# Automatic Discovery of Personal Action Contexts

Marielba Zacarias[1,3], H. Sofia Pinto[2], and Jose Tribolet[2,3]

[1] Universidade do Algarve
[2] Instituto Superior Tecnico
[3] INESC-Inov-CEO

**Abstract.** Enterprise modeling is an effective tool in designing and communicating enterprise operation and structure. The inclusion of the concept of context enables capturing and representing situated behaviors from actual execution. A context-based modeling approach is being developed to (1) capture and represent personal and inter-personal work patterns from action repositories, and (2) align them with enterprise activities and resources. This paper describes one activity of such approach; the discovery of personal action contexts and presents results on using clustering for their automatic discovery.

## 1 Introduction

Organizations communicate, document and understand their activity through models [1]. Enterprise modelling (EM) in Artificial Intelligence (AI) and Information Systems disciplines have been mainly used as communication tools to facilitate the design and implementation of business applications [2]. There are several EM frameworks available today. Some well-known generic frameworks are the Integrated Framework Architecture (IAF) [3], and the Open Group Architecture Framework (TOGAF)[4]. Within AI, two well known EO are the Enterprise Ontology proposed by [5] and the ontologies of the TOVE project [6].

IS/IA EM frameworks share two main characteristics. First, they allow representing different enterprise concerns in terms of perspectives. The most commonly depicted are *process*, *information*, *application* and *technology* perspectives [2]. Whereas the former describes enterprise *activities*, the remaining perspective describe its *resources*. Second, enterprise models are described with semi-formal or formal languages and most of them enable graphical representations. EM research and practice has shown that these languages reduce ambiguities and misunderstandings. The communicational power of enterprise models can be valuable tools in analyzing and (re)designing not only the organization, but also the behavior its members. Nonetheless, achieving this purpose entails overcoming some limitations. First, EM capture organization's design i.e. they use concepts (activity, resource) that are *abstractions* of what organizations do and as such, their relationship to daily actions is not always clear. Second, current EM frameworks offer an objective, 'aerial' view that assumes that activity and resource definitions are shared by all members of the organization.

The inclusion of the concept of context in enterprise models enables *situated* representations reflecting the behavior of given agents, in specific circumstances, from their particular viewpoints. A conceptual framework composed of an (1) architecture and ontology of organizational agents and contexts, and (2) a context-based methodological approach to build representations based on the ontology is proposed in [7, 8]. The goal of this framework is to enrich EM to capture and represent personal and inter-personal work patterns from actual actions and interactions, and to relate them with enterprise activities and resources. The methodological approach encompasses six activities; bootstrapping, action capture, context discovery, context analysis and context integration. This paper provides details of the context discovery activity and presents results in applying clustering as a means to support this activity.

The remainder of this paper is structured as follows; section 2 describes the fundamental concepts of our framework, section 3 summarizes the model acquisition activities giving more details of the context discovery activity, section 4 describes the usage of clustering in the automatic discovery of personal action contexts, section 5 gives our conclusions and future directions.

## 2 A Model Centered on Agents and Contexts

Figure 1 shows the fundamental concepts of this work, and the relationships between them. In this framework, activities and resources are design-related concepts. Agents, actions, roles, and contexts are related to execution. *Resources* are the entities (things, information or persons) relevant for the organization's operation. Resources are identified with nouns, and may be simple i.e. single items, or complex i.e. composed of several items. *Agents* are special resources with acting, coordination/management and adaptation capabilities. In our framework, agents may be human, automated or semi-automated, individual or collective. In this work, the concept of agent refers only to single individuals.

*Activities* are *abstractions* of what agents do. Activities use resources *(inputs)* and produce resources *(outputs)*. Activities are related to given organizational goals. Activities are specified with procedures i.e. the required steps to accomplish them. *Actions* are atomic acts performed by agents to change the state of resource-related items. Actions may communicative or non-communicative. *Interactions* are adjacent pairs of communicative actions exchanged between agents. Activities are accomplished with actions. However, due to the abstract nature of activities, the relationship between actions and activities will depend on the activity definition and a single action may be related to several activities.

