

Generative User-Experience Research for Developing Domain-specific Natural Language Processing Applications

Anastasia Zhukova^{1*}, Lukas von Sperl^{2†}, Christian E. Matt^{2†},
Bela Gipp¹

¹University of Göttingen, Germany.

²eschbach GmbH, Germany.

*Corresponding author(s). E-mail(s):

anastasia.zhukova@uni-goettingen.de;

Contributing authors: lukas.vonsperl@eschbach.com;
christian.matt@eschbach.com; gipp@uni-goettingen.de;

†These authors contributed equally to this work.

Abstract

User experience (UX) is a part of human-computer interaction (HCI) research and focuses on increasing intuitiveness, transparency, simplicity, and trust for the system users. Most UX research for machine learning (ML) or natural language processing (NLP) focuses on a data-driven methodology. It engages domain users mainly for usability evaluation. Moreover, more typical UX methods tailor the systems towards user usability, unlike learning about the user needs first. This paper proposes a new methodology for integrating generative UX research into developing domain NLP applications. Generative UX research employs domain users at the initial stages of prototype development, i.e., ideation and concept evaluation, and the last stage for evaluating system usefulness and user utility. The methodology emerged from and is evaluated on a case study about the full-cycle prototype development of a domain-specific semantic search for daily operations in the process industry. A key finding of our case study is that involving domain experts increases their interest and trust in the final NLP application. The combined UX+NLP research of the proposed method efficiently considers data- and user-driven opportunities and constraints, which can be crucial for developing NLP applications.

Keywords: generative research, natural language processing, human-computer interaction, user experience, domain-specific applications

1 Introduction

Natural Language Processing (NLP) has been recently extensively incorporated into industrial and domain applications. For example, NLP is used for speeding up processes, e.g., automation classification of types of customer feedback or filtering out spam emails, information extraction, e.g., named entity recognition to extract symptoms, diagnoses, and treatments from medical records, or auto-completing input forms with language models. Despite the broad integration, domain-specific NLP applications may require practicing more user-driven methodologies to address user needs with these applications.

Often, the data-driven approach falls short in exploring the needs of the domain users (Yang, 2018). On the one hand, domain users are often integrated into development at the late test phase to evaluate the usability of ML/NLP applications (Carney, 2019). Unlike user-driven software development, the development of NLP applications depends mainly on data availability or experimenting with machine learning (ML)/NLP trends and thus is a major driver of application development. On the other hand, the user-driven development of a domain-specific ML/NLP application in medicine showed that close collaboration with the domain users in the earlier stages increases the effectiveness of the final product (Yang, 2017). Therefore, integrating user experience (UX) and human-computer interaction (HCI) research into ML/NLP research addresses users' needs, fuses their expertise, and increases intuitiveness, transparency, simplicity, and trust for the system users (Boukhelifa et al, 2018; Paleyes et al, 2022).

This paper reports on developing a real-world domain NLP application for improving effectiveness and efficiency in daily plant operations in the process industry. A major project goal was to achieve high user acceptance for the final application, which required involving users from the initial design phase. That resulted in the proposed methodology for integrating generative research into a development cycle of domain NLP applications by combining data- and user-driven concept creation, prototype development, and prototype evaluation. Widely applied in design and software development, generative UX research is a methodology that aims to learn about domain users and their stories (Hanington, 2007; Anderson, 2022). To increase the credibility and utility of the final application, generative research engages the users in the early stage of development.

The proposed methodology combines data- and user-driven approaches to full-cycle UX-driven prototype development for NLP applications. Prototyping or developing the first minimal viable products (MVPs) most benefits from the methodology. The proposed methodology consists of the following phases. First, the methodology describes generative UX research steps to explore opportunities for ideas and application concepts. This phase relies on exploratory data analysis and contextual inquiries, i.e., user interviews and on-site observations, as a source of idea generation. Furthermore, the domain users validate a proposed concept and help define the most important features for prototype implementation. Second, the prototyping addresses the most important identified research challenges crucial for usability assessment. Lastly, the prototype evaluation measures user utility via contextual inquiries before and after interacting with the prototype.

Our team comprises UX researchers, data scientists, software developers, and a project leader because such a team composition facilitates a deeper understanding of the capabilities and limitations of the domain NLP application (Dove et al, 2017). Therefore, the paper may have specific interest for the individual practitioners with the above-mentioned roles or similar teams. Moreover, practitioners of user-driven research, e.g., customer research departments and quality assurance officers, could benefit from the proposed methodology and reported results.

The paper is structured as follows. Section 2 briefly reviews previous efforts and UX methodologies in creating ML applications. Section 3 describes a methodology of generative UX research, in which we propose to optimize user experience in the usually data-driven development of domain NLP applications. Our case study’s methodology covers a full-cycle NLP prototype development from generative research for concept ideation to user-driven prototype evaluation. Section 4 reports a case study about implementing the proposed methodology to prototype an industrial NLP application for the process industry. Section 5 discusses the impact of the generative UX research on developing our NLP prototype and summarizes the findings into a reusable methodology for applying generative research to increase the quality of NLP domain applications.

