

Mining processes, resource consumption and witnesses for task completion from logs

Deliverable D3.2

FFG – IKT der Zukunft
SHAPE Project
2014 – 845638



■ **Table 1** Document Information

Project acronym:	SHAPE
Project full title:	Safety-critical Human- & dAta-centric Process management in Engineering projects
Work package:	3
Document number:	3.2
Document title:	Mining processes, resource consumption and witnesses for task completion from logs
Version:	1
Delivery date:	01 October 2015 (M2)
Actual publication date:	_____
Dissemination level:	Public
Nature:	Report
Editor(s) / lead beneficiary:	WU Vienna
Author(s):	Saimir Bala, Cristina Cabanillas, Jan Mendling, Axel Polleres
Reviewer(s):	Claudio Di Ciccio, Tudor Ionescu

Contents

1	Introduction	1
2	Implementation of a project mining technique in Camunda	2
3	Project mining with annotations	7
3.1	Activity labels	7
3.2	Resource and resources type	8
3.2.1	Roles of resources based on file	9
3.2.2	Annotating which is the highest node	10
3.3	Understanding change weight	10
3.4	Project depth	11
3.5	Text recognition	11
4	Mining teamwork from event logs	12
5	Summary and future work	14
	References	14

1 Introduction

The Deliverable 3.2 of the SHAPE project¹ reports work performed under Task 3.2 *Gather resource information about available resources from enterprise repositories of Work Package 3 Automatic Discovery of Processes and Resources and Constraints from Unstructured Data*. Its main aim is to further progress on mining document repositories by identifying the next steps that must be done towards a novel process discovery technique that makes use of unstructured and structured data. Concomitant to this deliverable is the milestone **(M2)** *Mining processes, resource consumption and witnesses for task completion from logs*.

The previous deliverable D3.1 Cabanillas et al. [2015] provided an overview on text mining and process mining methodologies from a general perspective as well as a mining approach for a specific type of processes, namely, *project-oriented processes*, also known as *project mining*. This deliverable delves into the applicability of such methods on real project data and studies potential extensions to increase their application scope. In particular, the deliverable explains how:

- The project mining technique introduced in D3.1 has been implemented and integrated into the Business Process Management Systems (BPMSs) Camunda².
- The project mining approach can be improved by annotating the information obtained from the Version Control Systems (VCSs) logs according to several criteria in order to increase the precision of the description of the activities mined. For example, text mining techniques can provide insights from the messages attached to SVN commits.
- The process mining technique described in D3.1 can be extended to take into account not only process activities performed by one single human resource but also collaborative activities performed by a team.

The rest of the deliverable is organized as follows. Section 2 describes the implementation of project mining as a library that can be run in the Camunda BPMS engine. Section 3 gives details on the future directions for enriching and improving our project mining technique. Section 4 outlines our idea to mine team compositions from event logs for collaborative work in processes. Finally, Section 5 closes this deliverable and points out the next steps to be carried out.

¹ <https://ai.wu.ac.at/shape-project/>

² <http://camunda.org/>

2 Implementation of a project mining technique in Camunda

In D3.1 [Cabanillas et al. \[2015\]](#) we introduced a project mining technique that works as follows.

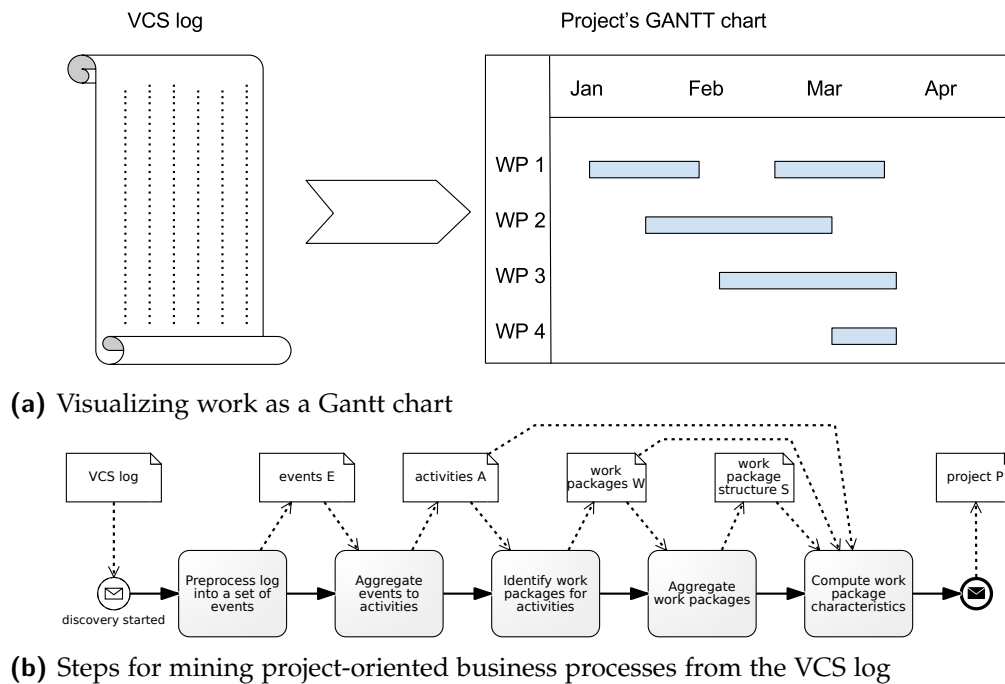
- First, it takes as an input a VCS log from the projects repository.
- Then it extracts the events from the commits. At the same time, it builds project's tree structure from the paths that are present in the log.
- The next step associates the events to each of the objects in the tree structure.
- Afterwards the events are aggregated into activities by using a threshold based technique described in [Baier et al. \[2014\]](#). Furthermore, activities starting time is estimated by taking into account the commit frequency of the contained events
- The obtained activities are further aggregated or collapsed to respond users' need for a coarser or finer grained view on the work, respectively.
- As a final step, a coverage indicator is computed that shows how efficiently the time available for a task is used by actual work.

