

Context Capture in Software Development

Bruno Antunes*, Francisco Correia and Paulo Gomes

Knowledge and Intelligent Systems Laboratory
Cognitive and Media Systems Group
Centre for Informatics and Systems of the University of Coimbra
Coimbra, Portugal
{bema, fcorreia, pgomes}@dei.uc.pt
<http://www.cisuc.uc.pt/>

Abstract. The context of a software developer is something hard to define and capture, as it represents a complex network of elements across different dimensions that are not limited to the work developed on an IDE. We propose the definition of a software developer context model that takes into account all the dimensions that characterize the work environment of the developer. We are especially focused on what the software developer context encompasses at the project level and how it can be captured. The experimental work done so far show that useful context information can be extracted from project management tools. The extraction, analysis and availability of this context information can be used to enrich the work environment of developers with additional knowledge to support their work.

Key words: Context, Software Development, Knowledge Management.

1 Introduction

The term context has an intuitive meaning for humans, but due to this intuitive connotation it remains vague and generalist. Furthermore, the interest in the many roles of context comes from different fields such as literature, philosophy, linguistics and computer science, with each field proposing its own view of context [1]. The term context typically refers to the set of circumstances and facts that surround the center of interest, providing additional information and increasing understanding.

The context-aware computing concept was first introduced by Schilit and Theimer [2], where they refer to context as “*location of use, the collection of nearby people and objects, as well as the changes to those objects over time*”. In a similar way, Brown et al. [3] define context as location, identities of the people around the user, the time of day, season, temperature, etc. In a more generic definition, Dey and Abowd [4] define context as “*any information that can be used to characterize the situation of an entity. An entity is a person,*

* Supported by FCT grant SFRH/BD/43336/2008.

place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves”.

In software development, the context of a developer can be viewed as a rich and complex network of elements across different dimensions that are not limited to the work developed on an IDE (Integrated Development Environment). Due to the difficulty on approaching such challenge, there is not a unique notion of what it really covers and how it can be truly exploited. With the increasing dimension of software systems, software development projects have grown in complexity and size, as well as in the number of functionalities and technologies involved. During their work, software developers need to cope with a large amount of contextual information that is typically not captured and processed in order to enrich their work environment.

Our aim is the definition of a software developer context model that takes into account all the dimensions that characterize the work environment of the developer. We propose that these dimensions can be represented as a layered model with four main layers: personal, project, organization and domain. Also, we believe that a context model needs to be analyzed from different perspectives: capture, modeling, representation and application. This way, each layer of the proposed context model will be founded in a definition of what context capture, modeling, representation and application should be for that layer.

This work is especially focused on the project layer of the software developer context model. We give a definition of what the context model encompasses at the project layer and present some experimental work on the context capture perspective.

The remaining of the paper starts with an overview of the software developer context model we envision. In section 3 we describe the current work on context capture, some preliminary experimentation and the prototype developed. An overview of related work is given in section 4. Finally, section 5 concludes the paper and point out some directions for future work.

2 Context Model

The software developer context model we propose takes into account all the dimensions that comprise the software developer work environment. This way, we have identified four main dimensions: personal, project, organization and domain. As shown in figure 1, these dimensions form a layered model and will be described from four different perspectives: context capture, context modeling, context representation and context application.

The *personal layer* represents the context of the work a developer has at hands at any point in time, which can be defined as a set of tasks. In order to accomplish these tasks, the developer has to deal with various kinds of resources at the same time, such as source code files, specification documents, bug reports, etc. These resources are typically dispersed through different places and systems, although being connected by a set of explicit and implicit relations that exist




Fig. 1. The software developer context model layers and perspectives.

between them. At this level the context model represents the resources that are important for the tasks the developer is working on.