2 Related Work

Domain-driven design (DDD) is a software development methodology that understands, conceptualizes, and models complex domains into applications closely aligned with business needs (Evans, 2004). The core of DDD lies in close interaction between developers, stakeholders, and users to study and implement the business logic of the underlying processes. According to DDD, the building blocks of software should reflect the physical world as closely as possible in terms of both functionality and naming (e.g., (Millett and Tune, 2015; Khononov, 2021; Vernon, 2013; Hippchen et al, 2017, 2019)). The methodology emphasizes understanding a domain by understanding and utilizing its ubiquitous language by all development participants in all planning, development, and refinement stages. Moreover, DDD enables dynamic adjustments of a domain model to changes in real life through an iterative process of collecting feedback from domain users.

User experience research has become a requirement for practical and effective software development in addition to the domain understanding (Fronemann and Peissner, 2014; Law and van Schaik, 2010). UX puts perspective onto users, thus defining functionality and look of new products or features and ensuring a need in new products (Goodman et al, 2012). UX research methods analyze how users interact with a system of interest and explore what users miss or find inconvenient. Quantitative (e.g., questionnaires) and qualitative (e.g., interviews) methods are used to collect and analyze user feedback (Rohrer, 2014). Structured interviews with domain-specific probe questions help UX elicit knowledge about the domain essential for better understanding system requirements (Hoffman et al, 1995).

The rapid growth of ML applications evolved from automating routine processes (e.g., spam filtering) to user-centered applications (i.e., automatic information extraction based on a specific user request). Although ML applications put users into the loop to ensure the usability of the application (Yang, 2017; Dove et al, 2017), most of the ML user-oriented development is product-focused, i.e., the rather later stages engage users to evaluate product usability (Visser et al, 2005). UX-driven research embeds user perspective evenly during the development of ML-powered applications, i.e., at the stages of requirement understanding (Carmona et al, 2018), data collection (Park et al, 2021), implementation (Carney, 2020), and evaluation (Cambo and Gergle, 2018). The UX research explores multiple user interaction and engagement aspects in the development process. UX research increases the effectiveness and communicability of ML applications according to the user’s mental maps (Agner et al, 2020), creates workflows to enable effective joint creativity process of users and ML models (Kayacik et al, 2019), ensures the best user experience from the services (Dove et al, 2017), incorporates expert knowledge into improving the ML applications (Amershi et al, 2014; Dudley and Kristensson, 2018), and engages users into accuracy evaluation (Cambo and Gergle, 2018; Boukhelifa et al, 2018).

Although users participate in multiple steps of the development process, a deeper user-centered perspective is required to distill users’ real needs (Dove et al, 2017). Most UX studies for ML applications follow a one- or two-step product-centered process (Visser et al, 2005; Yang, 2018): (1) interact with the users to learn about a domain, (2) perform usability evaluation of an ML application. Very little research integrates the domain users into the first step to explore their perspectives for idea generation and concept analysis. Exploring user perspectives would include learning about the domain users, their roles in this domain, how they use a currently existing system (if any), and what needs and pain points the users may have. Analysis of this information could evolve the ML development into a user-centered viewpoint.

Generative research, often applied in design and software development, aims to learn domain user stories and personal experiences through exploratory, generative, and evaluative methods (Hanington, 2007). Generative research targets a deeper understanding of user needs and develops product concepts based on the discovered user goals, utility, and underlying reasons for their actions (Anderson, 2022). Generative research is conceptually close to the practical design thinking methodology often applied in industry Brown et al (2008); Lindberg et al (2011); Dell’Era et al (2020). Although the ideation process in both methodologies is user-centered, generative research focuses more on exploring domain users in their environment and deriving prototype ideas based on immersive domain learning.

3 Methodology

In this section, we propose a methodology of generative UX research for optimizing a data-driven domain-specific NLP application toward higher user satisfaction. Generative research encourages user-focused development by getting to know the users better and explores opportunities for an ML/NLP application by learning users’ stories, pain points, and needs (Visser et al, 2005; Anderson, 2022). The methodology

is grounded on a domain-driven design (DDD) (Evans, 2004) of deep domain understanding through close collaboration between developers and domain experts. The iterative knowledge exchange about domain know-how, problem-solving requirements, and implementation constraints leads to effective domain modeling, i.e., a conceptual representation of the domain at consideration. An effective and knowledge-rich domain model reflects internal and external business processes most fully. Moreover, a model should be clear to both domain experts to provide feedback and developers to act as a foundation for implementation. The generative research puts domain experts even more into perspective by switching from business processes to domain experts' workflows and their daily routines.

Although generative research has already been employed in design and software development, it lacks coverage of full-cycle development of ML/NLP applications. For example, Anderson (2022) focused on how to learn about user needs and experiences. Hanington (2007) explored how to create a product concept based on the learned information from users and employed the users to evaluate this concept. Similarly, design thinking proposes a user-centered process of iterative ideating and prototyping (Brown et al, 2008; Lindberg et al, 2011; Dell'Era et al, 2020). Carney (2019) employed users to evaluate a prototype before developing a product for the production-ready state. Moreover, these methodologies do not consider an aspect of data-driven research that should be considered when creating an NLP application. The proposed methodology combines these previous efforts on generative research, design thinking, and domain understanding through exploratory data analysis required in ML/NLP applications.

