MALEKC:

Machine Learning for Engineering Knowledge Capture

Författare: Dag Bergsjö
Datum: 2019-11-29
Projekt inom: Machine Learning

Fordonsstrategisk Forskning och Innovation

Contents

1 Summary

The purpose of this research is to enable a paradigm shift in the automotive industry, by expanding from an experienced-based development process towards a proactive information-centric product management system where machine learning is used, to predict decision points in development process. The focus of this research is on how to create knowledge from existing data sources and utilize that efficiently in the development process. The targeted solution is to utilize machine learning algorithms and smart assistants, to identify the right knowledge in the right time to the right individual, which in the context of this project is an engineer or service technician making an "uninformed" decision. Chalmers and Wingquist Laboratory (WQ) is the main applicant. Fraunhofer Chalmers Research Center for industrial mathematics (FCC) is the driving partner of the Machine Learning knowledge. And Rejmes Transportfordon AB and AB Volvo are the main industrial participants and primary users of the research result.

Product development companies are collecting data in form of Engineering Change Requests for logged design issues, tests and product iterations. These documents are rich in unstructured data (e.g., free text) and previous research has pointed out that product developers find current IT systems lacking capabilities to accurately retrieve relevant documents with unstructured data. In this research we demonstrate a method using Natural Language Processing (NLP) and document clustering algorithms in order to find structurally or contextually related documents from databases containing Engineering Change Request documents. The aim was to radically decrease the time needed to search for and then organise the search results by utilizing such NLP algorithms to effectively search for related engineering documents and further create labelled clusters out of these documents.

This project is a continuation of the project Propid that aimed to investigate the potential for Big Data analytics within a product development setting, but it also focused on identifying potential future needs. The project is further supported by two internal industrial PhD students at AB Volvo.

2 Sammanfattning på Svenska

Syftet med detta forskningsprojekt är att möjliggöra ett paradigmskifte inom fordonsindustrin genom att gå ifrån en utvecklingsprocess där personberoende erfarenhetsbaserad kunskap har stort genomslag (och ofta glöms bort) till en proaktiv datadriven produktutvecklingsprocess där även kunskap baserad på maskininlärning används för att förutsäga och undvika potentiella problem. Den potentiella lösningen är att utnyttja maskininlärningsalgoritmer och smarta assistenter för att identifiera rätt kunskap och sedan förse rätt person med rätt kunskap i rätt tid. Rätt person i detta samband är en ingenjör eller servicetekniker som tar ett beslut.
Chalmers och Wingquist Laboratory (WQ) är huvudsökare. Fraunhofer Chalmers Research Center (FCC) är drivkraften gällande maskinlärning. Medan Rejmes Transportfordon AB och AB Volvo är de huvudsakliga industriella deltagarna och primära användare av forskningsresultatet.

3 Background

This project focuses on a future proactive information centric product management system where machine learning is used, to predict decision points in development process and to provide underlying and situation adapted knowledge customized for those decision points.
The focus is hence, not how to improve machine learning algorithms or data collection methods, but how to create knowledge from existing data sources and utilize knowledge derived from machine learning algorithms by utilizing smart assistants efficiently in the development process.

Data driven product development (From Process to PLM)
The academic area of product lifecycle management (PLM) focuses on making sure that data and information created during one phase of the product lifecycle (e.g. development) can be used to support the work in another phase (e.g. manufacturing preparation or even manufacturing operation). In practice the concept of PLM has led mainly to focus on a feed forward flow of digital product representations (e.g. 3D models) from development to preparation of manufacturing or aftermarket systems for a specific product. There is however a vision in the academic visions of PLM to also feed data and information back to enable learning in the development organization by being able to connect cause-and-effect chains between product/process decisions and their outcomes in terms of product qualities or process qualities. In practice however this remains a vision because the data flows which are generated in the different lifecycle phases (production, usage, sales, aftermarket) are disconnected and feedback concerning poor product or process decisions reaches those who were involved with them only when serious quality problems are detected through the quality functions.

One of the major root-causes to the problem of disconnected data flows lies in the existing database oriented technologies that underpin the systems used to manage the data during development (PLM), supply chain (ERP) and manufacturing (MES) and later in aftermarket. The data and information models employed in each of these sources are usually adapted to the specific needs of the processes they and now also each product individual is generating data (according its own logic that is adapted to the process of operating the product itself). This means that there are challenges in merging the data flows by merging the databases because the underlying logics cannot be merged. Information about product changes, including structured data (e.g., numerical data, categorical data and timestamps) mixed with unstructured data (e.g., text inserted by engineers), is stored in databases to keep track of evolving design solutions, verification and decision documents ((Pikosz and Malmqvist 1998)). Large PD projects can contain tens of thousands of product change documents (Arnarsson, Gustavsson et al. 2017), also known as Engineering Change Requests (ECRs) (Jarratt, Eckert et al. 2011). The ECR process focuses on achieving a rapid and quality-assured process from the identification of issues until resolution. With the fast development of data collection and storage, companies are now faced with large volumes of complex data from multiple sources (Wu, Zhu et al. 2013). Around 80% of organizational data is in unstructured form (Grimes 2008) and that poses the challenge of how to extract meaning from this vast amount of text (Kobayashi, Mol et al. 2018). Arnarsson et al. (2017) have identified the need of product developers to perform advanced searches within databases for efficient problem investigation and resolution. The problems have been that i) search functionality is limited to mostly considering structured data and not taking into account the rich information stored in unstructured data (e.g., text data) and ii) current search results return up to hundreds of reports back even when filters are applied, such as timeframe, vehicle types and keywords. It follows that it is very time-consuming to identify the most relevant documents.

Knowledge creation and reuse (management)
From the perspective of a knowledge-oriented company that strives for continuous learning and improvement this is a serious problem. The problem exists in most traditional product developing and producing companies which develop and sell their products on a global scale and where the

product production, usage and maintenance is spread across large geographic distances and dispersed across many different actors.