The *project layer* focuses on the context of the project, or projects, in which the developer is somehow involved. A software development project is an aggregation of a team, a set of resources and a combination of explicit and implicit knowledge that keeps the project running. The team is responsible for accomplishing tasks, which end up consuming and producing resources. The relations that exist between people and resources are the glue that makes everything work. The project layer represents the people and resources, as well as their relations, of the software development projects where the developer is included.

The *organization layer* takes into account the organization context to which the developer belongs. Similarly to a project, an organization is made up of people, resources and their relations, but in a much more complex network. While in a project the people and resources are necessarily connected due to the requisites of their work, in a software development organization these projects easily become separate islands. The knowledge and competences developed in each project may be of interest in other projects and valuable synergies can be created when this information is available. The organization layer represents the organizational context that surrounds a developer.

The *domain layer* takes into account the knowledge domain, or domains, in which the developer works. This layer goes beyond the project and organization levels and includes a set of knowledge sources that stand out of these spheres. Nowadays, a typical developer uses the Internet to search information and to keep informed of the advances in the technologies s/he works with. These actions are based on services and communities, such as source code repositories, development forums, news websites, blogs, etc. These knowledge sources cannot be detached

from the developer context and are integrated in the domain layer of our context model. For instance, due to the dynamic nature of the software development field, the developer must be able to gather knowledge from sources that go beyond the limits of the organization, either to follow the technological evolution or to search for help whenever needed.

The four context dimensions described before can be described through four different perspectives: *context capture*, which represents the sources of context information and the way this information is gathered, in order to build the context model; *context modeling*, which represents the different dimensions, entities and aspects of the context information (conceptual model); *context representation*, which represents the technologies and data structures used to represent the context model (implementation); and *context application*, which represents how the context model is used and the objectives behind its use.

3 Context Capture

Our current work is focused on the project layer of our developer context model, and we will discuss our work at this level from the different perspectives we have presented before.

Concerning context capture, the main sources of contextual information that feed up the developer context model at the project level are project management tools. These tools store a big amount of explicit and implicit information about the resources produced during a software development project, how the people involved relate with these resources and how the resources relate to each other. We are focusing on two types of tools: *Version Control Systems* (VCS) and *Issue Tracking Systems* (ITS). As shown in figure 2, the former deals with resources and their changes, the second deals with tasks. These systems store valuable information about how developers, tasks and resources relate and how these relations evolve over time. We are especially interested in revision logs and tasks. Briefly described, a revision log tell us that a set of changes were applied by a developer to a set of resources in the repository. A task commonly represents a problem report and the respective fix.

The network of developers, resources, tasks and their relations will be used to build our context model at the project level. This way, the context model of the developer, from a project point of view, will be modeled as a set of implicit and explicit relations, as shown in figure 3. The lines with a filled pattern represent the explicit relations and those with a dotted pattern represent the implicit ones. The developers are explicitly related with revisions, as they are the ones who commit the revisions, and with tasks, as each task is assigned to a developer. The relation between tasks and resources is not explicit, but we believe it can be identified by analyzing the information that describe tasks and revisions. The proximity between developers can be inferred by analyzing the resources were they share work. Also, the resources can be implicitly related by analyzing how often they are changed together.

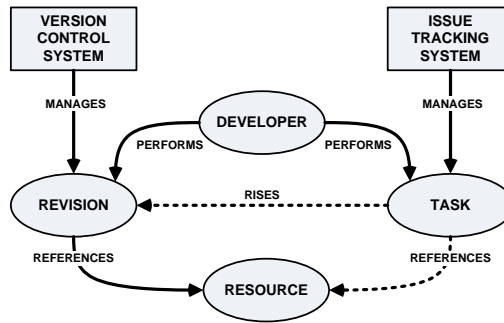


Fig. 2. The project layer relevant entities and their roles.

In order to extract relations from the information stored in project management tools, that information is previously imported and stored locally for analysis. The prototype developed uses a database to store both the imported data and extracted relations. In the next phase, we intend to represent these relations and connection elements in an ontology [5], which will gradually evolve to a global developer context model ontology. We believe that representing the context model in an ontology and formalising it using the Semantic Web [6] technologies promote knowledge sharing and reusability, since these technologies are standards accepted by the scientific community.