The methodology for full-cycle prototype development through generative research combines UX+ML/NLP to (1) explore ideas for an NLP application based on the stories of domain users, (2) collect feedback on a concept of a proposed NLP solution to their identified needs or pain points, (3) conduct a user-driven evaluation of a prototype from perspectives of user experience with a new NLP solution and accuracy estimation. Figure 1 depicts our proposed methodology to develop and evaluate a domain-specific NLP application. The methodology consists of three steps:

1. *Exploration*, i.e., design ideas for a prototype through exploratory data analysis, interviews, and observation of the domain users in their environment.
2. *Prototyping* to develop a fully working prototype with back- and front-end.
3. Prototype *evaluation*, i.e., evaluation of the change in user experience and prototype accuracy assessment.

Exploring ideas for an NLP prototype through *generative UX research* aims to collect the necessary information to understand a domain in general, get familiar with the domain data, and learn about the user's needs. We tailor a methodology for concept creation via generative research proposed by Hanington (2007) to a more specific use-case of concept ideation for a domain-specific industry application. A proposed generative research for NLP consists of three steps:

1. *Exploratory data analysis*: Learn about the domain system and data and create hypotheses we could probe during user interviews.

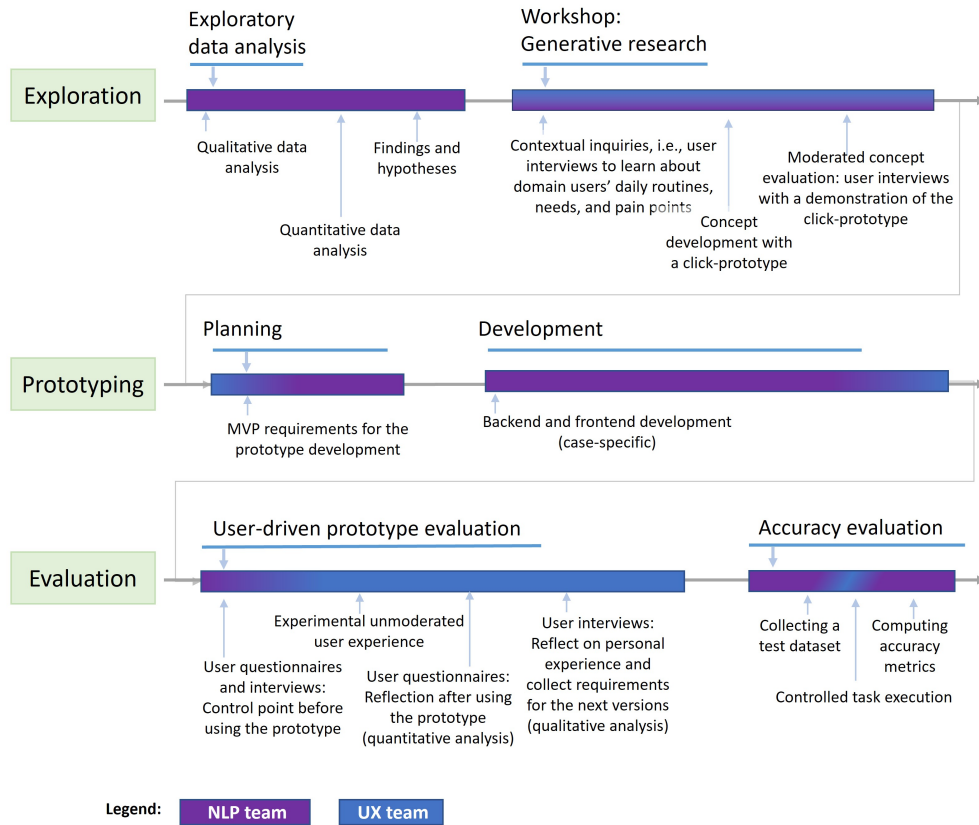


Fig. 1 A timeline for the proposed methodology of generative UX research for prototyping domain NLP applications used in our case study. The methodology consists of three stages: exploration (data exploration and generative research), prototyping (development of the prototype and evaluation (assessment of accuracy and user utility)). The phases show the approximate share of the workload of the UX and NLP teams. For the majority of the phases, both teams need to be involved. Especially during the generative research phase, an equal involvement of the two teams enables deep domain understanding.

2. *Contextual inquiries*: (1) Conduct user interviews in a focus group to learn about their daily routines and systematize user interviews to identify patterns, and (2) Generate ideas to propose a concept for an NLP application or feature.
3. *Concept evaluation*: Perform moderated concept evaluation of a click-dummy, i.e., enable supervised interaction with a “click-dummy” prototype to test the usability and usefulness of the proposed concept.

The user-driven prototype evaluation aims to assess prototype accuracy and how user utility changes after interacting with the prototype. Unlike concept evaluation at the stage of generative UX research, the prototype is evaluated in an open environment. The combination of the two perspectives (i.e., accuracy and user utility) determines

the success level of the prototype in the first step toward a system for shared domain knowledge. The evaluation consists of three steps (Carney, 2019):

1. *Control*: Use questionnaires and user interviews to identify an initial state of user experience, i.e., how users search for the information before independently interacting with the prototype;
2. *Experimental*: Provide the prototype to interact with it through provided tasks and openly throughout daily tasks as they feel fit;
3. *Reflection*: Reuse the same questionnaire and interview scripts to identify how the user experience changes and learn if the system works as expected or surprises them and provides insights.