A short and condensed format for knowledge reuse is essential to eliminate unnecessary activities and to ensure the applicability of the latest knowledge. The literature is consistent in stating that Engineering Checklists may be part of this condensed format; Kennedy et al. (2008) state that an Engineering Checklist is the principal Lean tool used by Toyota for knowledge reuse. Further, Morgan and Liker (2006) state that such checklists serve as reminders of the tasks that must be considered, decided and completed in a timely fashion and include design rules, standards, good and bad examples and other forms of empirical knowledge based on years of experience. The use of this type of format gets support also from the field of Cognitive Psychology which has proven that the human brain in decision-making mode cannot hold more than five-seven pieces of information simultaneously (Baddeley, Della Sala et al. 1996). An approach labeled "thin-slicing" (Ambady and Rosenthal 1992) used to populate checklists has been tested with promising results and greatly increases the transfer of knowledge from experts and improves decision-making of novices (Gladwell 2007). Thin-slicing as a knowledge capture methodology is well in line with checklists as tools with which to document and store knowledge that increases overall reusability.

However, simplified methods do not eliminate the need for constant supervision and instruction until the methodology becomes self-evident and universally accepted (Sendlhofer, Mosbacher et al. 2015). An engineering checklist that only tells the engineer what to do is easy to use, but is not reusable in providing a deeper understanding and facilitating the possibility of continuous improvements (Morgan and Liker 2006). Further, a common critique that needs to be refuted is that the engineering checklist needs to be kept brief and cannot consider all available knowledge on a certain subject (Kokkoniemi 2006), for example the solution provided by Catic and Malmqvist (2013) that the descriptions considering the "how's" and "why's" of the application of knowledge can be appropriate by excluded from the checklist and put into a reference document.

Machine Learning

Machine learning and data mining offer quantitative methods for performing data analysis with any system that is generating data. Once an overview of the data has been created, data can be identified and collected from the databases, and machine learning can be employed in order to learn the structure and extract insights from the historical data that companies are sitting on. The traditional way of analysis has been to try to find answers in data with manual exploration by, for example, users exporting data to a spread sheet software tool and exploring it according to their own methods. With the rapid development of easily accessible machine learning methods, there are opportunities to make these explorations more effective and precise by automating the work. Machine learning can help shed a light on a variety of complex issues that would not be obtainable through a "manual" inspection and analysis.

In the engineering design domain, early studies related to data analytics and machine learning include (Kuffner and Ullman 1991) who chartered out design engineers' needs for more design information than only standard design documents when developing complex products, and Reich (1997) who proposed a seven-step process for developing machine learning tools supporting civil engineering tasks. (Arnarsson, Malmqvist et al. 2016) showed how machine learning and visualization tools could be applied to exploring a database consisting of ECRs from a complex truck development project. The study demonstrated a process for compiling and cleaning the data along with methods for natural language processing, data visualization and exploration, and for pattern identification and analysis.

In this project, the main objective for using machine learning techniques is to aid the product developers in the development process by condensing and visualizing relevant data in order for them to be able to make a decision. The types of machine learning methods that will be used within the project are predictive models (e.g., for predicting how it will take before an issue is resolved), natural language processing methods (e.g., for parsing and understanding the input text written by the engineers), and Markov chain models (e.g., for analysing how deviation reports gets transferred between different report statuses). Furthermore, to identify structurally related information (e.g., all ECRs related to a particular part) and contextually related documents (e.g., ECRs dealing with a certain problem), dealing with similar issues or problems based on structured and unstructured data, a short list of the most relevant documents is produced. Recent developments in Natural Language Processing (NLP) have made it possible to search through huge databases of text documents to retrieve relevant items based on variety of search queries. By utilizing NLP methods such as search engine methodology or document embedding approaches, one can find documents (in this case ECRs) that are similar to each other by looking at the context of sequences of words together, instead of words in isolation. Previously, simpler methods relied on the notion that related documents share the exact same words and to find similar documents, one usually searches through the document set to find the ones that match best according to word count. More elaborate methods, such as doc2vec-methodology (Le and Mikolov 2014), utilizes the context of a document and automatically detects synonyms to words used in the search query.

Another benefit of these types of NLP methods is that they provide a similarity metric between documents, meaning that one can get a numeric value quantifying the degree of similarity between pairs of documents. This metric can be utilized for performing cluster analyses (Steinbach, Karypis et al. 2000) where one tries to group together related documents into clusters.

The main background for this project is based on the following three projects:

- *2015-06912,* PROPID Produktionsutveckling baserad på produktindividdata, FFI - Big Automotive Data Analytics (BADA)
- 2012-02175, Visualisering och IT i Produkt- och Produktionsutveckling (Vis-IT), FFI: Hållbar Produktion
- 2015-04509, Digilean, Produktion 2030

4 Approach and Goals

The goals for the project as stated in the project application was:

To create conditions to effectively identify and reuse knowledge by utilizing machine-learning algorithms on engineering change reports (ECR)

This goal has been central for the work performed and presented in this report. The goal has been further decomposed to also include ways of added validation and knowledge reuse using engineering Checksheets. We have divided the presentation into 4 distinct case studies towards the subject matter which all show promising results towards the automation of knowledge capture and validation.

This project aims to investigate and demonstrate machine learning to improve the business perspective of data driven product development. In such the project does not primarily aim to expand the knowledge of ML analytics and algorithms but rather how ML can be used to improve business processes, especially the learning process of automotive companies.

The idea is to prepare for the development of an intelligent assistant that utilizes machine learning algorithms to analyse and introduce relevant information from product deviation reports into engineering work processes. The data of the deviation reports is currently unstructured and contains numeric, categorical, and textual data. ML algorithms based on natural language processing (NLP) can interpret the unstructured data and construct hypotheses and form interview questions based on patterns from the database. These answers can be validated or improved by using, e.g., e-mail, this ues-case is depicted in figure 1. The assistant will then help the existing structure, within AB Volvo, of Knowledge Owners and users to update knowledge in repositories, based on thin slicing methodology known as Checksheets, a research area previously researched within the FFI financed project Vis-IT, (2012-02175) and later included into the knowledge management process and methodology at AB Volvo in 2015.