Fig. 3. The elements that compose the context model in the project layer.

The context information extracted at the project level will be used to inform the developer about the network that links people, resources and tasks on the project s/he works. This information can be prepared to facilitate consulting and presented to the developer easily accessible in her/his working environment.

3.1 Relation Extraction

We have implemented connectors that allowed us to collect all the desired information from the *Subversion* and *Bugzilla* systems, as they are among the most popular in use. Through the collected information, we could already perceive a set of explicit relations: which resources are created/changed by which developers, which tasks have been assigned to which developers (see relations number 1 and 2 in figure 3). There is also a set of implicit relations that would be valuable if disclosed.

Our approach to extract implicit relations between resources and tasks (see relation number 3 in figure 3) relies on the analysis of the text provided by revision messages, task summaries and task comments. It is common to find references to task identifiers in revision messages, denoting that the revision was made in the context of a specific task. Also, task summaries and comments commonly reference specific resources, either because a problem has been identified in a specific class or because a error stack trace is attached to the task summary to help developers identify the source of the problem. Taking this into account, we have defined three algorithms to find resource/task and task/revision relations:

- *Resource/Task (I)*. For each resource referenced in a revision, the respective name was searched in all task summaries. The search was performed using the file name and, in case it was a source code file, the full qualified name (package included) and the class name separately.
- *Resource/Task (II)*. For each resource referenced in a revision, the respective name was searched in task comments. This search was performed as described for the previous relation.
- *Task/Revision*. For each task, the respective identification number was searched in revision messages. The search was performed using common patterns such as “<id>”, “bug <id>” and “#<id>”.

The implicit relations between resources (see relation number 4 in figure 3) can be extracted by analyzing resources that are changed together very often. Revisions are often associated with specific goals, such as the implementation of a new feature or the correction of a bug. When developers commit revisions, they typically change a set of resources at a time, those that needed to be changed in order to accomplish a goal. When two resources are changed together in various revisions, this means that these resources are somehow related, even if they do not have any explicit relation between them, because when one of them is modified there is a high probability that the other also needs to be modified.

The proximity between developers (see relation number 5 in figure 3) can also be inferred by analyzing the resources they share work. Developers can share work when they commit revisions on the same resources or when they are assigned to tasks that are related to the same resources. This way, if two developers often make changes, or perform tasks, on the same resources, they are likely to be related.

Table 1. Number of extracted relations.

	Resource/Task (I)	Resource/Task (II)	Task/Revision
gEclipse	2527	31671	629
Subversive	208	9076	773

3.2 Preliminary Results

To validate the relation extraction algorithms, these were tested against two open-source projects from the Eclipse foundation:

- *gEclipse*. The gEclipse framework allows users and developers to access Computing Grids and Cloud Computing resources in a unified way, independently of the Grid middleware or Cloud Computing provider. Our analysis was performed over the work of 19 developers in a time window of approximately 3 years and 3 months, which included 3279 revisions and 1508 tasks with 7765 comments.
- *Subversive*. The Subversive project supports the development of an Eclipse plug-in for Subversion integration. Our analysis was performed over the work of 10 developers in a time window of approximately 3 years and 2 months, which included 1120 revisions and 1013 tasks with 3252 comments.

By applying the relation extraction algorithms to the information related with these two projects, we have gathered the results represented in table 1. The table shows the number of distinct relations extracted using each one of the algorithms in the two projects analyzed.

The results show that a large amount of implicit relations can be extracted from the analysis of the information associated to tasks and revisions. These relations complement the context model we are building by connecting tasks with related resources. With a more detailed analysis we have identified some problems with the algorithms creating some relations that do not correspond to an effective correlation between the two entities analyzed. These problems are mainly related with string matching inconsistencies and can be corrected with minor improvements in the way expressions are searched in the text.