*Roles* define the *observable* behaviour of agents. The specific role played by agents is determined by a given context. *Contexts* define situations continuously (re)created by agent actions and interactions. Contexts are related to (1) particular topics, problems, or other work situations, (2) the interacting agents, (3) time-related factors. This notion of context draws on human and social sciences [9–11]. In this fields, context is regarded as a fuzzy and dynamic entity

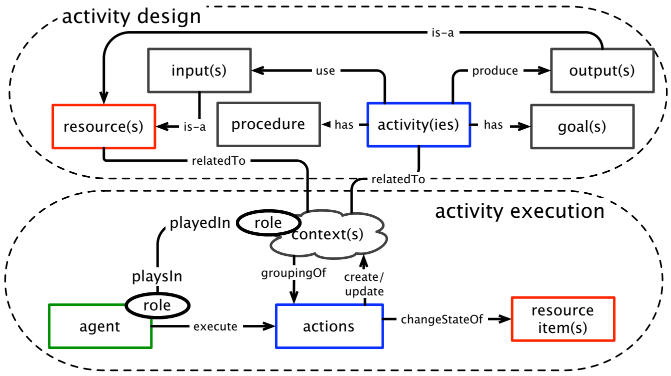


**Fig. 1.** Design and Execution Concepts

that emerges from interactions among entities. Contexts also reflect a particular viewpoint of a given agent.

Since activities are abstractions, the specific relationship of contexts and activities depend of the definition of activities. Hence, activity and context do not necessarily have the same boundaries. Firstly, whereas the execution of an activity may be related to several contexts and action streams, a context may be related to several activities. Secondly, due to the temporal nature of contexts, similar action streams performed at different times may be associated to the same activity but to a different context.

### 2.1 An Example Context

The "Prof. Smith's payment" context is created by the following (e-mail mediated) actions: (1) Prof. Smith *requests* the payment of a course, (2) Alice *requests* payment requirements, (3) Prof. Smith *explains* why he has not sent the requirements, *requests* an exception, and *promises* to deliver the payment requirements on date X, (5) Alice *analyses* Prof. Smith's request and seeks advice from her boss, and (6) Alice *informs* Prof. Smith of her decision. The boundaries of this context are defined by the "Prof. Smith's payment" topic or subject; the agents *Alice, Alice's boss* and *Prof. Smith*; and the date when these actions were executed.

The payment context may be related to several activities such as *pay professors, request payments*, and *prepare payment requirements*. Conversely, these activities are related to several other contexts (not depicted due to space limitations). Thirdly, actions use resource items that are not necessarily activity inputs or outputs. They may be *transient* resources, meaningful only within a particular context (e.g. reason for not sending requirements). Other resource items are *tools* supporting execution that are not part of activity design (e.g. e-mail).

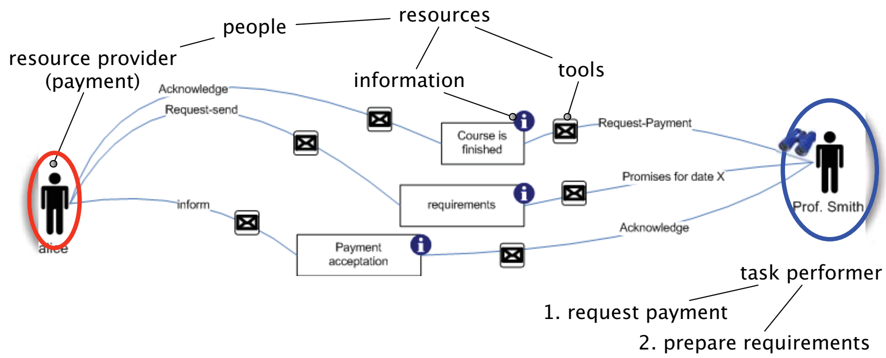


**Fig. 2.** Personal Context - Prof. Smith's view

*Personal contexts* reflect the *personal viewpoint* of an individual over a given context. Figure 2 shows the view that Prof. Smith has of the payment context. First, in personal contexts, owner agents (in this case, Prof. Smith) regard themselves as playing the role of *task performers*, while the remaining agents play resource provider or consumer roles. Second, the agent's view depends on the actions the agent is aware of. In this particular case, Alice requests advice from her boss about Prof. Smith's request. Professor Smith is not aware of the communicative actions between Alice and her boss. Hence, these interactions will not appear in Prof. Smith's view.

## 3 Model Acquisition

The fundamental concepts described in section 2 suggests firstly, that activity or resource-related behaviours of agents cannot be dissociated from their contexts of execution. Secondly, actual agent behaviours need to be captured from actions and interactions. Consequently, agent behavior is captured using a bottom-up and context-based approach, where we collect actions of a group of subjects, identify and analyse *action/interaction patterns* within given contexts.

