a way that hides the complexity of formal logic and feels like a natural and intuitive extension of the programming environment. We plan to evaluate the impact of such documentation tools on open source communities, particularly in the agile programming field.

Acknowledgment

We would like to thank Tudor Doru Gîrba, Roel Wuyts and the whole Pleiad research group for their feedback on an early draft of this article. We also thank all the people who participated in our experiment.

We gratefully thank Renato Cerro for revising an earlier version of this article.

This work has been partially funded by Program U-INICIA 11/06 VID 2011, grant U -INICIA 11/06, University of Chile, and FONDECYT project 1120094.

References

- [1] P. Clements, D. Garlan, R. Little, R. Nord, J. Stafford, Documenting software architectures: views and beyond, in: ICSE '03: Proceedings of the 25th International Conference on Software Engineering, IEEE Computer Society, Washington, DC, USA, 2003, pp. 740–741.
- [2] V. Jakobac, A. Egyed, N. Medvidovic, [Improving system understanding via interactive, tailorable, source code analysis](#), in: M. Cerioli (Ed.), Fundamental Approaches to Software Engineering, Vol. 3442 of Lecture Notes in Computer Science, Springer Berlin / Heidelberg, 2005, pp. 253–268, 10.1007/978-3-540-31984-9_19.
URL http://dx.doi.org/10.1007/978-3-540-31984-9_19
- [3] A. Marcus, J. I. Maletic, Recovering documentation-to-source-code traceability links using latent semantic indexing, in: ICSE '03: Proceedings of the 25th International Conference on Software Engineering, IEEE Computer Society, Washington, DC, USA, 2003, pp. 125–135.

- [4] R. Pierce, S. Tilley, Automatically connecting documentation to code with rose, in: SIGDOC '02: Proceedings of the 20th annual international conference on Computer documentation, ACM, New York, NY, USA, 2002, pp. 157–163. doi:10.1145/584955.584979.
- [5] L. Shi, H. Zhong, T. Xie, M. Li, An empirical study on evolution of api documentation, in: D. Giannakopoulou, F. Orejas (Eds.), Fundamental Approaches to Software Engineering, Vol. 6603 of Lecture Notes in Computer Science, Springer Berlin / Heidelberg, 2011, pp. 416–431. doi:10.1007/978-3-642-19811-3_29.
- [6] A. Wyner, K. Angelov, G. Barzdins, D. Damljanovic, B. Davis, N. Fuchs, S. Hoefler, K. Jones, K. Kaljurand, T. Kuhn, M. Luts, J. Pool, M. Rosner, R. Schwitter, J. Sowa, On controlled natural languages: Properties and prospects, in: N. E. Fuchs (Ed.), Proceedings of the Workshop on Controlled Natural Language (CNL 2009), Vol. 5972 of Lecture Notes in Computer Science, Springer, Berlin / Heidelberg, Germany, 2010, pp. 281–289.
- [7] T. Kuhn, A survey and classification of controlled natural languages, Computational Linguistics.
- [8] N. E. Fuchs, K. Kaljurand, T. Kuhn, Attempto Controlled English for knowledge representation, in: C. Baroglio, P. A. Bonatti, J. Małuszyński, M. Marchiori, A. Polleres, S. Schaffert (Eds.), Reasoning Web — 4th International Summer School 2008, no. 5224 in Lecture Notes in Computer Science, Springer, 2008, pp. 104–124.
- [9] J. L. De Coi, N. E. Fuchs, K. Kaljurand, T. Kuhn, Controlled English for reasoning on the Semantic Web, in: F. Bry, J. Małuszyński (Eds.), Semantic Techniques for the Web — The REWERSE Perspective, Vol. 5500 of Lecture Notes in Computer Science, Springer, 2009, pp. 276–308.
- [10] T. Kuhn, *AceWiki: A natural and expressive semantic wiki*, in: D. Degler, m. schraefel, J. Golbeck, A. Bernstein, L. Rutledge (Eds.), Proceedings of the Fifth International Workshop on Semantic Web User Interaction (SWUI 2008) — Exploring HCI Challenges, Vol. 543 of CEUR Workshop Proceedings, CEUR-WS, 2009.
URL http://ceur-ws.org/Vol-543/kuhn_swui2008.pdf
- [11] J. Sillito, G. C. Murphy, K. De Volder, *Questions programmers ask during software evolution tasks*, in: Proceedings of the 14th ACM SIGSOFT international symposium on Foundations of software engineering, SIGSOFT '06/FSE-14, ACM, New York, NY, USA, 2006, pp. 23–34. doi:10.1145/1181775.1181779.
URL <http://doi.acm.org/10.1145/1181775.1181779>
- [12] M.-A. D. Storey, K. Wong, H. A. Müller, How do program understanding tools affect how programmers understand programs?, in: I. Baxter,