Figure 1 clarifies the above description by showing in *a)* Figure 1a, the input and the output of project mining; and *b)* Figure 1b, the approach to obtain a Gantt chart from a VCS log, as a five-step process.

The approach has been implemented in Camunda, an open-source, extensible BPMSs. The software has been integrated into the architecture that is presented in Deliverable D6.1 [Ionescu \[2015\]](#), which provides further details about the general problem and a conceptual solution. Specifically, D6.1 presents an architecture based on three types of processes as depicted in Figure 2:

1. The first process is the main business process. Users execute activities as long as the process runs and generate artifacts. Artifacts are stored in a VCS.
2. The second process is the adaptation process. This process is able to perform changes in the main process to adapt to the evolution of the organization.
3. The third process is the mining process. As long as the main process runs, the mining process analyzes the artifacts that are being produced and eventually comes out with new insights on the running main business process. When a new insight is ready and shows a potential or existing risk, the adaptation process takes over to correct the behavior of the main process.

As said, the mining process makes use of external data, such as the VCS logs and possible additional sources of data. These sources are denoted as KB in the

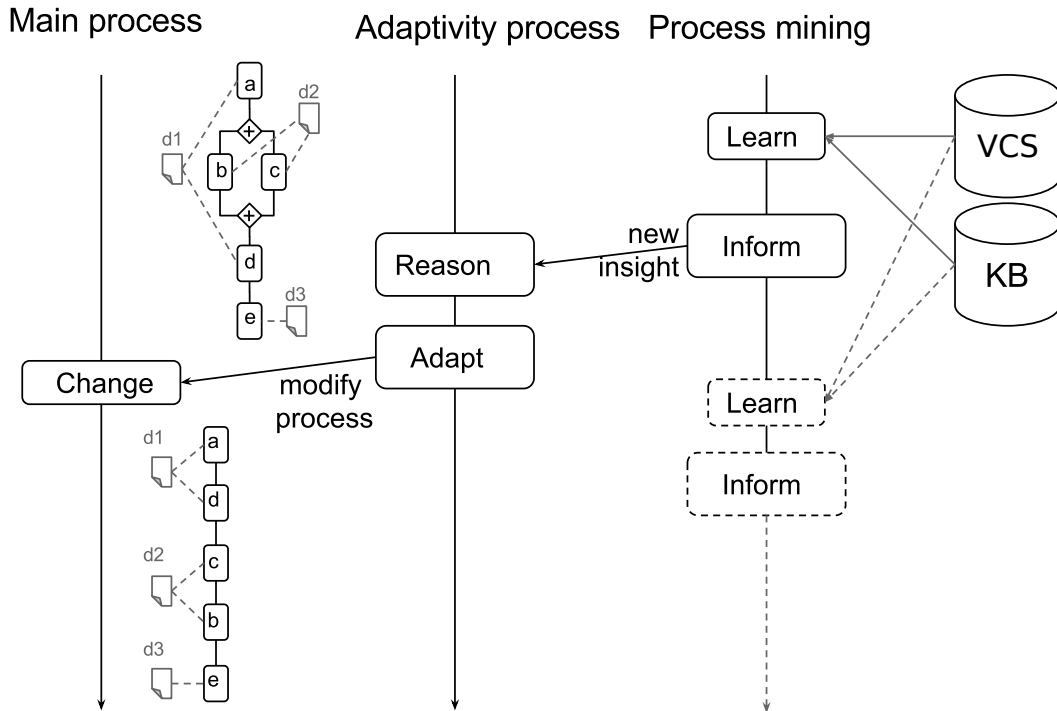


■ **Figure 1** Illustration of the project mining approach

67 picture and they may describe rules, regulations or behavioral patterns that the
 68 process must comply with. At the current state, the process mining task runs
 69 the project mining algorithm introduced in D3.1. In this architecture, the project
 70 mining technique helps to get useful insights on project-oriented business pro-
 71 cesses. To this end, it analyzes the artifacts that are produced during the lifetime
 72 of the project. The procedure has two main steps: *i) Learn* - runs an instance
 73 of the miner, and *ii) Inform* - communicates a new insight on the process to the
 74 adaptivity process.