Our approach encompasses six activities; (1) bootstrapping, (2) action capture, (3) context discovery, (4) Context visualization (5) context analysis and (6) context integration. In *bootstrapping*, the basic action types and resources to be registered are defined, and their meanings discussed. Ideally, action and resource definitions are registered. The *action capture* activity creates action repositories. Actions are registered along with its execution date, using a structure defined in [8]. Since individuals may perform several, unrelated tasks and switch among them, it is also necessary to identify action streams, that is, if an action follows an action previously registered, the number of the preceding action need to be registered. When relevant, object descriptions are complemented with the supporting resources used in performing each action (tools, information, people).

Table depicted in figure 3 shows how the action stream of the example context described in section 2, looks like once collected and structured.

| | | | | | Object Description\| Action | | | |
|---|---|---|---|---|---|---|---|---|
| nº | Date | foll. | Subject | Action Type | Receiver | Nested Action | Main resource-related items | ..supporting Resources (tools, people, information items) |
| ... | ... | ... | ... | ... | ... | ... | ... | ... |
| **20** | 1-Apr | 0 | Prof. Smith | **request** | Alice | **pay** | course X | e-mail |
| **21** | 1-Apr | 20 | Alice | **check** | | | Prof. Smith's payment requirements | excel, payment requirement records |
| **22** | 1-Apr | 21 | Alice | **request** | Prof. Smith | **send** | Prof. Smith's payment requirements | e-mail |
| **23** | 1-Apr | 22 | Alice | **inform** | | **pay** | will proceed when requirements are sent | e-mail |
| **24** | 1-Apr | 23 | Prof. Smith | **request** | Alice | **pay** | without payment requirements | e-mail |
| **25** | 1-Apr | 24 | Prof. Smith | **inform** | Alice | | reason for not sending requirements | |
| **26** | 1-Apr | 25 | Prof. Smith | **inform** | Alice | **promise** | requirements for date D | e-mail |
| **27** | 2-Apr | 26 | Alice | **analyze** | | | payment request and reason given | |
| **28** | 2-Apr | 0 | Alice | **request** | Boss | **analyze** | accept or reject Prof. Smith's request | telefone |
| **29** | 2-Apr | 28 | Alice's boss | **analyze** | | | payment request and reason given | |
| **30** | 2-Apr | 29 | Alice's boss | **suggest** | Alice | **accept** | payment request of Prof. Smith | telefone |
| **31** | 2-Apr | 30 | 28 | **inform** | Alice | | prof. Smith is a good professor | telefone |
| **32** | 2-Apr | 27 | Alice | **accept** | Prof. Smith | **pay** | course X | e-mail |
| **33** | 2-Apr | 32 | Alice | **order** | Luisa | **pay** | course X to Prof. Smith | |
| **34** | 2-Apr | 33 | Alice | **inform** | Prof. Smith | **pay** | is ordered | e-mail |

**Fig. 3.** Sample Action Log

*Context discovery* entails identifying, characterizing, and labeling personal action contexts. In this activity, personal contexts become 'entities' characterized by a set of keywords. *Context visualization* displays main context characteristics to their owners, for validation purposes. In *context-based analysis*, personal contexts are used as units of analysis for several purposes:

- Discovering personal work and resource usage patterns
- Discovering context switching behavior to capture and model human multitasking at work. Since personal contexts group together similar actions and resources, they provide a good estimation of work fragmentation.
- Linking inter-related pairs of personal contexts of two given individuals allows to uncover and characterize inter-personal contexts and further, context-based interaction networks.
- A proper association of actions with activities and resources require analyzing action groupings rather than single actions. Hence, the identification and characterization of personal action contexts allows to relate daily actions to organizational activities and resources.

In *Context integration*, context-based patterns are related to enterprise activities and resources. Those considered as good practices, trigger changes in current task/resource models. In this paper, we focus the *context discovery* activity. A detailed description of all activities is provided in [8].

### 3.1 Discovering Contexts

The identification of personal contexts entails identifying and labeling of meaningful action groupings of single individual. The *context discovery and visualizations* activities involve first, grouping action streams involving a similar set of resources. These groupings labeled, characterized are shown to their corresponding performers who may re-arrange them. Validated groupings define subject's *personal action contexts*. Each personal context is labeled by their owners.