The methodology combines data- and user-driven approaches that correspondingly originate from ML/NLP and UX research. A team of UX designers, data scientists, and developers ensures diverse and effective collaboration for each stage in the proposed pipeline (Yang et al, 2018). Discussing prototype ideas requires a balanced view of project aspects such as technical capabilities of NLP approaches and models, possible limitations (e.g., data or computation power limitations), and benefit to domain users (Yang, 2018; Carmona et al, 2018). The ultimate goal of the methodology is to increase user involvement and interest in an industrial application and increase system usefulness (Paleyes et al, 2022; Boukhelifa et al, 2018).

4 Case study

This case study describes implementing a domain-specific NLP application into a digital tool for capturing shift logs and managing plant operations in the chemical process industry. Shift logs usually contain problems and solutions arising during production, performed tasks, instructions from supervisors, reports about production goals, changes on production lines, and much more information about daily operations. As the system installed for a study participant plant had been used for several years, a large collection of text logs was available as a data source for NLP applications.

We consulted a user subgroup that closely interacts with the system daily to ideate how an NLP feature or application would look and evaluate which solution would be most effective for the domain system users. The main challenge was that the way the users intended to use the NLP application within the overall scope of tasks did not coincide with the data shape or form. Therefore, the goal of the case study was to devise a pilot idea based on the user perspective and derive an application that will assist domain users' daily tasks based on the data at hand. For this, we employed and implemented our proposed methodology of generative UX research for full-cycle prototyping domain NLP applications.

4.1 Exploration

4.1.1 Exploratory data analysis

A goal of domain understanding and exploratory data analysis is to learn about the area for the prototype and investigate data caveats that can either support or limit various stages of development. Our domain exploration steps included qualitative analysis, i.e., examining text data and workflows of domain experts, and quantitative analysis, i.e., statistical analysis of the text records and composition of domain vocabulary.

Qualitative analysis

We implemented an NLP feature for an existing system of electronic shift books, which is used for logging information about daily operations in the process industry, e.g., chemical or pharmaceutical production. Integrating a feature into an existing system involves conceptualizing difficulties, such as adhering to certain workflows and users' mind maps (Yang, 2018). We conducted a case study with one plant in the chemical domain that operates in the German language. A database with the plant's daily operations logs is used for data analysis and prototype development.

Records in shift books differ from regular text documents by their semi-structured nature, i.e., one record in a logging system contains both unstructured text fields and structured fields, which act as record attributes. Figure 2 depicts a simplified mock-up of the log structure. The assigned attributes can either complement the text data, e.g., text can report about an assigned functional location, or duplicate information in the text, e.g., a functional location is both mentioned in the text and is attached as an attribute. In some cases, the attributes remain empty, although they are mentioned in the text.

Time stamp	Functional locations	Product	Description
...
2021/08/01 10:04	Alpha-L1-R101-P5002 Pump 5002 50 l/min	Phenol	P5002 slowed down, high temp. & strange smell. -- replace leaking shut-off valve mat. no. 14231312 -- flushed towards R105.12 using C ₄ H ₈ O ₂ -- informed maint to TSR
...

Fig. 2 A simplified example of a record in a logging system for the daily operations in the process industry. A record is semi-structured, i.e., it has structured attributes of the records and unstructured text descriptions. The domain language contains professional terms, abbreviations, shortenings, chemical formulas, various IDs (e.g., order IDs), and plant machinery and equipment inventory codes. The example is in English for better readability.

We observed two significant signs of a domain language. First, records contain quite a diverse writing style. While some records contained complete sentences, some looked like handwritten notes about statuses and a list of tasks to address later. Second, the text contains numeric data, abbreviations, product names, special symbols,

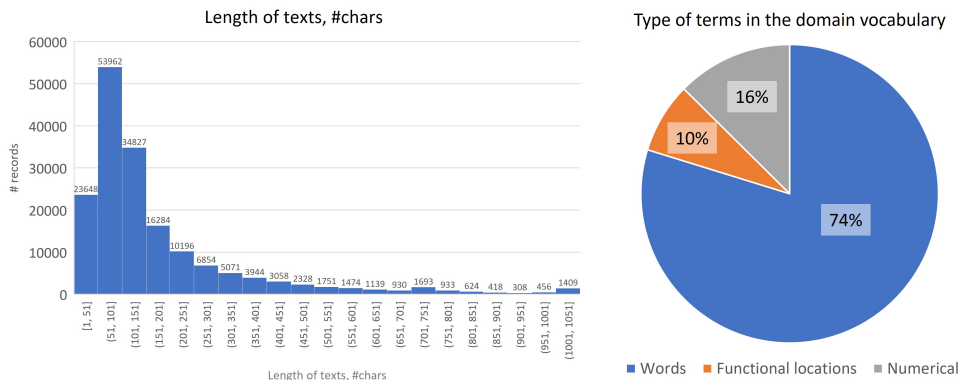


Fig. 3 Quantitative analysis of the domain text records of a logging system. Left: A histogram of the length of text records shows that most records are very short, i.e., smaller than 200 chars. Right: The domain vocabulary includes words and many codes of functional locations (i.e., machinery) that act like nouns. Additionally, the vocabulary contains a lot of digit-containing terms that refer to IDs or numerical values, e.g., measurements.

terminology related to chemistry, and encoded terms that act like words. For example, Figure 2 shows that the terms “P5002” and “R105.12” act as nouns that correspondingly refer to a pump and a reactor. But outside of this domain, knowing the meaning of these terms is impossible. Moreover, the abbreviation TSR, which stands for technical safety review, and “maint,” which stands for “maintenance,” illustrate the short word forms used in this domain.