75 Camunda is able to run services that are endpoints in a server. However, this
 76 has the disadvantage that we cannot make use of the visualization part which
 77 was implemented as a graphical user interface (GUI) component in Java. Such
 78 component allowed for user interaction with the elements of the directory (see
 79 [Cabanillas et al. \[2015\]](#)).

80 To achieve our goal of running the entire project mining approach in Camunda,
 81 we have separated the application logic part from the visualization part. Regard-
 82 ing the latter, we had the choose whether to adopt a representation of a Gantt
 83 chart that is exportable in a Web exchange format, or to make our component a
 84 library that can be called inside Camunda by a direct invocation. The first option
 85 would allow an external web application to interpret the encoded information and
 86 draw a Gantt chart accordingly. However, we opted for the second option because



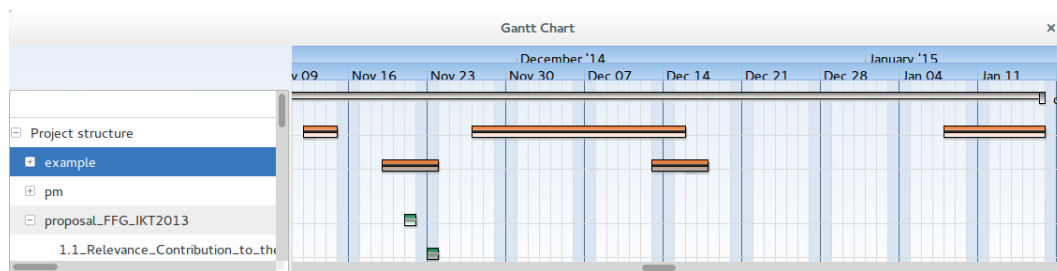
■ **Figure 2** Three process architecture from Deliverable D6.1

87 it does not require external services, thus allowing for a more integrated software
 88 environment. The choice was also guided by the design of the architecture shown
 89 in Figure 2 and discussed in more detail in deliverable Ionescu [2015]. It makes
 90 use of internal processes rather than calling them as external services. Therefore,
 91 we wrapped our project mining technique into a service. Nonetheless, the first
 92 option is still achievable with little modification.

93 As described above, the project mining technique takes as input a VCS log and
 94 a project-dependent threshold, and returns a representation of the Gantt chart. In
 95 order to still be able to exploit the aggregation functionalities available with the
 96 GUI-based implementation (see Figure 3), we compute all the sub-events aggrega-
 97 tions for each node of the project tree and use an *isAggregated* flag to indicate
 98 whether a node is expanded or collapsed by the user. All the information is stored
 99 as a payload on each tree node. This makes it possible to avoid recomputing on
 100 user's end. Granularity scaling can be done by hiding the events and showing
 101 only the activities of a user-chosen node when a user decides to aggregate on the
 102 node, and vice-versa when the user decides to collapse the subtree of the node.

103 A node representation in Java is an object with the attributes shown in Listing 1.

```
104     private String value;
105     private boolean isAggregated;
```

■ **Figure 3** The Gantt chart visualization from the GUI-based version of the project mining technique. This picture shows a part of the project tree structure on the left part and the corresponding activities or events on the right. The nodes *example* and *pm* are collapsed and therefore activities (*aggregation* step) are shown for them. On the contrary, the node *proposal_FFG_IKT2013* is expanded and fine granularity events are shown for it and its children.

```

106     private Activity activity;
107     private List<Event> eventList;
108     private Node parent;
109     private List<Node> childList;

```

■ **Listing 1** Java object that models an event

110 An event e is implemented as a Java object whose attributes are shown in
 111 Listing 2.

```

112     String id;
113     DateTime start;
114     DateTime end;
115     String fileID; //where did it happen
116     String type; // whether it's a Modify, Delete or Add
117     String author;
118     String commitID;
119     String comment;

```

■ **Listing 2** Java object that models an event

```

120     [[e1 e2 e3] [e4 e5] [e7 e8 e9]]

```

■ **Listing 3** An example of a project structure representation through a tree with events

121 Listing 3 illustrates the data structure that has been chosen for an *activity* ob-
 122 ject, specifically, a list of lists. The one in the listing is composed of 9 events
 123 that are grouped into 3 different chunks. This representation allows to express
 124 which events are aggregated from the aggregation step of the project mining algo-
 125 rithm. At the same time, such representation is easily exportable in formats like

6 Public Document

126 JavaScript Object Notation (JSON) or XML³.

127 The program outputs a Gantt chart as a Plain Old Java Object (POJO). A POJO
128 has interesting properties such as *a*) it can be persisted into a database (e.g. using
129 Hibernate⁴ or Java persistence API⁵) and *b*) is also translatable into other formats,
130 such as JSON⁶. In this case, we output a POJO that makes use only of object and
131 list entities .

```
132 +-[0] MiningSVN
133   +-[1] lib
134     +-[2] license4j-1.4.0.jar   A D
135   +-[1] src
136     +-[2] db
137       +-[3] jdbc
138         +-[4] TestMySqlConnection.java   A M D
139         +-[4] ConnectionFactory.java   A M
140     +-[2] gui
141       +-[3] nebula
142         +-[4] TreeConnectorExample.java   A D
143         +-[4] examles
144           +-[5] GanttSectionExample.java   A
145 +-[0] README.md   A M M M
```

■ **Listing 4** An example of a project structure representation through a tree with events

146 Listing 4 shows a part of the output of the program with its own Git change log
147 (MiningSVN.log) as input file. This is a simplified printout of the real output tree,
148 and shows only the type of the events that have occurred for each node, leaving
149 out the details on events and activities. The letters A, D and M stand, respectively,
150 for Added, Deleted and Modified. For instance, the file README.md has been added
151 once and modified three times.