The project consisted of four work packages. One that aimed to collect data from the different sources and fine tune the analysis. One that aimed to demonstrate how the output from the data analysis was used in a specific case to produce useful knowledge. One that aimed to demonstrate how the output from the data analysis could be used to optimize remanufacturing/reuse decisions for specific product individuals, and finally there was a project management and dissemination work package that was used to coordinate the other work packages and communication concerning results.

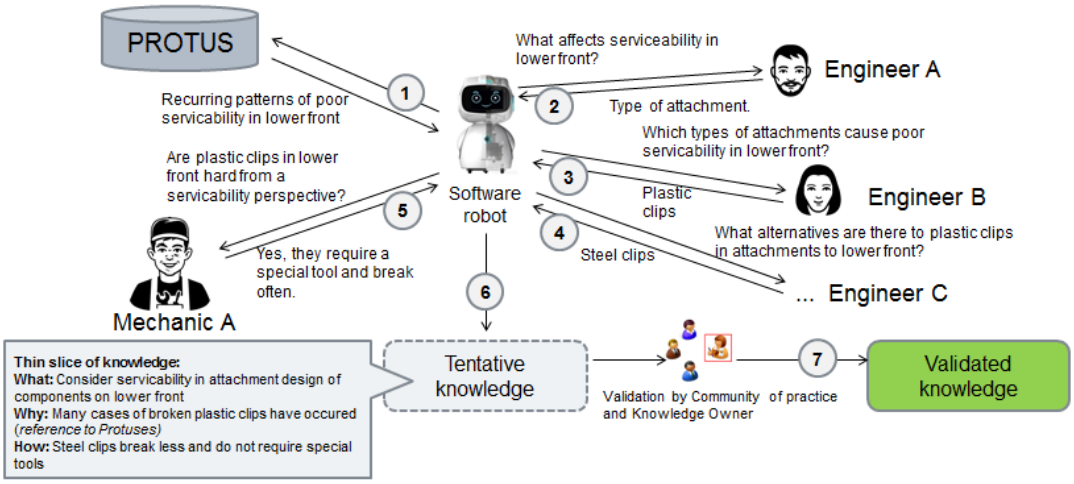

Figure 1Example of a use case for the Intelligent assistant further described in WP 3, Protus is the name of the database holding ECR at Volvo

5 Results

In this chapter, the results from the four different case studies performed are briefly presented. Case 1 and Case 2 relate manly to the Machine Learning algorithms developed end tested to generate (or find) knowledge, wheras case 3 mostly refer to the current practice of creating and reusing knowledge at the case companies. Case 4 presents experiments to validate knowledge by using an automated approach.

FCC has led the development of the demonstrator for cases 1 and 2 and the use cases have been refined from research at Volvo, with some support of Chalmers. Case 3 and 4 has mainly been led by Chalmers and have been tested and evaluated in both industrial and academic environments at Volvo. The cases 3 and 4 have further been supported by Rejmes Transportfordon.

5.1 Case I: Modelling industrial engineering change processes using the Design Structure Matrix for sequence analysis

The study is based on a product development database containing ECRs from an industrial commercial vehicle development project (see Figure 3). Compared to manually analysing a few ECRs at a time, the Markov chain DSM approach allows the analysis of large sets of ECRs, thus providing new, data-driven insights. Transitions of ECR states are mostly random and it is therefore suitable to use Markov chain as it can model randomly changing systems where it assumes that the future state only depends on the current state.

A Markov chain is a type of stochastic process where the transition probabilities between the available states fulfil the Markov property, which means that the probability of evolving from one state to another only depends on the current state. This implies that the process is, in a sense, memoryless and does not care about the history of the process. For the given case, when we consider an ECR as a discrete stochastic process, it means that when an ECR is in a specific state at a specific time point, the probability of being transferred from that state to another only depends on the state the ECR is in at that time point. This is because we assume that the Markov property (i.e., memoryless property) holds. The employed process for identifiying an cleaning data for analysis is depicted in Figure 2.

There are 15 ECR process states used in the Markov chain DSM that are organized into three categories. Each state has a numerical name and description where the first numerical digit indicates which states belong to the same category. One example of these chains is visualised in Figure 4 (For full detail refer to the paper described in chapter 6.2).

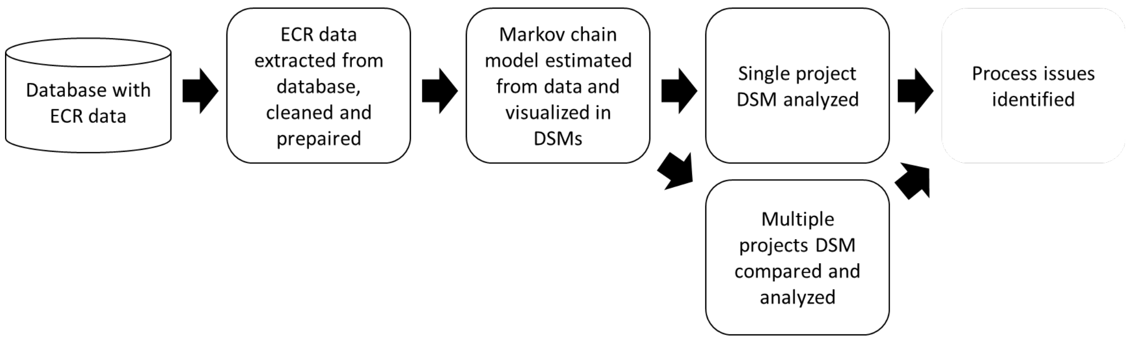

Figure 2Flowchart of the methodology.