Quantitative analysis

We performed quantitative data analysis to explore the qualitative analysis findings further and compare a general domain with the production-specific domain terminology. The analysis included the following tasks: (a) evaluate the length of the text records, (b) inspect text composition, and (c) explore the difference between the domain language and the general domain.

Figure 3 shows a histogram of the text lengths, where most records do not exceed 200 chars, i.e., the size of Twitter posts. A pie chart depicts the proportion of the types of terms based on the overall number of terms in a dataset for analysis. We observed a significant impact of the numeric values (10%) and encoded terms of the local domain (16%), i.e., abbreviations for the functional locations used in the production. Such terms require local dictionaries to enable decoding the terms for out-of-domain understanding.

To explore the prevalence of the domain terminology, we visualized terms from the local domain and general use in a scatter plot. For the local domain, we extracted 1000 most frequent nouns from 20,000 randomly selected text records. For the general domain, we used a dump of German Wikipedia ¹ and extracted the 1000 most frequent nouns from the randomly selected 20,000 sentences. We focused on nouns as

¹<https://github.com/GermanT5/wikipedia2corpus>




Fig. 4 Analysis of the domain language shows a significant difference to a general domain (Wikipedia): t-SNE projection of the vectorized domain and general vocabulary clearly shows spots with no overlap between the most frequent terms. Left plot: 1000 general and domain most frequent terms. Right plot: 1000 general and 200 domain most frequent terms. In the list of frequent domain terms, “Product_#” is an anonymized version of the real product names.

the most content-bearing terms. Figure 4 depicts a t-SNE projection of all extracted terms vectorized with fastText word embeddings (Mikolov et al, 2018) and shows that domain and general terms lean towards different parts of the plot. Moreover, on the right plot, we see that some frequent terms are unique and do not overlap with the general domain. For example, the terms “Trennschichtsonde,” “Einstiegsbereit,” and “Laderohr” do not occur in Wikipedia.

Findings and hypotheses

Our qualitative and quantitative analyses led to the following findings: (1) Text quality was quite uneven compared to the regular text, i.e., short texts of Twitter size often have the format of handwritten notes with incomplete sentences. (2) The text data contained many domain terms, i.e., specific terminology for the process and chemical industry, domain abbreviations, product names, and codes of functional locations that act like nouns. We hypothesized that text records with such uneven structure and writing style most likely have a short lifetime and are mainly used to keep track of the current processes and production state rather than as a source of knowledge collected over the long term.

4.1.2 Workshop: Generative research

The objective of generative research is to learn how domain users use the system daily and elicit ideas for an NLP-based solution to improve the effectiveness of their existing processes. We performed contextual inquiries, i.e., one-to-one interviews with representative domain users in their environment (Beyer and Holtzblatt, 1997). We

organized a week-long workshop at a plant to ensure an immersed environment for the research team (Soini et al, 2005; Beyer and Holtzblatt, 1997). The workshop’s goal was to better understand experts’ processes and workflows through observation and interaction with them (Evans, 2004). The observed processes were to be organized into explanatory models to be later distilled into ideas for the prototype implementation.

Contextual inquiries: User interviews

The guided interviews consisted of open-ended questions to learn about the stories of the interviewees’ daily routines. Interviewees in a focus group have expert knowledge of the logging system and represent most of the roles interacting with it, thus ensuring different perspectives of the contextual inquiries. The focus group consisted of seven persons with multiple roles in production: a plant coordinator, three shift leaders, a laboratory manager, a production assistant, and a second plant manager.

The interviews followed two objectives: first, to learn how the interviewees use the logging system during their shifts, and second, to observe if they challenge our hypotheses derived after the exploratory data analysis. The personal interviews contained the following questions and inquiries:

1. Please tell us about your shift. Which events do you usually have during your average work day?
2. How do you use the logging system in your daily tasks? On which occasions? What information would you be interested in there?
3. How would you describe the quality of the logged text? Is there any information that you may lack?

We summarized the interviews as follows: (1) a timeline of a shift that generalizes three stories of the shift leaders, (2) an overview of answers to the last two questions from all interviewees. These forms of summarization followed two different goals. First, we need to find the patterns in the user stories and identify which challenges or problems the interviewees mention at each step that we can turn into ideas for an NLP application. Second, to get an overview of the diverse opinions that can enable a broader picture for generating ideas.

Figure 5 generalizes an average shift of a shift leader. We identify three phases of a shift, i.e., beginning, middle, and end of the shift, which determine which tasks a shift leader needs to perform. For the steps in the phases, we indicated a motivation for a step and which mentioned aspects we could turn into opportunities. For example, at the beginning of a shift, a shift leader reads through the logs of the past 24 hours to collect an overview of what happened when they were not on duty. We learned that if not enough information was logged, they talked to their colleagues, and the information from the oral communication remained unrecorded. We noted that recording the additional information from communication was an opportunity to improve data completeness in the logging system. We identified multiple cases throughout a shift, which indicated that more data could be logged to assist in more informed daily operations. Additionally, we learned that shift leaders rarely use the system to research a solution for solving incidents. The most common way to find information of interest is to read through the recent records or use a string-match search to find specific