152 This version of our technique can also be run as a standalone component from
153 the command line. An example call that takes as input the file MiningSVN.log and
154 uses a threshold of 7 days to aggregate is encoded as shown in Listing 5.

```
155 java -cp project-mining.jar main.ProjectMiner -f MiningSVN.
156 log -t 7
```

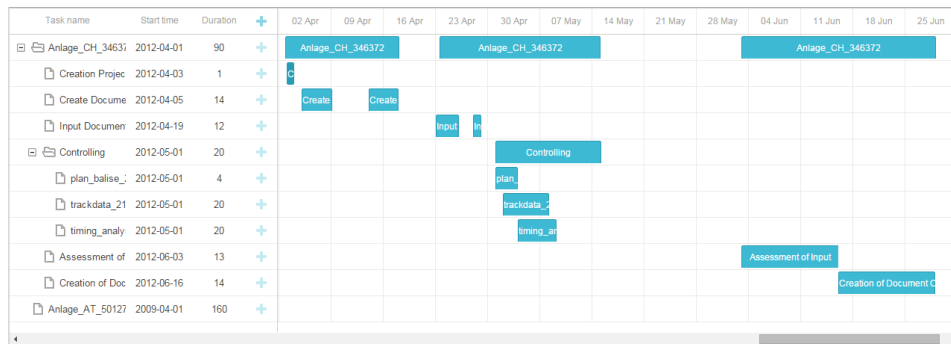
■ **Listing 5** Command line call to the project mining algorithm

³ <http://www.w3.org/XML/>

⁴ <http://hibernate.org/>

⁵ <http://www.oracle.com/technetwork/java/javaee/tech/persistence-jsp-140049.html>

⁶ <http://www.json.org/>



■ **Figure 4** Visualization of the Gantt chart in a browser

157 The library was included in the context of the deliverable D6.1 Ionescu [2015].
 158 Figure 4 shows a Javascript implementation that uses the dhtmlxGantt⁷ library to
 159 display the results of the project miner in a web browser.

160 The tool is open source and can be tried under the link <https://github.com/s41m1r/MiningCVS/tree/application-logic/MiningSVN>.
 161

162 **3 Project mining with annotations**

163 This section outlines further developments on the approach for mining project-
 164 oriented business processes presented in D3.1 (Cabanillas et al. [2015]) and pub-
 165 lished in Bala et al. [2015]. As already outlined as future work in the previous
 166 deliverable, we aim at improving our mining technique in order to fully exploit
 167 the data that we can obtain from the software repository logs. We have identified
 168 7 directions to get further insights on the process being discovered, which serve
 169 for annotating and classifying the work into activities in a more precise way. The
 170 following subsections describe each of the 7 points we are currently focusing on
 171 for current and future development.

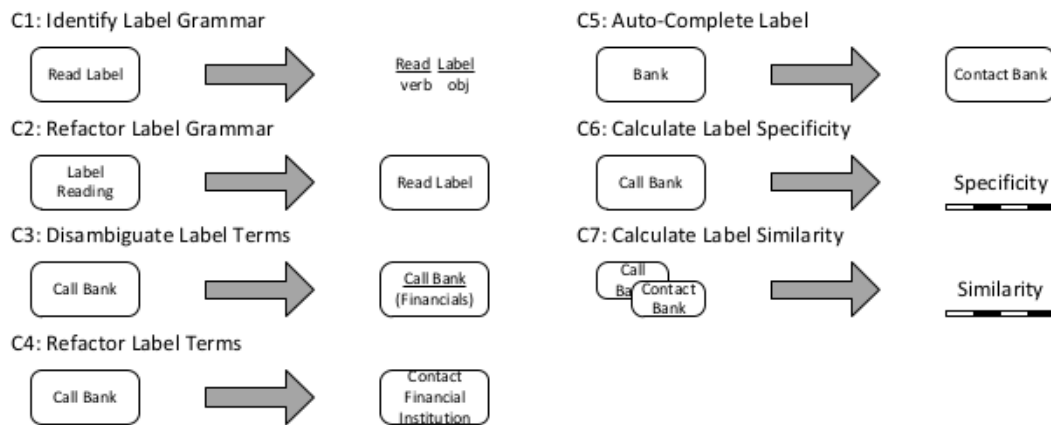
172 **3.1 Activity labels**

173 Following up on mining project repositories, the first thing we need to investigate
 174 is how to identify business process activities from the data available in Subver-
 175 sion (SVN). The comments entered by users in the commits can help towards this
 176 purpose, but still, there is the need to understand to which business process activ-
 177 ity a text refers. Difficulties grow when considering real-world cases where often
 178 users do not follow any given pattern when it comes to documenting their work