- A. Quilici, C. Verhoef (Eds.), Proceedings Fourth Working Conference on Reverse Engineering, IEEE Computer Society, 1997, pp. 12–21.
- [13] B. de Alwis, G. C. Murphy, Answering conceptual queries with ferret, in: Proceedings of the 30th International Conference on Software Engineering (ICSE), ACM, New York, NY, USA, 2008, pp. 21–30. doi:[10.1145/1368088.1368092](https://doi.org/10.1145/1368088.1368092).
- [14] P. C. Rigby, M. P. Robillard, [Discovering essential code elements in informal documentation](#), in: Proceedings of the 2013 International Conference on Software Engineering, ICSE '13, IEEE Press, Piscataway, NJ, USA, 2013, pp. 832–841.
URL <http://dl.acm.org/citation.cfm?id=2486788.2486897>
- [15] H. Zhong, Z. Su, [Detecting api documentation errors](#), in: Proceedings of the 2013 ACM SIGPLAN international conference on Object oriented programming systems languages & applications, OOPSLA '13, ACM, New York, NY, USA, 2013, pp. 803–816. doi:[10.1145/2509136.2509523](https://doi.org/10.1145/2509136.2509523).
URL <http://doi.acm.org/10.1145/2509136.2509523>
- [16] R. Wuyts, S. Ducasse, [Symbiotic reflection between an object-oriented and a logic programming language](#), in: ECOOP 2001 International Workshop on MultiParadigm Programming with Object-Oriented Languages, 2001.
URL <http://scg.unibe.ch/archive/papers/Wuyt01a.pdf>
- [17] M. Würsch, G. Ghezzi, G. Reif, H. C. Gall, Supporting developers with natural language queries, in: ICSE '10: Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering, ACM, New York, NY, USA, 2010, pp. 165–174. doi:<http://doi.acm.org/10.1145/1806799.1806827>.
- [18] W. W. Royce, Managing the development of large software systems, in: Proc. IEEE Wescon, 1970, pp. 1–9.
- [19] C. J. Stettina, W. Heijstek, [Necessary and neglected?: an empirical study of internal documentation in agile software development teams](#), in: Proceedings of the 29th ACM international conference on Design of communication, SIGDOC '11, ACM, New York, NY, USA, 2011, pp. 159–166. doi:[10.1145/2038476.2038509](https://doi.org/10.1145/2038476.2038509).
URL <http://doi.acm.org/10.1145/2038476.2038509>
- [20] D. E. Knuth, Literate Programming, Stanford, California: Center for the Study of Language and Information, 1992.
- [21] S. C. B. de Souza, N. Anquetil, K. M. de Oliveira, [A study of the documentation essential to software maintenance](#), in: Proceedings of the 23rd annual international conference on Design of communication: documenting & designing for pervasive information, SIGDOC '05, ACM, New York, NY, USA, 2005, pp. 68–75. doi:[10.1145/1085313.1085331](https://doi.org/10.1145/1085313.1085331).
URL <http://doi.acm.org/10.1145/1085313.1085331>