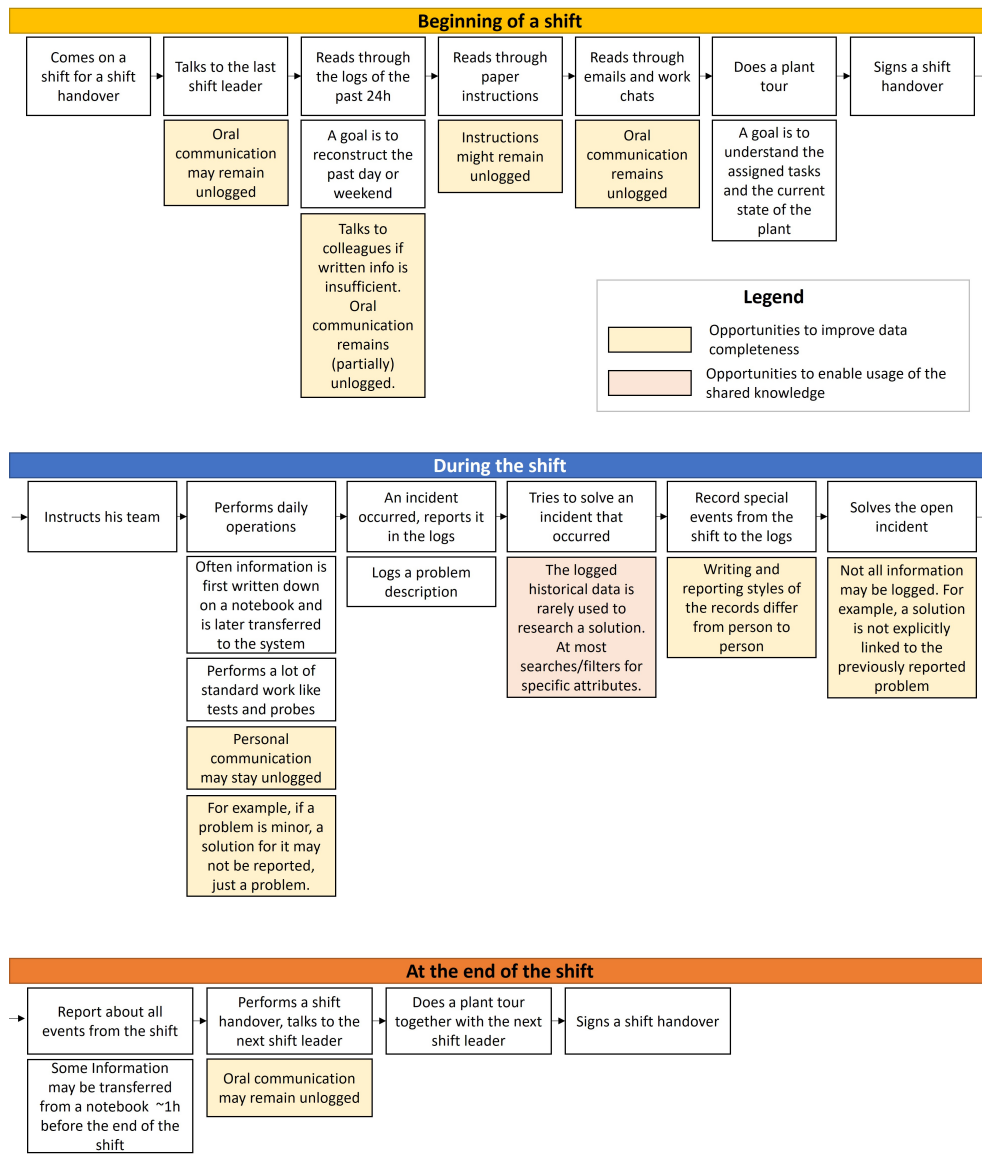


Fig. 5 A generalization of a shift from a shift leader resulting from our interviews that acts as a domain model for our user study. The context inquiries identified three phases of a shift and the main tasks and events that may occur during the day. We provide a short description and outline the challenges of the steps in all phases, which we can turn into opportunities for an NLP application.

keywords, e.g., product names or order IDs. This aspect would be an opportunity to facilitate knowledge sharing in the logging system.