Conceptually, contexts are updated by each executed action. However, visualizing an exact, up-to-date depiction of all contexts is not practical, neither necesary. Hence, this work does not aim at displaying exact depictions of every on-going context. Rather, the goal is to display approximate pictures of them. This approximate picture is given by a set of key features that defined as *context keywords*, which allow characterizing contexts. Through these keywords, it is possible to classify individual actions into their corresponding context(s).

*Identifying Agents.* Identifying agent keywords in personal action contexts is straightforward. Since personal contexts are defined by actions performed by a single individual, it is represented by the *subject* of a particular action grouping. For example, in the *prof. Smith's payment context*, the agent is 'Alice'.

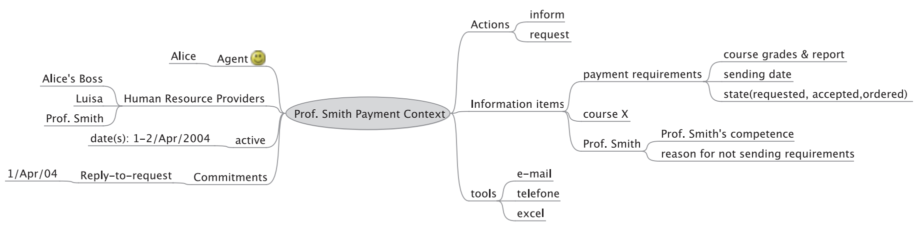


**Fig. 4.** Payment Context "Keywords" (Alice)

*Identifying predominant action types.* Being the most recurrent action types of the corresponding action log (Fig. 3), *inform*, and *request* are selected as action keywords(fig. 4). Though the identification of action keywords typically corresponds to the most frequent action types, some unfrequent actions may be included. When logs are collected within restricted time intervals, relevant actions not frequently executed during that particular interval may be left out. Hence, the keyword selection must ultimately be decided by each context owner.

*Identifying predominant resource-related items.* The identification of resource-related keywords follows the same logic of action keyword identification. Most frequent resource-related items will typically become resource-related keywords. However, as in the case of action types, the final selection is ultimately decided by the context owner. In figure 4, 'payment requirements', and 'Prof. Smith' are identified as keywords because they group together several items related to them. Though not appearing several times, 'course X' was selected because Alice identified it as a distinguishing item of this context. The tools 'e-mail', 'telefone', and 'excel' were also included to indicate the most common supporting tools. Finding resource-related candidate keywords is more difficult than finding frequent action types. Whereas action types are identified through single and previously separated words (verbs), resource related items are embedded within non-structured object descriptions and they may be represented by noun or noun phrases.

*Identifying time-related keywords.* The date(s) when contexts are typically active refer to the date interval(s) grouping the greatest number of actions. The definition of this keyword will vary according criteria defined by both the observers and observed subjects. In the example, this keyword reflects the execution date of all actions (1-2 Apr). The commitments handled are directly related to the typical actions. Being *request*, a recurrent action within this context, it is inferred that the *reply-to-request* commitment had to be frequently handled.

## 4 Using Clustering to Discover Personal Contexts

Clustering is a *data mining* technique applied to discover groups of related data objects based on attribute similarity within data sets, without any prior knowledge of the group definition [12]. Clustering techniques apply when instances are to be divided in natural groups. These clusters presumably reflect some mechanism at work in the domain from which instances are drawn, causing some instances to bear a stronger resemblance to one another than with the remaining instances. Since the clustering underlying mechanisms are rarely known, the choice of clustering approaches is usually dictated by the available tools.

There are three main clustering methods; (1) Nearest Neighbor rule known as K-means, (2) incremental and (3) statistical clustering. Statistical or probabilistic clustering assigns instances to clusters probabilistically. Probabilistic clustering supposes identifying the probability density functions of data source. Each distribution governs the attribute values of a different cluster. An efficient representation of the probability density function is the *mixture model*, which asserts that the data is a combination of k individual component densities, corresponding to the k clusters [13]. When the class of an instance is known, the cluster distribution gives the probability of an instance having a certain attribute value set. Since data records may belong to all k clusters but with different probability, the mixture model allows overlapping clusters. A well-known implementation of the mixture model is the Expectation-Maximization (EM) algorithm [13]. Due to its probabilistic nature, arbitrary shaped clustering can be effectively represented by the choice of suitable component density functions (poisson, spherical or non-spherical Gaussians for numeric attributes, and binomial or multinomial distribution for categorical or discrete data).