3.3 Prototype

We have developed a prototype, in the form of an Eclipse plug-in, to show how the context information can be integrated into the an IDE and used to help developers. In figure 4 we show a screenshot of the prototype, where we can see an Eclipse View named “Context” that shows context information related to a specific resource. Each time the developer opens a resource, this view is updated with a list of developers, resources and tasks that are related with that resource through the relations we have described before. This way the developer can easily gather information about what resources are likely to be related with

that resource, what other tasks affected that resource in the past, and what other developers may be of help if extra information is needed. The availability of this information inside the IDE, where developers perform most of their work, increases developers awareness and reduces their effort on finding information that would be hidden and dispersed otherwise.

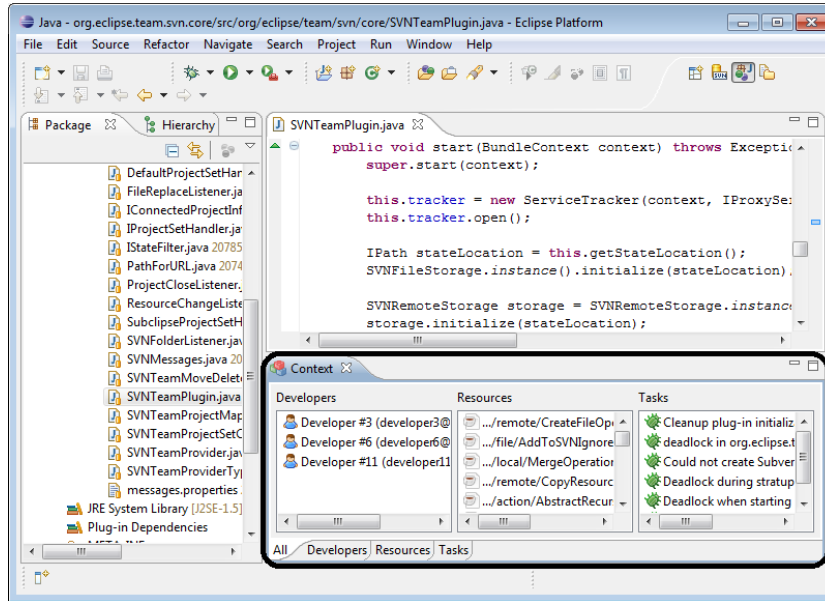


Fig. 4. The context plugin for Eclipse.

4 Related Work

The EPOS (Evolving Personal to Organizational Knowledge Spaces) project, presented by Schwarz [7], aims to build organizational structured knowledge from information and structures owned by the elements of the organization. The world of a knowledge worker is defined as containing document-like objects, objects used for classification and applications. This work environment is taken into account when modeling the user context, which comprises information about various aspects, including currently or recently read documents, relevant topics, relevant persons, relevant projects, etc. The context model is formalized using RDFS [8], and each user's context is an up-to-date instance of this context model. The context information is gathered and modeled through user observation and/or user feedback. The gathering and elicitation of contextual information is done in two ways: context-polling, by requesting a snapshot

of the user's current context; and context-listening, by subscribing the Context Service to be informed of every change in the user's context. Being a developer a knowledge worker, much of the concepts referenced in this work apply to the software development domain, but the specificities of this domain demand for a context model adapted to the reality of the work environment of a software developer.

Modularity is in the basis of the development of complex software systems and largely used to support a programmer's tasks, but not always help the programmer finding the desired information or delimit the point of interest for a specific task. Based on this, Kersten and Murphy [9] have been working on a model for representing tasks and their context. The task context is derived from an interaction history that comprises a sequence of interaction events representing operations performed on a software program's artifact. They then use the information in a task context either to help focus the information displayed in the IDE, or to automate the retrieval of relevant information for completing a task. The focus of this work is the task and the knowledge elements present in the IDE that are more relevant for the fulfillment of that task. Our approach aims to define a context model that goes beyond the IDE and explores the knowledge provided by the different systems that support the software development process.