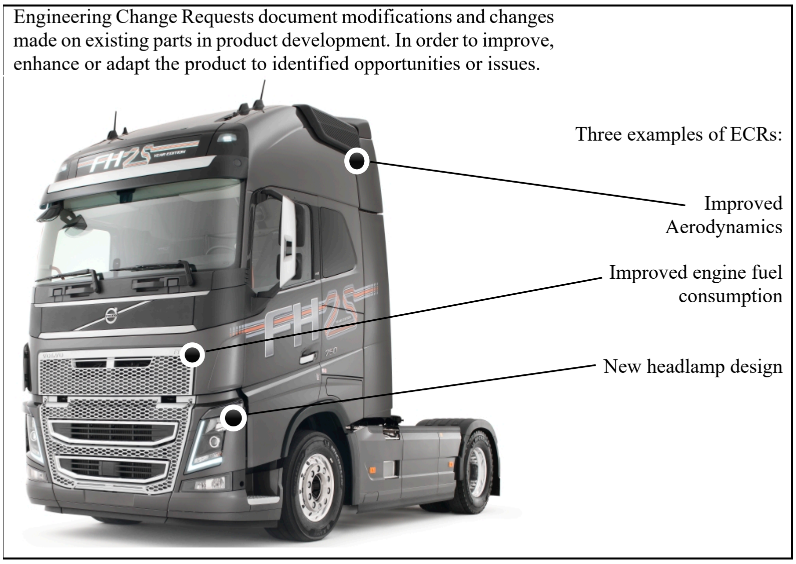

Figure 3Example of an industrial commercial vehicle that generates ECRs during product development with examples of ECRs and their location on the vehicle

Comparison of projects
Table 1 below shows the interaction of projects and their pattern for benchmarking where we observe the following:

- Projects have between two to three patterns each.
- P3 and P4 are unique for specific projects.
- P2 is common for all.
- P1 is common for all, except project C.

Figure 4 presents the average DSM for all the four projects in this study. Notably, only patterns P1 and P2 are visible in the figure. However, the probability of state 10 to sequence to states 19 and 3 is almost the same since pattern P1 should be the dominant sequence, and P4 is a process deviation. Pattern P4 appears, therefore, to be isolated to Project C but on the edge of taking place in other projects. Meanwhile, pattern P3 is not visible in the average DSM and seems to be isolated to Project B since the probability of sequences in the iteration is so small that they fade out in a larger dataset.

Table 1. Interaction of projects and patterns.

		Pattern			
		P1	**P2**	**P3**	**P4**
Project	**A**	X	X		
	B	X	X	X	
	C		X		X
	D	X	X		

All projects	1	10	11	13	15	19	2	21	22	3	31	35	36	37	39	#
1 ECR created	P	99.7										P				7619
10 ECR distributed						33.8				33.4	10.0	2.4				7739
11 ECR on hold		3.5			3.5					60.7	21.4					28
13 ECR from external project						11.7	5.8			11.7	23.5	11.7				17
15 ECR without solving responsible		9.2				23.1	2.5			24.7	8.2	7.2				194
19 ECR with solving responsible										53.1	19.5	2.7			2.8	2914
2 ECR re-issued						5.0		2.8		25.7	12.1	5.0			4.2	140
21 ECR incomplete		6.9				5.0	9.0			29.8	7.6	4.0			2.1	275
22 ECR not approved										27.2	16.8	8.1	2.1		5.7	851
3 Identification of solution											45.3	9.1			6.8	5157
31 Assessment of solution										4.0		8.9	3.6	5.1	36.5	4539
35 Decision on assessed solution									3.9	12.6	17.9		4.7		5.0	1335
36 Solution approved										4.0	7.0	2.2		2.2	32.5	440
37 Testing solution										2.9	5.8	3.1			16.0	441
39 Solution ready											2.5					2665

Figure 4 Sequence visualized for all the four projects in one DSM with 7,619 ECRs.

Conclusions regarding Case I

This study clearly shows that modelling ECR data with the Markov chain DSM can be done, and that the results are useful for analysing the sequences of ECRs. We used four different project DSMs that, in turn, identified four patterns as key results. First, the DSM identified the most common sequence for states within an ECR. Second, there is a tendency to skip state 3 (i.e., identification of solution) and sequence straight to state 31 (i.e., assessment of solution), which is a process deviation. Third, it is possible to identify iterations occurring in the sequence process, which means that we are able to see loops in the DSM where ECRs sequence back to an earlier point. Notably, the DSM can be used to point out iterations; as such, we can ask the engineers why they occur and how we can possibly prevent states from sequencing back to improve how we work with ECRs. Last, the model helps to identify deviation from the most common sequence of ECRs as there is a tendency to skip state 19 (i.e., ECR with solving responsible), which is the most common way of working, and sequence instead to state 3 (i.e., identification of solution).

The key results show that the Markov chain DSM is beneficial in analysing the ECR process. It is possible to closely examine the sequences and improve the handling process for ECRs so that they follow the intended way of working. Accordingly, certain sequences should be made mandatory to ensure the correct process flow as recommended by the key users.

5.2 Case II: Natural language processing methods for knowledge management

The aim of the work outlined in this Case is thus to develop a method and test search queries that can be further used for clustering ECR documents, with the aim to discover new patterns and knowledge. Therefore, relationship between documents that will support engineers in making better decisions in future projects were identified. Specifically, we addressed two research questions:

1. *Can NLP and document clustering algorithms be utilized for grouping Engineering Change Requests?*
2. *What benefits do product developers obtain from automated document clustering?*

Previous study at the company presented in Arnarsson *et al.* (2017), interviewed twenty individuals with diverse experience throughout the PD process in order to identify the information needs of product developers. As a result of this study it was identified an opportunity to perform integrated searches across databases that should include structured and unstructured data; however, in the case explored this was not currently accomplished.

Results

The search application considered in this project consists of three modules; the Elasticsearch, the doc2vec model, and the clustering application. These modules are tied together by a frontend search service that can pass queries to both the Elasticsearch and doc2vec modules and then summarize the results using the clustering application.