### 4.1 Case Study

The usage of clustering in discovering contexts was tested in a case involving a software development team of a commercial bank. The team was integrated by 4 programmers (Gonçalo, Carla, Catarina, Alexandre) and the project leader (Mariana). During the observation, the team worked on applications to manage *Suppliers, Claims, Mail, Evictions*, and *Marketing Campaigns*. In this case, a three-week observation was conducted, where over 650 actions were collected, and grouped in personal contexts by the observed subjects. Figure 5 show the contexts identified by the subject Carla during the observation period, along with its identifying keywords.

| Ctx | Actions | Human Resources | Information Items | Tools |
|---|---|---|---|---|
| c1 | program | pedro, mariana | common services application | google, msdn, sqlserver, visual studio dotnet |
| c2 | answer, discuss, help | alexandre, mariana | mail application, suppliers application | mail application, suppliers-app sw |
| c3 | accept, assist | catarina, mariana, alexandre, goncalo | team meeting, project status, resource distribution plan | e-mail |

**Fig. 5.** Personal Contexts of the Subject Carla

### 4.2 Data Preparation

The format of the data to be clustered was integrated by a record set with the following columns: (1) *number* that reflects the action chronological order, (2) *day* of occurrence, (3) *agent-sender* performing the action, or sending a communicative action, (4) *agent-receiver* of the communicative action,(5) *type* of the action performed, (6) *nested action* embedded in communicative actions, (7) *description* of the object of the action. In the case of communicative actions, it describes the object of the embedded action within it, (8) *tools* applicational or technological items used in performing actions, (9) *Information* items, (10) *human resources* used in performing actions.

This format was well fitted for manual context identification. It enabled a visual appreciation of action attributes looking at one or two of lines. Nonetheless, its reuse for automated means required rethinking not only the attributes to be included, but also their format and the extension of their domain, that is, the set of values of each attribute. One essential step in preparing data was thus to devise means of extracting attribute values from free-form textual fields.

*Attribute Selection:* The attributes were selected (1) accordingly to the definition of personal contexts and (2) from the experience acquired from the manual selection process. Three types of attributes were selected; *action, resource*, and *date-related attributes*. Recurrent action types form part of the characteristics of personal contexts. This assumption was demonstrated in the manual process, where in each grouping it was possible to identify a couple of predominant action types. Thus, the `action type` was selected. Though relevant, the `nested action` was not included due to the high number of non-communicative actions, where this field is empty. Since the action groupings are based on resource similarity, all fields containing information related to resources were included, whether this information referred to toos, information or human resources. Hence, the `agent-receiver, description, tools` and `human resources` attributes were selected. The `information items` attribute was not selected because it was noted in the manual process that the items included in this column were already present in the `description` colummn and it was thus, redundant. Finally, the `day` and `number` attributes were selected in order to discriminate clusters with similar action types and resources, but performed in different time intervals and which according to the definition of contexts, may belong to different contexts.

| Keyword Te rm | frequency |
|---|---|
| suppliers application | 192 |
| Claims application | 105 |
| Team meeting | 58 |
| evictions web service | 42 |
| Mail application | 31 |

**Fig. 6.** Some recurrent description items

**Data cleansing and transformation:** A second iteration of cleansing and transformation was accomplished for the attributes selected, using Sql Server Integration Services (SSIS) ®. Summarizing, it encompassed the following steps: (1) Worksheet data was transferred to a Sql Server database table. (2) In order to guarantee consistent and logic attribute value sets, dictionaries were created for the *receiver, action, description, tools* and *human resource* attributes. These dictionaries were built using a term extraction service to identify recurrent nouns and/or noun phrases from the textual data of the corresponding fields, along with the frequency of each term. The table in figure 6 shows a sample of the dictionary created for the description column. (3) The dictionaries created were analyzed to detect remaining errors or inconsistencies using a fuzzy grouping service that applies fuzzy logic to group similar terms. (4) Detected typo errors and name inconsistencies were removed from the action log, using simple data base operations. (5) The final *description* attribute dictionary was analyzed to determine the most meaningful keywords. This analysis resulted in the selection of description terms formed by noun phrases with a frequency $\geq 5$.