In the same line of task management and recovery, Parnin and Gorg [10] propose an approach for capturing the context relevant for a task from a programmer's interactions with an IDE, which is then used to aid the programmer recovering the mental state associated with a task and to facilitate the exploration of source code using recommendation systems. Their approach is focused on analyzing the interactions of the programmer with the source code, in order to create techniques for supporting recovery and exploration. Again, this approach is largely restricted to the IDE and the developer interaction with it.

With the belief that customized information retrieval facilities can be used to support the reuse of software components, Henrich and Morgenroth [11] propose a framework that enables the search for potentially useful artifacts during software development. Their approach exploits both the relationships between the artifacts and the working context of the developer. The context information is used to refine the search for similar artifacts, as well as to trigger the search process itself. The context information model is represented with RDF [12] statements and covers several dimensions: the user context, the working context, and the interaction context. While the focus here is in software reuse, our approach focuses on the information captured during the process development.

5 Conclusions

We have presented our approach of a software developer context model. Our context model is based on a layered structure, taking into account four main dimensions of the work environment of a software developer: personal, project, organization and domain.

The current work is focused on the project layer of the software developer context model. We have discussed this layer in more detail and presented preliminary experimentation on the context capture perspective. The results show that it is possible to relate tasks and revisions/resources using simple relation extraction algorithms. These relations were then used in a plug-in for Eclipse to unveil relevant information to the developer.

As future work we plan to improve the prototype we have developed with better visualization, search and filtering functionality. We also want to explore the use of ontologies to represent the developer context model. The remaining layers of the context model will be addressed iteratively, as an extent to the work already developed. Finally, we intend to test our approach with developers working in with real world projects.

References

1. Mostefaoui, G.K., Pasquier-Rocha, J., Brezillon, P.: Context-aware computing: A guide for the pervasive computing community. In: Proceedings of the IEEE/ACS International Conference on Pervasive Services, ICPS 2004. (2004) 39–48
2. Schilit, B., Theimer, M.: Disseminating active map information to mobile hosts. *IEEE Network* (1994) 22–32
3. Brown, P.J., Bovey, J.D., Chen, X.: Context-aware applications: From the laboratory to the marketplace. *Personal Communications, IEEE* 4 (1997) 58–64
4. Dey, A.K., Abowd, G.D.: Towards a better understanding of context and context-awareness. In: CHI 2000 Workshop on the What, Who, Where, When, and How of Context-Awareness, The Hague, The Netherlands (2000)
5. Zuniga, G.L.: Ontology: Its transformation from philosophy to information systems. In: Proceedings of the International Conference on Formal Ontology in Information Systems, ACM Press (2001) 187–197
6. Berners-Lee, T., Hendler, J., Lassila, O.: The semantic web. *Scientific American* 284 (2001) 34–43
7. Schwarz, S.: A context model for personal knowledge management applications. In: Modeling and Retrieval of Context, 2nd International Workshop (MRC 2005), Edinburgh, UK (2005)
8. Brickley, D., Guha, R.V.: Rdf vocabulary description language 1.0: Rdf schema (2004) Published: W3C Recommendation.
9. Kersten, M., Murphy, G.C.: Using task context to improve programmer productivity. In: Proceedings of the 14th ACM SIGSOFT International Symposium on Foundations of Software Engineering, Portland, Oregon, USA, ACM (2006) 1–11
10. Parnin, C., Gorg, C.: Building usage contexts during program comprehension. In: Proceedings of the 14th IEEE International Conference on Program Comprehension (ICPC'06). (2006) 13–22
11. Henrich, A., Morgenroth, K.: Supporting collaborative software development by context-aware information retrieval facilities. In: DEXA Workshops, IEEE Computer Society (2003) 249–253
12. Miller, E., Manola, F.: Rdf primer (2004) Published: W3C Recommendation.