In the study 8,000 documents were available to match against queries for NLP and document clustering. To evaluate unsupervised methods (clustering) one needs to know some ground truth describing document clusters. This does not exist in the database we consider here, so the approach was instead to test a few queries and let a domain expert validate whether or not the clusters differs and whether or not the documents connected to each cluster match the topics describing the clusters. Three queries were tested, resulting in five clusters for each query and each cluster had seven labels describing the cluster. Number of labels and clusters was predefined. The labels are automatically generated by the document clustering algorithm (LDA) and should summarize the main key words connected to the clusters.

Based on the NLP and document clustering algorithms tested, clusters of similar documents can be identified automatically. Traditionally, product developers would type in a search query in the form of structured or unstructured text. The search results would then be displayed in an ordered list by a selected variable that could then be filtered by structured text. The benefit with the approach employed in this study is that a product developer can type in a search query and the traditional results list of documents will be further clustered into groups. This could then further save time in searching for similar documents as they would appear together in respective clusters. There are a few use cases for this:

- Knowledge management: When making design guidelines with best practice designs, this approach could help to identify related design areas on similar topics to be further documented for future designs.
- Inexperienced product developers: There is often a risk that inexperienced product developers do not have the knowledge of previous work or where to locate these documents. This could therefore be used by them to enable an insight into similar product features that have previously been designed, so these designs can be taken into consideration when working on a new product or re-design of an already existing product.
- Design issues: When product developers are faced with design issues, whether it is related to product quality or during development, they could leverage a search like this to quickly identify previous designs in the database to use for root cause analysis. The amount of documented designs makes this often a challenging task, which could be supported by identifying clusters of related designs.

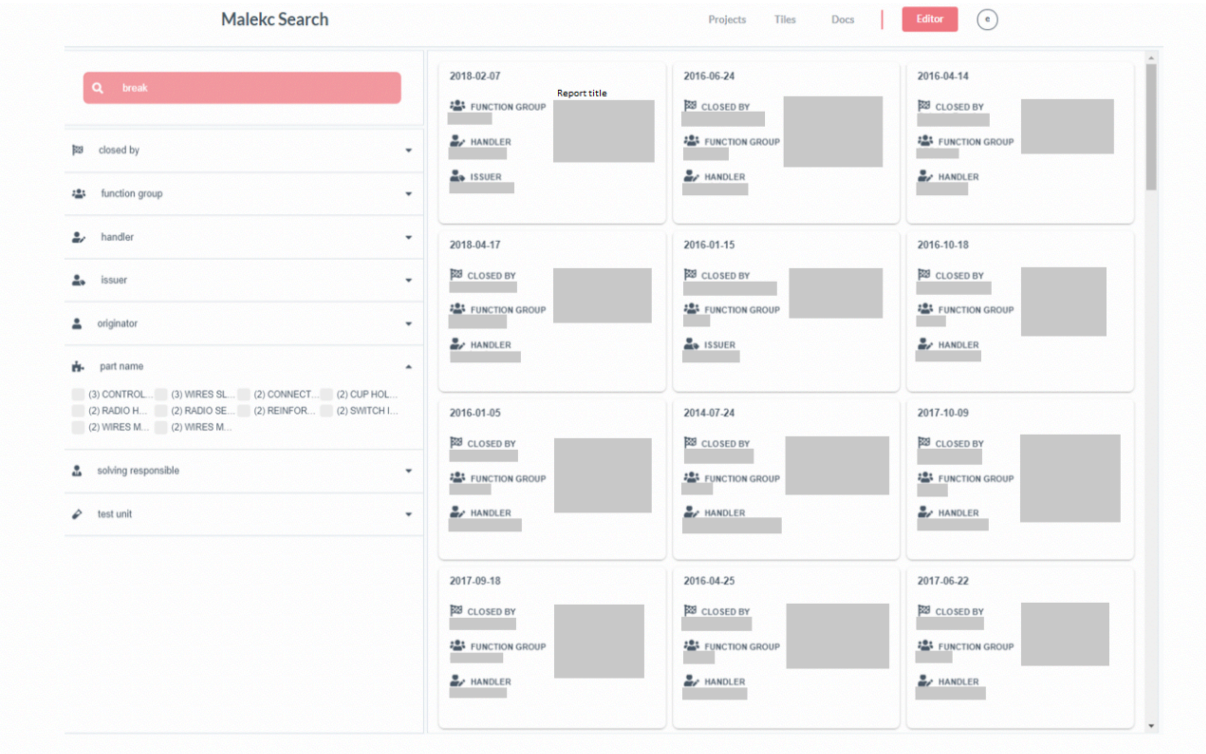

Figure 5 The graphical interface of the search prototype, run in a webbrowser)

Conclusions regarding case II

It can be seen that the NLP and document clustering algorithms work in the cases tested, and provide useful user input. However, further tests are needed to explore its limits. Conclusions regarding the NLP and document clustering approaches are that they seem to work well together in finding similar documents and clustering them together, using an ECR database.

Lastly, in order to achieve better results from the document clustering methods one needs to perform more and better data cleaning. Even though standard NLP cleaning (stop word removal, lemmatization) is performed, one needs to perform more domain specific cleaning of the ECRs in order to really extract the meaning of the free text in the ECRs. Such cleaning might include tasks such as automatic spelling correction and domain specific synonym normalization.

A domain knowledge expert at the case company evaluated the results and confirmed that the algorithms applied managed to find relevant document clusters given the queries tested.

5.3 Case III: Exploring knowledge management in practice

This case is actually a collection of several studies performed in relation to the other activities in the project. They are however central for the approach to capture and validate knowledge.

Methodology
The main part of the research was carried out in conjunction with several (4) master thesis project at AB Volvo to map knowledge areas and potential questions needed to be asked to identify unique and specialized knowledge within a specific field.