### 4.3 Modeling

This section summarizes the results of the modeling phase. A probabilistic clustering approach was selected among other approaches due to (1) the characteristics of the problem domain and (2) the mixed nature of the attributes. Regarding the problem domain, since contexts are regarded as entities with no clear-cut boundaries, accepting the existence of overlapping clusters, where a single action is associated to several contexts with a certain probability, seems the more appropriate approach. Regarding the attribute nature, the data set collected involves some numeric attributes. However, most of them involve textual data, for which probability clustering is adequate. The tool use for clustering was the Microsoft Clustering Algorithm of Sql Server Analysis Services ®.

The algorithm requires several parameters that includes the clustering method, number of clusters to be found and a group of tuning parameters. The first parameter was set for a simple version of probabilistic clustering, adequate for small data sets. Since the goal is to explore clustering as means for automating the discovery o personal contexts, the clustering process was accomplished separately for each individual and the cluster count parameter was set to the value corresponding to the number of contexts identified by each individual in the manual context identification process. For the tuning parameters, default values were used since the size of the data set did not require different values.

### 4.4 Clustering Results

Figure 7 depicts the cluster diagram for the subject *carla*, once analyzed and labelled. The diagram shows three clusters, where cluster 'commonSservicesApplic' is darker because this cluster groups a greater number of actions than the remaining clusters. The link between clusters 'DevelopmentSupport' and cluster 'TeamMeetings' indicates that there is a greater similarity between them than between these two clusters than with the former cluster. Understanding these clusters requires looking at the cluster characteristics, and comparing them with the keywords of manually identified clusters.

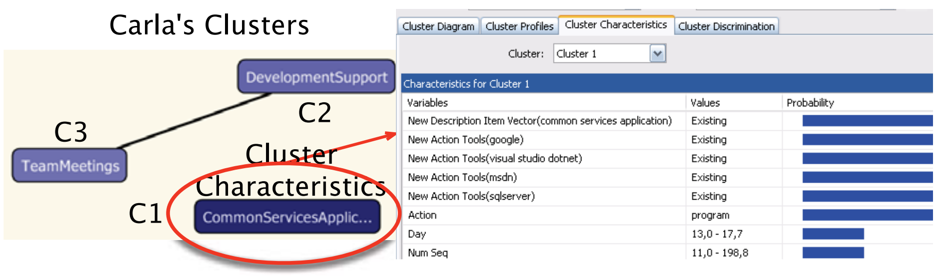


**Fig. 7.** Clustering Results

Figure 7 depicts the top attributes of the cluster 'Common Services Application' in decreasing probability. The top characteristics of cluster 1 (fig. 7) include the description item *common services application*, the action *program*, and the tool terms *google, msdn, visual studio dotnet* and *sqlserver*. These characteristics are very similar to the keywords identified for the personal context *common* *services application*. This similarity allowed labelling cluster 1 as the 'common-ServicesApp..' (c1) in figure 5. Likewise, cluster 2 characteristics were similar to the development support context (c2), and cluster 3 to the team meetings context.

### 4.5 Evaluating Clustering Results

The clustering process was evaluated both qualitative and quantitatively. The qualitative evaluation involved a visual comparison of the characteristics of manually identified contexts and the clusters produced by the algortithm. As a result, a correspondence matrix relating manual and automatic clusters was built. Whereas in several cases a one-to-one mapping was possible, in others, this mapping could not be established. In order to obtain a more precise accuracy measure, another comparison matrix was built. In this matrix, the rows identify the manually identified contexts and columns identified the clusters produced by the algorithm for each subject. The cluster with (1) the greatest number of actions belonging to a same manual context, and (2) similar characteristics to the manual context, was considered as the correct cluster. Actions of the same context scattered in other clusters were considered as an incorrect classification.

Accuracy was estimated adding the number of correctly grouped actions and dividing them by total number of actions. Accuracy estimates for each cluster, individual and overall accuracy were calculated. At a cluster level of each subject, programming and team meetings contexts exhibit more accuracy. This is consistent with the qualitative evaluation since those contexts were identified most easily than others. At an individual level, accuracy ranged from 0.89 (Carla) to 0.56 (Mariana). Accuracy diminished as the number of contexts increased. The overall accuracy estimate (0.71) indicates that over 70% of the total actions were correctly grouped. Figure 8 depicts the accuracy estimates of Carla's clusters.