⁷ <http://dhtmlx.com/docs/products/dhtmlxGantt/>

179 progress. Even though guidelines are available such as the ones in [Mendling et al.](#)
 180 [\[2010\]](#), companies often use their own naming conventions. Therefore, a first chal-
 181 lenge that we would need to face in SHAPE is the *activity identification* from the
 182 natural language text of the comments. For this purpose, entity-recognition algo-
 183 rithms ([Etzioni et al. \[2005\]](#), [Ritter et al. \[2011\]](#)) can be firstly used in order to extract
 184 only the interesting part of the sentence. Afterwards, the extracted sentences can
 185 be refined into verb-object labels by applying automatic refactoring from [Leopold](#)
 186 [et al. \[2010\]](#).

187 The work of [Mendling et al. \[2015\]](#) raises the attention for the semantic aspects
 188 of text labels and natural language descriptions. The authors identify 25 chal-
 189 lenges for semantic process modeling. Figure 5 shows seven of the 25 challenges
 190 that are specifically related to the semantic of activity labels.



■ **Figure 5** Challenges related to the semantic of labels. Picture from [Mendling et al. \[2015\]](#)

191 3.2 Resource and resources type

192 We can mine resources and resource types out of document repositories. Given a
 193 VCS log, resources are present in the commit events. However their role remains
 194 unknown, as log files do not store any additional information about the resource
 195 besides the ID, which often consists of merely an email address. Therefore, we
 196 want to tackle the question: *how can we understand resource types?* Roles can be
 197 classified into two different types

- 198 ■ *Organizational roles.* For instance, Software Analyst or Project Manager. Such
 199 roles are typically defined in an organizational model describing the "static"
 200 structure of the company.

201 ■ *Functional roles*. These roles are more dynamic than the aforementioned roles,
202 since they may change for different projects, e.g., a writer or a reviewer when
203 preparing a project deliverable.

204 Organizational roles have received much interest from literature (Leitner et al.
205 [2013], Baumgrass and Strembeck [2013] and process mining of the organizational
206 perspective techniques have been implemented (Zhao and Zhao [2014], Song and
207 van der Aalst [2008]). In recent work Schönig et al. [2015] present a novel approach
208 that uses declarative mining and is able to deal with agile processes.

209 Efforts have also been made towards mining the functional roles of resources
210 from software repositories. Anvik and Murphy [2007] compare two approaches
211 that make use, respectively, of the source repository logs and of carbon-copy lists
212 (CC:), comments, and bug resolver information from reports in bug networks. The
213 evaluation was carried out by project experts from the Eclipse Platform. Results
214 show an improvement with respect to a previous approach by Anvik et al. [2006]
215 that took into consideration only the software repository files. The work of Yu
216 and Ramaswamy [2007] takes into account the interaction between users of a VCS
217 in order to group them into clusters. This allows distinguishing, for example,
218 between *core* members, who interact often; and *associate* members, who interact
219 more rarely. Tests on ORAC-DR⁸ and MediaWiki⁹ showed that this approach has a
220 predictive accuracy (percentage of correctly classified samples) of at least 94% and
221 a coverage (percentage of a developer roles that can be predicted by the rule) of at
222 least 80%.

223 3.2.1 Roles of resources based on file

224 VCS repositories can also be used to identify roles. There is typically a connection
225 between the type of the produced artifact and the role of the user who creates or
226 accesses it. Policies, handbooks and guidelines can be further used to gather more
227 insights on which resource can or may produce which artifact. For instance, the
228 phrase “*The project manager is in charge of writing the meeting minutes*” gives infor-
229 mation on both the artifact (i.e. `minutes.doc`, `minutes.txt`, etc) and the role of
230 the resource who produces it. Furthermore, by knowing a-priori which resource
231 can produce what artifact, it is possible to understand whether a role violation is
232 going on (i.e. some resource is working on a task (s)he was not assigned to).