Analysis of knowledge capture questions
The top 30 questions which provided the most actionable knowledge was derived from interviews at Volvo GTT (Global Trucks Technology). In all 300 questions was identified, as of specific interest to be able to indemnify and specify knowledge connected to different process stages and subsystems of the product development process, this list is presented in master thesis 1 (chapter 6.2).

Analysis of implementation and adoption of the Checksheet format
In the pre-study of the project as well as during the fall of 2017 a major study was performed at AB Volvo to research the use and adoption of the Checksheet format for knowledge reuse which had started about 5 years earlier, however in a slow moving fashion. In the basic form, the checksheet is a simplistic methodology focusing on the What, why and how of knowledge elements. It was previously identified as a good way of presenting knowledge for reuse, however in order to validate this approach a series of interviews was performed. The result from these studies have been published in JP 1. More details regarding the evaluation at Volvo can be read in Daniel Stenholms PhD thesis, listed in chapter 6.2.

Analysis on technology areas (Specifically Batteries)
Specifically, master thesis 3 (chapter 6.2) focused on battery design in order to further investigate the needs of a new and emerging knowledge area, which included many knowledge gaps needing to be filled. The study involved both internal and external interviews and benchmarking. A study visit at Rejmes was conducted to highlight problems with current design regarding batteries.

Analysis on knowledge sharing formats in and between factories
These studies are mainly presented in Master thesis 2 and 4 (where 4 was a follow up on the barriers identified in 2). The major barriers towards factory wasts was identified and used as input into new factory design. A specific knowledge transfer case with high relevance for future dissemination.

Conclusions regarding case III
One of the knowledge formats previously utilized and that have been implemented at AB Volvo during the past five years is the CheckSheet format. It was also one of the initial goals of the project to evaluate how the knowledge reuse format has been used, how it can be improved and to what extent the ML experiments described in Case 1 and 2 could improve the practice.
The work connected to this theme found the current approach to document knowledge in the form of Checksheets is still viable. Not every project or technology area has come to implement Checksheets in the same way, and the level of success varies also among different departments and knowledge areas.
The thesis projects performed in conjunction with the industrial PhD students at Volvo further highlighted the need of a flexible knowledge format that can be maintained with a minimal amount of resources by the Communities of Practice. In the next case, Case IV the approach to validate captured knowledge is presented.

5.4 Case IV: Tests to automate knowledge validation

As the search and cluster algorithms started to return useful information regarding similarity reports a need was developed to validate the potential knowledge identified. Three experiments was conducted in order to identify how engineers would react and respond to such automated questions. The researchers had a general belief that the risk of feeling annoyed and not responding properly to the questions was a big risk, that eventually also could disqualify the attempts.

The first two experiments was conducted on students at master level, and the third (and final) experiment on experienced engineers working at Volvo GTO. All experiments took place in Fall av 2019.

The setting of the two first experiments was a laboratory exercise on knowledge management where a Lego car was built. In the first experiment participants was given a questionnaire with questions of what they had learned during the lab. The final question was an "opt-in" to answer follow up questions on email (aimed to be validation questions). The process of the experiment is illustrated in Figure 6. The outcome of the first experiment was not convincing. Of 46 students performing the task and asked to be included in the follow up questioning, 18 opted-in to be able to answer questions, 7 students actually answered the question where 4 provided useful and meaningful answers, 2 answers were in direct opposition of what was considered "best practice" at the time, whereas 1 answer was too general to be considered useful.

For the second experiment the setting was somewhat similar, but this time individuals was set in specific Communities of Practice (COP) to discuss the questions, after filling in the same form as in experiment 1. The outcome of experiment 2 showed that the CoP practice generated more useful knowledge than the automated emails did. However, taking a lot more time from the engineering work. Still an observation was that students seemed more engaged in the knowledge creation and seemed not to mind adding the extra time.

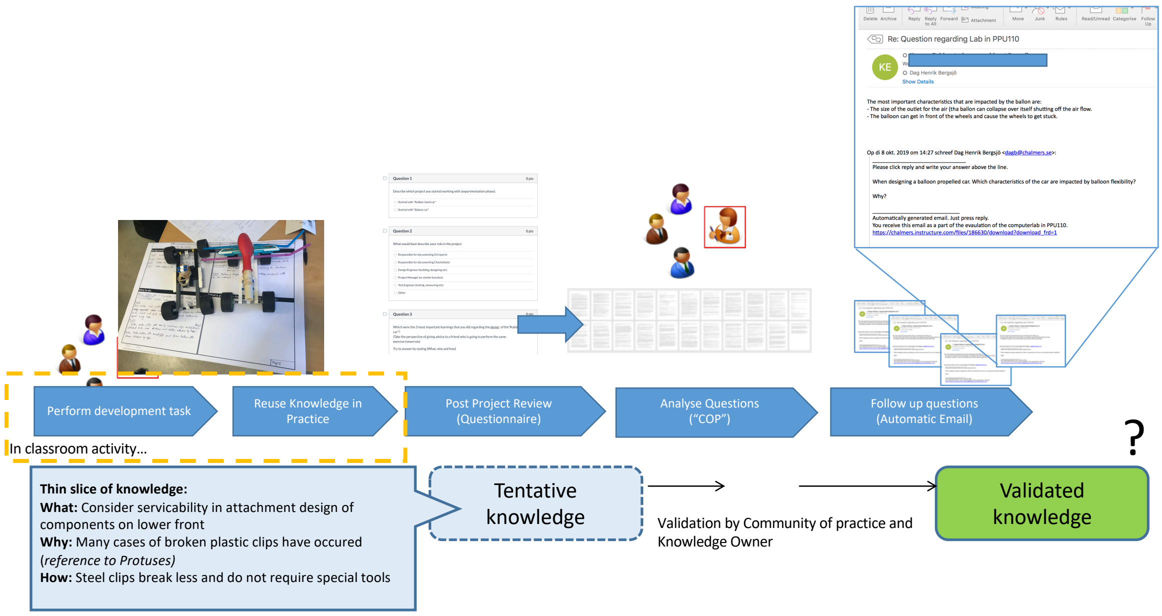

Figure 6Illustration of the process for experiment 1 and 2 in an accademic setting. From left, a development task is performed followed by a knowledge transfer activitiy. Step 3 is the post project review where information is gatherd through a questionnaire. Step 4 involves discussion of technical questions and input from the post project reviews to document knowledge. The final step is the validation of knowledge through the use of emailed questions.