- [22] R. P. Buse, W. R. Weimer, [Automatically documenting program changes](#), in: Proceedings of the IEEE/ACM international conference on Automated software engineering, ASE '10, ACM, New York, NY, USA, 2010, pp. 33–42. doi:10.1145/1858996.1859005.
URL <http://doi.acm.org/10.1145/1858996.1859005>
- [23] M. Badger, [Dokuwiki — a practical open source knowledge base solution](#), Enterprise Open Source Magazine 4 (10).
URL <http://opensource.sys-con.com/node/318853>
- [24] A. Bacchelli, M. D'Ambros, M. Lanza, Are popular classes more defect prone?, in: Proceedings of FASE 2010 (13th International Conference on Fundamental Approaches to Software Engineering), 2010, pp. 59–73.
- [25] G. Sridhara, L. Pollock, K. Vijay-Shanker, Generating parameter comments and integrating with method summaries, in: Program Comprehension (ICPC), 2011 IEEE 19th International Conference on, 2011, pp. 71–80. doi:10.1109/ICPC.2011.28.
- [26] R. P. Buse, W. R. Weimer, [Automatic documentation inference for exceptions](#), in: Proceedings of the 2008 international symposium on Software testing and analysis, ISSTA '08, ACM, New York, NY, USA, 2008, pp. 273–282. doi:10.1145/1390630.1390664.
URL <http://doi.acm.org/10.1145/1390630.1390664>
- [27] T. Kuhn, [How controlled English can improve semantic wikis](#), in: C. Lange, S. Schaffert, H. Skaf-Molli, M. Völkel (Eds.), Proceedings of the Forth Semantic Wiki Workshop (SemWiki 2009), Vol. 464 of CEUR Workshop Proceedings, CEUR-WS, 2009.
URL <http://ceur-ws.org/Vol-464/paper-03.pdf>
- [28] M. Krötzsch, D. Vrandečić, M. Völkel, H. Haller, R. Studer, Semantic Wikipedia, Web Semantics: Science, Services and Agents on the World Wide Web 5 (4) (2007) 251–261. doi:10.1016/j.websem.2007.09.001.
- [29] S. Schaffert, J. Eder, S. Grünwald, T. Kurz, M. Radulescu, R. Sint, S. Stroka, [KiWi — a platform for semantic social software](#), in: C. Lange, S. Schaffert, H. Skaf-Molli, M. Völkel (Eds.), Proceedings of the Fourth Workshop on Semantic Wikis — The Semantic Wiki Web, Vol. 464 of CEUR Workshop Proceedings, CEUR-WS, 2009.
URL <http://ceur-ws.org/Vol-464/paper-17.pdf>
- [30] K. Kaljurand, T. Kuhn, A multilingual semantic wiki based on Attempto Controlled English and Grammatical Framework, in: Proceedings of the 10th Extended Semantic Web Conference (ESWC 2013), Springer, 2013.
- [31] C. Bock, A. Fokoue, P. Haase, R. Hoekstra, I. Horrocks, A. Ruttenberg, U. Sattler, M. Smith, [OWL 2 Web Ontology Language — structural specification and functional-style syntax](#), W3C Candidate Recommendation,