⁸ <http://www.oracdr.org/oracdr>

⁹ <https://www.mediawiki.org/wiki/MediaWiki>

233 3.2.2 Annotating which is the highest node

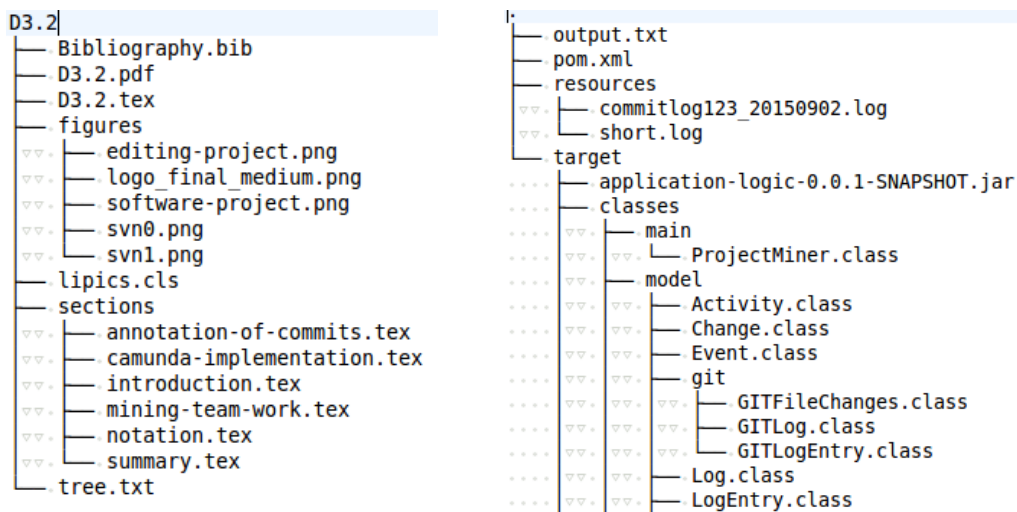
234 Another annotation challenge is to map each resource to the right hierarchy level
235 (s)he belongs to. This data can be easily understood from the VCS, but requires
236 the project structure to respect the real access hierarchy according to the resource
237 role. In this case we are able to discriminate between resources that operate on a
238 higher level on the project tree and those who operate in depth. Such separations
239 give us more information about the role of the resource. In fact, more resources
240 such as managers or project leaders operate on the top of the tree hierarchy and
241 they have more access and more visibility on the entire project. On the contrary
242 resources such as programmers usually commit their changes in depth and have
243 less access to the higher levels of the project hierarchy.

244 3.3 Understanding change weight

245 An important issue in project discovery is to understand the type of change events
246 that the project undergoes. The challenge that we want to address in future work
247 can be formulated as the question *how big is the change?*. This translates to under-
248 standing what type of events are actually happening in the project repository. It
249 is possible to gather more insights into the project evolution if we look not only
250 in which part of the project directory the changes happened, but also how much
251 modification occurred in single files. This can be done by measuring the effec-
252 tive change of the files themselves. Many VCSs such as Git¹⁰ or SVN¹¹ have *diff*
253 commands incorporated that allow the user to see the differences between two
254 different versions of a document. We plan to use such information for retrieving
255 the file change, measured in terms of Lines of Code (LoC) or megabytes. Further-
256 more LoC is also a complexity metric, which in turn may lead to more changes.
257 In this sense, complexity can also be related to tasks/effort. [Rothenberger and](#)
258 [Hershauer \[1999\]](#) argue that LoC can be an appropriate measure for development
259 effort. Especially when: i) the subject domain does not change; ii) the code is
260 written in the same programming language; and iii) the software development
261 company follows a rigid methodology. This is normally the case with SHAPE,
262 where engineers follow a handbook and have a defined working methodology.
263 Nevertheless, approaches to achieve better insights that go beyond LoC measures
264 have also been addressed by [Canfora et al. \[2007\]](#), who are able to infer changes
265 between two document versions even when the LoC count does not change.

¹⁰ <https://git-scm.com/>

¹¹ <https://subversion.apache.org/>



(a) An editing project structure layout

(b) An excerpt from a software development project structure

■ **Figure 6** The structure of a project reflects also its type

3.4 Project depth

266


267 The project tree gives further information on the resources as well as on the type
 268 of project. For example, software development projects tend to present a very
 269 deep tree structure. Figure 6 shows the LaTeX layout of writing this manuscript
 270 and an excerpt from the project mining software, respectively on the left side
 271 and on the right side. It is easy to see how the structure is almost linear in the
 272 writing project (Figure 6a) as the maximum tree depth is 2, and how it drastically
 273 changes including deep-nested files up to 5 levels of depth, when it comes to a
 274 software project (Figure 6b). We can exploit the information that comes from the
 275 project depth to characterize certain types of work that strongly affect the tree
 276 depth. We plan to combine this information with other sources of evidence, such
 277 as filenames, for discovering the type or work.

3.5 Text recognition

278

279 It is possible to gather information on the meaning of the various parts of the
 280 project by also looking at common best practices that users adopt when using a
 281 VCS. For instance, SVN best practices¹² recommend the idea of a project root
 282 folder. Figure 7 illustrates that. A "project root" contains exactly three subdirec-
 283 tories: /trunk, /branches, and /tags. A repository may contain one or more

¹² <https://svn.apache.org/repos/asf/subversion/trunk/doc/user/svn-best-practices.html>



■ **Figure 7** Recommended repository layout for SVN projects. The trunk, tags, and branches trio is sometimes referred to as “the TTB directories”.

284 project roots. The *trunk* is the directory under which the main project develop-
 285 ment occurs; *branches* is the directory in which various named branches of the
 286 main development line can be created; and *tags* is a collection of tree snapshots
 287 that are created, and perhaps destroyed, but never changed. Typically, a change in
 288 the repository reflects a single purpose: the fixing of a specific bug, the addition of
 289 a new feature, or any other particular task with a specified goal. Furthermore, it is
 290 usual for several projects to use clear guidelines and handbooks. Such handbooks
 291 suggest, for instance, how names should be used for the different artifacts that are
 292 produced within the project. Using this information, we can associate a meaning
 293 to the changes that are made to the project. This may lead to an improved activity
 294 recognition.