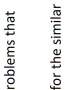
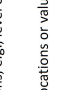
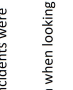
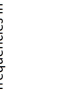
Roles of the interviewees	First round of interviews: Contextual inquiries		Second round of interviews: Concept evaluation
	Information used from the logbook	Data quality	How would you use a system as a source of shared knowledge?
 <p>Shift leader</p>	<ul style="list-style-type: none"> Uses the logbook to get context and more detailed reasons compared to the maintenance books Measurements, level of danger, signatures for repair applications Specific events from specific report, type of event, or functional location Only critical events are interesting over the past 30 days Information about the deviations will be retrieved by filtering items for a specific date and functional locations and then by reading through the matching logs Information about specific experiments Cases with deviations, special analytics Information about the probes, but it is often only emailed Past events are important but hard to read without aggregation Reads logs of daily operations and lab results. Instructions based on the products and special details are highly required for production Need to systematically read the logs to maintain the value of the logged information 	<ul style="list-style-type: none"> Some information is only available in the systems with recipe books Some information needs to be regularly transferred from the paper notes to the log system Some important information, e.g., about products, can come from the file-sharing systems Long entries are of low interest because they take a lot of time to read Tasks are not created during the morning meetings based on the reported information, i.e., it is passed over the other communication channels A guideline on how to report about specific types is missing, thus, data quality varies Not all deviations are logged Reasons/causes of problems can be missing, and it gets hard to understand why a problem occurred and how to prevent it in future More information from the external sources, e.g., excel or pdf files, need to be smoothly integrated into the system More information should be documented because the oral information can be misunderstood Some solutions may not be reported, e.g., for small recurring problems 	<ul style="list-style-type: none"> Find solutions for the problems that occurred in the past Aggregate information for the similar production cases Identify cases when raw materials were not processed in the right time and speed Find quality and safety risks that caused deviations in the processes Have an overview of a full story about an issue reported by multiple departments Identify quality and safety risks that deviations can cause in the processes Check levels of deviations, e.g., level of a mixture in a tank Check for leakages at locations or values of measures Find the context of problems and solutions, e.g., location, what happened, and a value of deviation Identify cases of motor failures, changed filters, warning alarms Find reasons why the incidents were closed Check a level of a batch when looking for a quality problem Find which problems occurred to specific products Read about the reasons why a problem or incident occurred Statistics about the statuses of the recurring tasks, e.g., frequencies in a specific time frame
 <p>Laboratory manager</p>	<ul style="list-style-type: none"> Semi-complete data quality Unclear or "lost" information is discussed in person. Usually, this information is not logged Some paperwork is not logged in the system Quality of the information about the shifts can depend on a reporting person If there is not enough information, check the maintenance scheduling system, talk in person, or consult with a production manager Some instructions and directives are missing Some problems can be reported too shortly or are too difficult to comprehend – oral communication may be required Some information is only communicated personally (oral, email, work chat) More information needs to be collected and shared between the departments Some information stays in the personal communication (oral, emails, work chats) Some events with problems are labeled as info and not as non-critical problems. Therefore, the criticality of the information is underestimated. 		
 <p>Plant manager</p>	<ul style="list-style-type: none"> Information about the probes, but it is often only emailed Past events are important but hard to read without aggregation Reads logs of daily operations and lab results. Instructions based on the products and special details are highly required for production Need to systematically read the logs to maintain the value of the logged information 		
 <p>Plant coordinator</p>			
 <p>Production Assistant</p>			

Fig. 6 An overview of the user replies for the questions from two rounds of interviews: to help create an idea for an NLP prototype and evaluate a proposed prototype concept. From the interviewees with different roles, we learned about the diversity of use cases of the logging system. Although the interviewees described a few cases of data incompleteness, they agreed that using the information from the logged records as shared knowledge can support them daily.

Figure 6 overviews how the interviewees use the logging system and their feedback about the data quality. On the one hand, the interviewees indicated a diverse spectrum of information they use from the system, e.g., acquiring more context to the information from other systems running in the production, staying informed about deviation cases, reading about results of specific experiments, or learning only about critical events over the past 30 days. On the other hand, the interviewees agreed on the problems in the data quality and described it as semi-complete. For example, the interviewees revealed that they often require oral communication to clarify unclear logged information. Additionally, they may have needed to refer to external sources of information to understand the full picture (e.g., emails, work chats, or attached files). Moreover, the reporting style was person-dependent and sometimes lacked a reporting standard for specific events. The diverse writing styles may have depended on the type of reported event, i.e., reporting about some events required a higher level of detail, whereas reporting about others required only short messages due to using other systems for these events.

The interviews confirmed our hypotheses about the high data quality variation and the logged information’s short lifespan. We learned from the interviewees that some records contained very detailed information, and some were scarce. We also learned that depending on the criticality of the events, interest in the records varies on average between 24 hours and 30 days. If the interviewees needed to find any information, they applied available filters to the logged records and read through a filtered list. Occasionally, they use a string-match search to find records containing specific keywords they have previously seen in the text.

Concept development with a click-dummy prototype

Our goal for the NLP application was to improve the exchange of information, i.e., logging and retrieving crucial information. We expected the effective and efficient information exchange to yield higher production efficiency and shorter plant downtime. From the interviews, we learned about the common problems of data completeness, varying text quality, and the short life span of the records. To align with this goal and address the identified problems, we proposed turning the logging system of daily operations into a source of shared domain knowledge (Figure 7). As of the current state, only the most recent logged information is important for the users. We assumed that an NLP-driven domain knowledge processing system would enable leveraging the collected knowledge over the long run. Ultimately, users should be encouraged to log more information and improve their domain expertise by learning from the logged experience.

As a first stage, we proposed a semantic search to enable efficient and effective access to the full history of daily operations. Our hypothesis was an efficient search engine for records should motivate users to log more complete and structured information. If the unstructured records were retrieved effectively, the new records with clear and complete information would have more confidence to be found later on demand. Figure 7 depicts a goal of improving the data completeness in the system by logging more information, e.g., from personal oral and written communication. If the discussions’ key points can be later retrieved on demand, users get more motivated to