However, experiments on students in a similar environment can be regarded as TRL 5 or 6, and gave valuable information in how to pose questions and stimulate validation of the knowledge a real experiment on real engineers was the only way to see how real engineers with a hectic schedule and other priorities would react. This experiment would then reflect a TRL level of 7-8 in the case of realistic environment.

For the study, eleven engineers were selected from two departments. They first received an introduction letter where it was explained that they would receive an email from a robot, Tinker, which would ask them three questions, which they should spend between 5-10 minutes to answer in written text and then return to the robot. The subject was regarding performing workshops, as this is a general task that is performed across organisations in engineering departments.

The first question could be one of four:
- 1a: What are the most important aspects to consider when performing a workshop?
- 1b: Which tools and methods do you regularly use when performing a workshop?
- 1c: What are the best ways to visualise workshop results for stakeholders and decision makers?
- 1d: How can you tell that a workshop was successful?

Question 1 was sent to three people, question 2 to another three people, question 3 to another three people and question 2 to the last two persons. The answering rate was 100%.

Table 2: Summary of responses

		Technology	Process	People
What	Aspects	0	2	3
	Method	3	2	0
	Visualise	2	1	2
	Success	0	2	2
Why	Aspects	0	2	3
	Method	1	2	0
	Visualise	0	2	1
	Success	0	0	2

Table 2 shows in which category the responses where addressing the knowledge when asked questions of "what" to think of, do or visualise. Second half shows in which category the responses where addressing the knowledge when asked questions of "why" they are proposing certain ways of performing a workshop.

Conclusions regarding Case IV

The intelligent assistant was prototyped using a series of experiments described in Case IV. The promising results identified was that engineers tended to give a higher response rate than students did, and that the engineering response rate was 100% (11 of 11). These show good preliminary results in order to continue to implement one of the commercially available "intelligent assistant"-frameworks e.g. from IBM or Microsoft. Even though the responses from the industrial test was better than expected there are hurdles to cross. E.g. the knowledge asked for was general, whereas spearhead knowledge may be more difficult to ask questions about and may be easily misunderstood. Future studies have to find a balance when a knowledge validation assistant may be useful and when traditional discussions among experts are more efficient. A TRL level around 4 was promised in the application and we are confident that this has been reached (may even be a 5 or 6 depending on the type of knowledge and people asked).

6 Publications

6.1 Knowledge and Dissemination of Results

Planned dissemintation of Results	Mark With X	Comment
Increase knowledge within the area	X	Focus has been to increase knowledge about the potential and future use.
Transfer to other advanced technical development projects.		Not yet, but likely in the next 3-6 months.
Transfer to product development projects.		Not yet specifically although the general checksheet methodology is used daily in product development.
Market introduction		
Being used in policy development.	(N/A)	

6.2 Publications

There are 4 articles (2 journal (JP) and 2 conference (CP) papers currently published).

JP1 Stenholm, D., D. C. Stig, L. Ivansen and D. Bergsjö (2019). "A framework of practices supporting the reuse of technological knowledge." Environment Systems and Decisions: 1-18.

JP2 Corin Stig, D. and D. Bergsjö (2019). "Engineering challenges of intrafirm technology reuse." Systems Engineering **22**(3): 243-254.

CP1 Arnarsson, I. Ö., O. Frost, E. Gustavsson, D. Stenholm, M. Jirstrand and J. Malmqvist (2019). Supporting Knowledge Re-Use with Effective Searches of Related Engineering Documents-A Comparison of Search Engine and Natural Language Processing-Based Algorithms. Proceedings of the Design Society: International Conference on Engineering Design, Cambridge University Press.

CP2 Arnarsson, Í. Ö., E. Gustavsson, J. Malmqvist and M. Jirstrand (2018). ANALYSIS OF ENGINEERING CHANGE REQUESTS USING MARKOV CHAINS. DS 92: Proceedings of the DESIGN 2018 15th International Design Conference.

3 additional journal articles have been submitted but have not yet been published (one is ready for print, one is accepted and one is awaiting initial feedback):

1) "Natural language processing methods for knowledge management"
2) "Modelling industrial engineering change processes using the Design Structure Matrix for sequence analysis: A comparison of multiple projects"
3) "A Lean Framework For Reusing Knowledge Introducing Engineering Checksheets"

PhD Thesis:
Daniel Stenholm, "Reuse of Engineering Knowledge: Perspectives on Experience-Based Codified Knowledge in Incremental Product Development", Chalmers 2018

The PhD Theiss of Ivar Örn Arnasson is expected to be published within 3 months (early 2020).

The following 4 master thesis projects/reports have been conducted in relation to the project at AB Volvo:

1) SALELKAR LAKSHMI PRADIP - Towards an Automated Knowledge Management System – An evaluation of questions to ask when validating newly generated knowledge, Chalmers 2017

2) ASHWIN SATHYAN CHEISLER MACHADO, Quantification of Losses for a single product flow in End-to-End Supply chain, Chalmers 2018

3) Mikaela Collijn Emma Johansson Design for Assembly and Disassembly of Battery Packs, Chalmers 2019

4) JONNY BLOMBERG MARTIN HÅKANSSON, Knowledge Management in the Procurement Process, Chalmers 2019

7 Conclusions and Future Research

If we revisit the initial goal stated that was: “To create conditions to effectively identify and reuse knowledge by utilizing machine-learning algorithms on engineering change reports (ECR)”.