295 4 Mining teamwork from event logs

296 In D3.1 [Cabanillas et al. \[2015\]](#) we introduced an approach to mine the organiza-
 297 tional perspective of business processes to represent resource-aware process mod-
 298 els from the information stored in event logs [Schönig et al. \[2015\]](#). In short, the
 299 approach discovered resource assignment rules describing (i) the conditions that
 300 the resources of the organizations must meet in order to participate in the process
 301 activities, and (ii) resource-related conditions that constrain the execution order
 302 of some activities, e.g., some types of resources may be forced to execute activi-
 303 ties in a specific order, whereas other types of resources are not subject to such
 304 restrictions. These rules were extracted by confronting the information from past
 305 executions and the organizational model of the company (i.e., the positions, roles,
 306 organizational units, etcetera). That work was afterwards extended to include a
 307 post-processing step that simplified the generated process model by pruning out
 308 unnecessary, redundant resource assignment rules [Schönig et al. \[2015\]](#).

309 The aforementioned approach focused on activities that were performed by
310 one single resource, i.e., individual work. We are currently working on extending
311 the technique to also take into consideration activities that are collaboratively per-
312 formed, i.e., which involve the participation of several resources. The goal is to
313 mine the characteristics of the team members in such collaborative activities, e.g.,
314 to discover a team composed of a resource with a certain role R , a resource with
315 expertise on a specific software application S and a resource from a certain organ-
316 isational unit U . Similarly to the previous approach, the framework to discover
317 the characteristics of the team members is constituted by the workflow resource
318 patterns [Russell et al. \[2005\]](#), specifically the following creation patterns:

- 319 ■ *Direct Distribution*, i.e., the ability to specify at design time the identity of the
320 resource that will execute a task.
- 321 ■ *Role-Based Distribution*, i.e., the ability to specify at design time that a task can
322 only be executed by resources that correspond to a given role.
- 323 ■ *Organisational Distribution*, i.e., the ability to offer or allocate activity instances
324 to resources based on their organisational position or organisational units and
325 on their relationship with other resources.
- 326 ■ *Capability-Based Distribution*, i.e., the ability to offer or allocate instances of an
327 activity to resources based on their specific capabilities.

328 Following the same approach as for individual work mining, the mining of
329 team compositions will consists of three steps:

- 330 1. A pre-processing phase will filter the teams found in the event log to later
331 process only those that are frequent, i.e., the combination of team members that
332 occur at least a certain amount of times according to a established threshold.
333 The rest of member combinations will be considered exceptions and will not
334 be taking into account for the characterization of the teams that are allowed to
335 execute a certain collaborative activity.
- 336 2. The aforementioned creation patterns will be checked within the different
337 teams for each activity in order to find out which rules apply for the team
338 members, i.e., how the team allowed to execute an activity is characterized.
339 This characterization is not accurate, since it is only able to determine the
340 different creation patterns that occur within a team, but not how they are dis-
341 tributed among the team members. This implies that if we have, e.g., a team
342 composed of three members and we discover three rules *direct(activity, Peter)*,
343 *role(activity, Coordinator)* and *capability(activity, speaksEnglish)*, at this point
344 we do not know whether there must be one person with each of these charac-
345 teristics, or there must be several members meeting one or several conditions.

346 3. A two-step post-processing phase will increase the precision of the information
347 about the team composition by specifying the distribution of the discovered
348 characteristics among the team members in order to bridge the previous gap.

349 **5 Summary and future work**

350 This deliverable sets the field for future developments on extracting process knowl-
351 edge from unstructured and semi-structured data. We have identified and ex-
352 plained seven possible new insights that we can obtain from combining data from
353 repositories with semantic annotation from text. This document also contains a
354 novel mining technique for resource-perspective mining. This approach allows
355 for discovering teamwork. As last, we have also modified the original GUI based
356 application into a Java library that can be run inside Camunda.

357 For future work we aim at a publication on improving the project discovery
358 technique using user comments from commits. We plan on using the points that
359 we identified earlier in this deliverable as a starting base for the next dissemi-
360 nation. We are also working on the three-step approach for mining team com-
361 positions that we outlined in Section 4, which we plan to submit as a journal
362 publication in the following weeks.