| Context Description | Cluster 1 | Cluster 2 | Cluster 3 | *Total Context* | **Success Rate** |
|---|---|---|---|---|---|
| Common Services Application Programming | ***20*** | 1 | | *21* | **0.95** |
| Development Support | | ***8*** | 2 | *10* | **0.80** |
| Team Meetings | | 1 | ***11*** | *12* | **0.92** |
| *Total Cluster* | *20* | *10* | *13* | *43* | ***0.89*** |

**Fig. 8.** Carla's clusters Evaluation Results

## 5 Conclusions and Outlook

This paper describes the results of using clustering techniques in the automatic discovery of personal contexts from manually captured actions. A relevant part of data preparation was the extraction of recurrent noun phrases from action descriptions, which served as the attribute values of the action description attribute. As indicated by all case study subjects, action descriptions were taken into account in the manual identification of personal contexts. Hence, including description items as part of clustering attributes was essential.

A probabilistic clustering algorithm was applied to the actions of each subject, separately. A qualitative and quantitative evaluation procedure was defined to compare manually identified contexts with clusters produced by the algorithm. Whereas the automatic process produced acceptable starting point clusters in almost all cases, in the case of the team leader (Mariana), who handled a greater number of contexts, clustering quality requires further improvement. Cluster quality can be improved (1) including automatic capture mechanisms that would reduce data inconsistency and errors, (2) determining attribute dependencies, and (3) determining attribute degree of relevance for context detection. It is noteworthy that some manual groupings are too fine-grained, regarding the number of actions registered. Moreover, some groupings obey to user knowledge not explicit in action data and thus, cannot be found through automated techniques. Hence, a successful usage of this technique in supporting context identification requires taking into account their limitations.

Currently, the benefits of our approach have been validated in small case studies. Automated support for all activities is required in order to employ it in wider settings and during longer periods. We are researching in proper means of automating action capture, as well as context discovery and analysis activities. Action capture prototypes are being tested in real settings.The present work represents a first effort in automating the context discovery activity. The selection of the clustering method and tool was motivated by the importance of using technologies readily available for enterprises. Nonetheless, other means need to be explored, in order to determine their advantages and inconveniences.

## References

1. Caetano, A., Silva, A., Tribolet, J.: Object-oriented business process modeling with roles. In: 7th International Conference on Information Systems Implementation Modeling. (2004)
2. Shekkerman, J.: How to Survive in the Jungle of Enterprise Architecture Frameworks. Trafford (2004)
3. Cap Gemini.: The Integrated Architecture Framework IAF http://www.capgemini.com/services/soa/ent_architecture/iaf/ (2007)
4. Open Group.: The Open Group architectural framework (TOGAF) - version 8.1 http://www.opengroup.org/architecture/togaf8-doc/arch/ (2003)
5. Uschold, M., King. M., Moralee, S., Zorgios, Y.: The Enterprise Ontology In: Knowledge Engineering Review, 13(1) (1998) 31–89
6. Fox, M.S.: The TOVE Ontology Project http://www.eil.utoronto.ca/enterprise-modelling/tove/index.html (2002)
7. Zacarias, M., Magalhães, R., Caetano, A., Pinto, H., Tribolet, J.: Towards Organizational Self-Awareness: An Initial Architecture and Ontology. In: Ontologies for Business Interactions. IDEA Group (2008) 101–120

8. Zacarias, M., Serendero, P., Pinto, H.S., Tribolet, J.: Capturing and modeling work practice: A context-based approach. Revue d'Intelligence Artificialle (In Press)
9. Kokinov, B.: A dynamic approach to context modeling. In Brezillon, P., Abu-Hakima, S., eds.: Working Notes of the IJCAI'95 Workshop on Modelling Context in Knowledge Representation and Working Notes of the IJCAI'95 Workshop on Modelling Context in Knowledge Representation and Reasoning. (1995) 59–72
10. Latour, B.: Reassembling the Social : An Introduction to Actor-Network. Oxford University Press (2005)
11. Engeström, Y., Miettinen, R., Punamäki, R.: Perspectives on Activity Theory. Cambridge University Press, Cambridge MA, USA (2005)
12. Witten, I.H., Frank, E.: Data Mining. Morgan Kaufmann Publishers (2000)
13. Bradley, P., Fayyad, U., Reina, C.: Scaling EM (expectation - maximization) clustering to large databases. Technical Report MSR-TR-98-35, Microsoft Research (1999)