We can say that the main intent of this goal has been met by the studies presented as Case I and II above. We have utilized Markov chains to identify patterns and NLP to create an efficient search engine to identify a cross reference knowledge in engineering change reports ECR.

The intelligent assistant was prototyped using a series of experiments described in Case IV It was further identified that engineers tended to give a higher response rate than students did, and that the engineering response rate was 100% (11 of 11). These show good preliminary results in order to continue the study and implement one of the commercially available “intelligent assistant”-frameworks. Even though the responses from the industrial test was better than expected there are hurdles to cross. E.g. the knowledge asked for was quite general, whereas spearhead knowledge may be more difficult to ask questions about and may be easily misunderstood. Future studies have to find a balance when a knowledge validation assistant may be useful and when traditional discussions among experts are more efficient.

Future work:

There are still work to be done when it comes to evaluate and deploy the solutions presented in this report. Enclosed experiments, ever how relevant, using real engineering data are still experiments. There is a long way to go to further develop and implement these approaches as engineering tools available to engineers in everyday work. At the most we can classify the results from a TRL7-8 maturity, mostly based on the fact that they were performed on real data in a relevant industrial environment.

Regarding the knowledge validation (intelligent assistant) this study shows that further research is needed in order to evaluate which solutions that may be useful, the main contribution in this project is hence a broader knowledge regarding how automated knowledge validation may be performed and what the potential barriers are

This project was based on previous project called Propid. In Fall 2017 MALEKC started focused on building knowledge from deviation reports in product development.

8 Participants and Contacts

The project has been carried out in collaboration between Chamers, Fraunhofer Chalmers Research Center for industrial mathematics (FCC), AB Volvo and Rejmes Transportfordon AB. The consortium made it possible to explore problems that origin in development or production, and which arise in the marketplace. The arrangement as such was valuable and gave many insights that could not be obtained without an extended value chain. Case studies regarding data collection, analysis and discussion were made possible thanks to this approach. The following researchers have been involved in the project:

From Chalmers:
Project Leader, Associate Professor, Dag Bergsjö, Chalmers, dagb@chalmers.se
Professor Johan Malmqvist, johan.malmqvist@chalmers.se
Dr Daniel Stenholm daniel.stenholm@chalmers.se

From FCC:
Associate Professor, Mats Jirstrand mats@fcc.chalmers.se
Dr. Emil Gustavsson emil.gustavsson@fcc.chalmers.se
Otto Frost otto.frost@fcc.chalmers.se

From Volvo:
Industrial PhD student, Ivar Örn Arnarsson ivar.arnarsson@volvo.com
Industrial PhD student, Malin Hane Hagström, malin.hane.hagstrom@volvo.com
Dr Amer Catic, amer.catic@volvo.com

Fron Rejmes Transportfordon AB:
Service Manager: Mattias Fredriksson, mattias.fredriksson@rejmesbil.se

References

Ambady, N. and R. Rosenthal (1992). "Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis." Psychological bulletin **111**(2): 256.

Arnarsson, Í. Ö., E. Gustavsson, J. Malmqvist and M. Jirstrand (2017). Design analytics is the answer, but what questions would product developers like to have answered? DS 87-7 Proceedings of the 21st International Conference on Engineering Design (ICED 17) Vol 7: Design Theory and Research Methodology, Vancouver, Canada, 21-25.08. 2017.

Arnarsson, I. Ö., J. Malmqvist, E. Gustavsson and M. Jirstrand (2016). "Towards big-data analysis of deviation and error reports in product development projects." DS 85-2: Proceedings of NordDesign 2016, Volume 2, Trondheim, Norway, 10th-12th August 2016: 083-092.

Baddeley, A., S. Della Sala, T. Robbins and A. Baddeley (1996). "Working memory and executive control [and discussion]." Philosophical Transactions of the Royal Society B: Biological Sciences **351**(1346): 1397-1404.

Catic, A. and J. Malmqvist (2013). "Effective method for creating engineering checklists." Journal of Engineering Design **24**(6): 453-475.

Gladwell, M. (2007). Blink: The power of thinking without thinking, Back Bay Books.

Grimes, S. (2008). "Unstructured data and the 80 percent rule." Carabridge Bridgepoints: 10.

Jarratt, T., C. M. Eckert, N. H. Caldwell and P. J. Clarkson (2011). "Engineering change: an overview and perspective on the literature." Research in engineering design **22**(2): 103-124.

Kennedy, M. (2008). Ready, set, dominate: implement Toyota's set-based learning for developing products and nobody can catch you, Oaklea Press.

Kobayashi, V. B., S. T. Mol, H. A. Berkers, G. Kismihók and D. N. Den Hartog (2018). "Text mining in organizational research." Organizational research methods **21**(3): 733-765.

Kokkoniemi, J. (2006). Experiences from generating checklists. Proceedings of the IASTED International Conference on Knowledge Sharing and Collaborative Engineering (KSCE2006), St. Thomas, Virgin Islands, USA.

Kuffner, T. A. and D. G. Ullman (1991). "The information requests of mechanical design engineers." Design studies **12**(1): 42-50.

Le, Q. and T. Mikolov (2014). Distributed representations of sentences and documents. International conference on machine learning.

Morgan, J. M. and J. K. Liker (2006). The Toyota product development system. Seattle, WA, Productivity Press.

Pikosz, P. and J. A. Malmqvist (1998). Comparative Study of Engineering Change Management in Three Swedish Engineering Companies. Proceedings of DETC'98, Atlanta, GA, USA, Paper No DET98/EIM-5684.

Sendlhofer, G., et al. (2015). "Implementation of a surgical safety checklist: interventions to optimize the process and hints to increase compliance." PLoS One **10**(2): e0116926.

Steinbach, M., G. Karypis and V. Kumar (2000). A comparison of document clustering techniques, KDD workshop on text mining, Boston.

Wu, X., X. Zhu, G.-Q. Wu and W. Ding (2013). "Data mining with big data." IEEE transactions on knowledge and data engineering **26**(1): 97-107.