363 **References**

- 364 John Anvik and Gail C Murphy. Determining implementation expertise from bug
365 reports. In *Mining Software Repositories, 2007. ICSE Workshops MSR'07. Fourth*
366 *International Workshop on*, pages 2–2. IEEE, 2007.
- 367 John Anvik, Lyndon Hiew, and Gail C Murphy. Who should fix this bug? In
368 *Proceedings of the 28th international conference on Software engineering*, pages 361–
369 370. ACM, 2006.
- 370 Thomas Baier, Jan Mendling, and Mathias Weske. Bridging abstraction layers
371 in process mining. *Information Systems*, 46:123 – 139, 2014. ISSN 0306-4379.
372 <http://dx.doi.org/10.1016/j.is.2014.04.004>. URL <http://www.sciencedirect.com/science/article/pii/S0306437914000714>.
373
- 374 Saimir Bala, Cristina Cabanillas, Jan Mendling, Andreas Rogge-Solti, and Axel
375 Polleres. Mining project-oriented business processes. In Hamid Reza Motahari-
376 Nezhad, Jan Recker, and Matthias Weidlich, editors, *Business Process Manage-*
377 *ment*, volume 9253 of *Lecture Notes in Computer Science*, pages 425–440. Springer
378 International Publishing, 2015. ISBN 978-3-319-23062-7. [10.1007/978-3-319-](https://doi.org/10.1007/978-3-319-23063-4_28)
379 [23063-4_28](https://doi.org/10.1007/978-3-319-23063-4_28).
- 380 Anne Baumgrass and Mark Strembeck. Bridging the gap between role mining and
381 role engineering via migration guides. *information security technical report*, 17(4):
382 148–172, 2013.
- 383 Cristina Cabanillas, Jan Mendling, Axel Polleres, and Saimir Bala. Requirements
384 for process, resource and compliance rules extraction from text. Technical Re-
385 port 1, 2015.
- 386 Gerardo Canfora, Luigi Cerulo, and Massimiliano Di Penta. Identifying changed
387 source code lines from version repositories. In *MSR*, volume 7, page 14, 2007.
- 388 Oren Etzioni, Michael Cafarella, Doug Downey, Ana-Maria Popescu, Tal Shaked,
389 Stephen Soderland, Daniel S Weld, and Alexander Yates. Unsupervised named-
390 entity extraction from the web: An experimental study. *Artificial Intelligence*, 165
391 (1):91–134, 2005.
- 392 Tudor Ionescu. Workflow-processing and verification for safety-critical engineer-
393 ing – conceptual architecture. Deliverable D6.1 (M2), 2015.
- 394 Maria Leitner, Anne Baumgrass, Sigrid Schefer-Wenzl, Stefanie Rinderle-Ma, and
395 Mark Strembeck. A case study on the suitability of process mining to produce
396 current-state rbac models. In *Business Process Management Workshops*, pages 719–
397 724. Springer, 2013.
- 398 Henrik Leopold, Sergey Smirnov, and Jan Mendling. Refactoring of process model
399 activity labels. In ChristinaJ. Hopfe, Yacine Rezgoui, Elisabeth MÃ@tais, Alun

- 400 Preece, and Haijiang Li, editors, *Natural Language Processing and Information Sys-*
401 *tems*, volume 6177 of *Lecture Notes in Computer Science*, pages 268–276. Springer
402 Berlin Heidelberg, 2010. ISBN 978-3-642-13880-5. [10.1007/978-3-642-13881-2_28](https://doi.org/10.1007/978-3-642-13881-2_28).
- 403 Jan Mendling, Hajo A Reijers, and Wil MP van der Aalst. Seven process modeling
404 guidelines (7pmg). *Information and Software Technology*, 52(2):127–136, 2010.
- 405 Jan Mendling, Henrik Leopold, and Fabian Pittke. 25 Challenges of semantic
406 process modeling. *IJISEBC*, 1(1):78–94, 2015.
- 407 Alan Ritter, Sam Clark, Mausam, and Oren Etzioni. Named entity recognition in
408 tweets: An experimental study. In *Proceedings of the Conference on Empirical Meth-*
409 *ods in Natural Language Processing, EMNLP '11*, pages 1524–1534, Stroudsburg,
410 PA, USA, 2011. Association for Computational Linguistics. ISBN 978-1-937284-
411 11-4. URL <http://dl.acm.org/citation.cfm?id=2145432.2145595>.
- 412 Marcus A Rothenberger and James C Hershauer. A software reuse measure: Mon-
413 itoring an enterprise-level model driven development process. *Information &*
414 *Management*, 35(5):283–293, 1999.
- 415 Nick Russell, Wil M. P. van der Aalst, Arthur H. M. ter Hofstede, and David
416 Edmond. Workflow Resource Patterns: Identification, Representation and Tool
417 Support. In *CAiSE*, pages 216–232, 2005.
- 418 Stefan Schönig, Cristina Cabanillas, Stefan Jablonski, and Jan Mendling. A Frame-
419 work for Efficiently Mining the Organisational Perspective of Business Pro-
420 cesses. *Decision Support Systems*, page Under review, 2015.
- 421 Stefan Schönig, Cristina Cabanillas, Stefan Jablonski, and Jan Mendling. Mining
422 the organisational perspective in agile business processes. In *Enterprise, Business-*
423 *Process and Information Systems Modeling*, pages 37–52. Springer, 2015.
- 424 Minseok Song and Wil MP van der Aalst. Towards comprehensive support for
425 organizational mining. *Decision Support Systems*, 46(1):300–317, 2008.
- 426 Ligu Yu and Srin Ramaswamy. Mining cvs repositories to understand open-
427 source project developer roles. In *Proceedings of the Fourth International Workshop*
428 *on Mining Software Repositories*, page 8. IEEE Computer Society, 2007.
- 429 Weidong Zhao and Xudong Zhao. Process mining from the organizational per-
430 spective. In *Foundations of Intelligent Systems*, pages 701–708. Springer, 2014.