- <http://www.w3.org/TR/2009/CR-owl2-syntax-20090611/> (June 2009).
URL <http://www.w3.org/TR/2009/CR-owl2-syntax-20090611/>
- [32] D. Tsarkov, I. Horrocks, FaCT++ Description Logic reasoner: System description, in: Automated Reasoning — Proceedings of the Third International Joint Conference IJCAR 2006, Vol. 4130 of Lecture Notes in Computer Science, Springer, 2006, pp. 292–297.
 - [33] M. Meyer, T. Gırba, M. Lungu, [Mondrian: An agile visualization framework](#), in: ACM Symposium on Software Visualization (SoftVis'06), ACM Press, New York, NY, USA, 2006, pp. 135–144. doi:10.1145/1148493.1148513.
URL <http://scg.unibe.ch/archive/papers/Meye06aMondrian.pdf>
 - [34] A. Lienhard, A. Kuhn, O. Greevy, [Rapid prototyping of visualizations using mondrian](#), in: Proceedings IEEE International Workshop on Visualizing Software for Understanding (Vissoft'07), IEEE Computer Society, Los Alamitos, CA, USA, 2007, pp. 67–70. doi:10.1109/VISSOF.2007.4290702.
URL <http://scg.unibe.ch/archive/papers/Lien07bMondrian.pdf>
 - [35] A. Bergel, S. Ducasse, O. Nierstrasz, R. Wuyts, [Classboxes: Controlling visibility of class extensions](#), Journal of Computer Languages, Systems and Structures 31 (3-4) (2005) 107–126. doi:10.1016/j.cl.2004.11.002.
URL <http://scg.unibe.ch/archive/papers/Berg05aclassboxesJournal.pdf>
 - [36] R. Wuyts, Declarative reasoning about the structure of object-oriented systems, in: Technology of Object-Oriented Languages, Proceedings of TOOLS-26'98, IEEE, 1998, pp. 112–124.
 - [37] C. Kramer, L. Prechelt, Design recovery by automated search for structural design patterns in object-oriented software, in: Proceedings of the 3rd Working Conference on Reverse Engineering (WCRE '96), IEEE Computer Society, Washington, DC, USA, 1996, pp. 208–215.
 - [38] J. Fabry, T. Mens, Language-independent detection of object-oriented design patterns, Journal of Computer Languages, Systems and Structures 30 (1-2) (2004) 21–33.
 - [39] T. Kuhn, An evaluation framework for controlled natural languages, in: N. E. Fuchs (Ed.), Proceedings of the Workshop on Controlled Natural Language (CNL 2009), Vol. 5972 of Lecture Notes in Computer Science, Springer, Berlin / Heidelberg, Germany, 2010, pp. 1–20.
 - [40] S. Demeyer, S. Tichelaar, S. Ducasse, FAMIX 2.1 — The FAMOOS Information Exchange Model, Tech. rep., University of Bern (2001).
 - [41] IEEE recommended practice for software design descriptions, IEEE standard 1016, IEEE Computer Society (1998).

- [42] T. Kuhn, S. Höfler, Coral: Corpus access in controlled language, *Corpora* 7 (2) (2012) 187–206.
- [43] T. Kuhn, [The understandability of OWL statements in controlled English](#), *Semantic Web* 4 (1) (2013) 101–115.
URL <http://www.semantic-web-journal.net/content/understandability-owl-statements-controlled-english>
- [44] K. M. Mens, R. Wuyts, T. D’Hondt, Declaratively codifying software architectures using virtual software classifications, *Technology of Object-Oriented Languages*, International Conference on (1999) 33.
- [45] R. Wuyts, A logic meta-programming approach to support the co-evolution of object-oriented design and implementation, Ph.D. thesis, Department of Computer Science, Vrije Universiteit Brussel (January 2001).
- [46] G. C. Murphy, D. Notkin, K. Sullivan, Software reflexion models: bridging the gap between source and high-level models, in: *Proceedings of the 3rd ACM SIGSOFT symposium on Foundations of software engineering*, ACM, New York, NY, USA, 1995, pp. 18–28.
- [47] J. Aldrich, C. Chambers, D. Notkin, ArchJava: connecting software architecture to implementation, in: *Proceedings of the 24th International Conference on Software Engineering*, ACM, New York, NY, USA, 2002, pp. 187–197.
- [48] N. Sangal, E. Jordan, V. Sinha, D. Jackson, Using dependency models to manage complex software architecture, in: *Proceedings of the 20th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*, ACM, New York, NY, USA, 2005, pp. 167–176.
- [49] K. Mens, A. Kellens, F. Pluquet, R. Wuyts, Co-evolving code and design with intensional views, *Computer Languages, Systems and Structures* 32 (2006) 140–156.
- [50] P. Kaplanski, [Modeling object oriented systems via controlled English verbalization of Description Logic](#), in: M. Rosner, N. E. Fuchs (Eds.), *Pre-Proceedings of the Second Workshop on Controlled Natural Languages (CNL 2010)*, Vol. 622 of CEUR Workshop Proceedings, CEUR-WS, 2010.
URL <http://ceur-ws.org/Vol-622/paper17.pdf>
- [51] M. Kimmig, M. Monperrus, M. Mezini, [Querying source code with natural language](#), in: *Proceedings of the 26th IEEE/ACM International Conference On Automated Software Engineering*, 2011, pp. 376–379.
URL <http://www.monperrus.net/martin/Querying-Source-Code-with-Natural-Language.pdf>