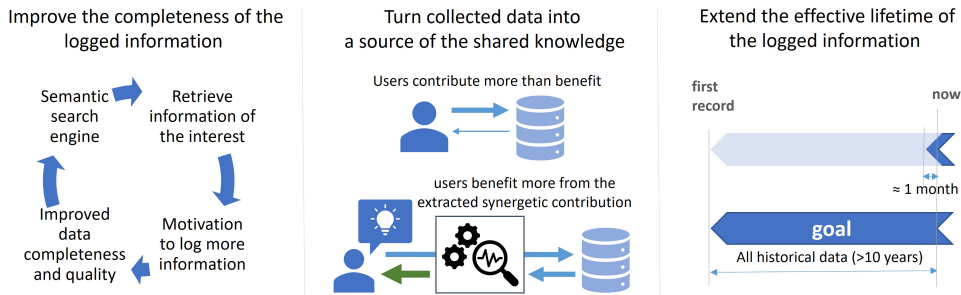


Fig. 7 The historical data of the plant operation contains a lot of know-how information on daily operations, but the use of it is quite limited due to the inefficient search. By implementing efficient logging semantic search and allowing users to find more complex records, the project aims to motivate the logging of more detailed descriptions of daily operations, thus improving the completeness and quality of input data.

transfer this information into the logging system. Representation of the domain language was a core challenge for this domain NLP application. For the prototype phase, we required a feasible solution to design, implement, and test within a short period but also efficient and robust to perform user studies for the prototype evaluation. Section 4.2 reports more detail on this consideration.

To evaluate our concept, we performed Wizard of Oz prototyping (Browne, 2019; Carney, 2019), i.e., developed a “click-dummy” prototype, which mimics the functionality of the front-end of a semantic search engine. We adhere to the UX principle of “fail fast” and collect feedback on the proposed concept in the earliest possible stages of development.

A click-dummy prototype aims to follow two users’ mental models: the logging system and commonly used search engines, e.g., Google. Figure 8 shows that the prototype contains the visual attributes of a search field, a list of the resulting records, and such additional features as highlighting the matching excerpts and a similarity score. Moreover, its overall visual design is familiar to the users by resembling the overall design of the logging system.

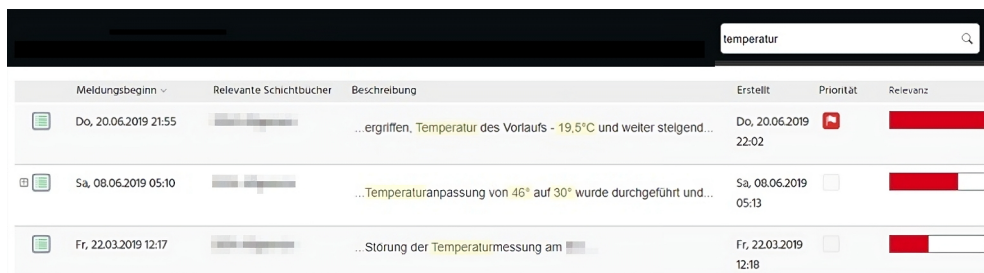


Fig. 8 Based on the opportunities learned from the user interviews, we created a “click-dummy” prototype for a moderated concept evaluation through user interviews (Wizard of Oz prototyping). We demonstrated this prototype to the interviewees and inquired if and how they would use it.

Moderated concept evaluation

We conducted a second round of interviews where we presented the click-dummy prototype and evaluated our concept via guided interviews and moderated usability sessions. In the one-to-one interviews, we aimed to learn if providing efficient access to historical data would allow performing daily tasks in a more informed and effective manner:

1. Would an effective search engine assist you daily? If yes, how and when? Which queries would you devise?
2. Imagine that you have a system that could provide effective access to the expertise and shared knowledge about plant operations. Such a system could answer domain questions, extract information, and identify patterns in the text data. Which inquiries would you have for this system?

All interviewees agreed that efficient search would enable them to access the previously logged knowledge, e.g., solutions to problems. Moreover, they highlighted that a search application would provide additional assistance on a night shift and could assist in planning product launches. The shared ideas for the queries helped us to create lists of queries for the later prototype evaluation (see Section 4.3). Figure 6, the last column, overviews the interviewees' replies on how they would use a system of shared domain knowledge. For example, they highlighted a need to access solutions to previously occurring problems. Moreover, they are expected to retrieve information with a deep understanding. For example, for a query "quality state of product_A on Monday," they would expect to see information about the batch level on that day. Additionally, they would explore information in an aggregating way, e.g., assemble problems that occurred to a specific product.

4.2 Prototyping

The exploration phase resulted in a decision to develop a semantic search engine prototype for documents with uneven text structure and domain-specific language. The prototype's main challenge is handling the domain-specific language, including professional terms, abbreviations, shortenings, chemical formulas, various IDs, and plant machinery and equipment inventory codes. For the prototype development, we followed a *UX principle "fail fast, fail often"* (Dove et al, 2017) and implemented simpler approaches that require less time and resources.

4.2.1 Planning

MVP requirements for the prototype development

From the exploration phase and workshop, we deduct the following main requirements for the prototype: The search engine of the prototype shall be designed in a typical IR process architecture: (1) indexing of text records from a SQL database, (2) retrieving relevant records to a given query². The most challenging requirement for the MVP is to tackle the problem described in the domain-specific language and technical style of the text records. It should use context expansion (i.e., adding definitions and descriptions

²We use the terms "documents" to refer to the records of the logging system, i.e., short documents.

of technical codes to the documents) and encoding documents vector word representation capable of handling out-of-vocabulary domain words, e.g., fastText. Furthermore, the MVP is supposed to be seamlessly integrated into the main software application.

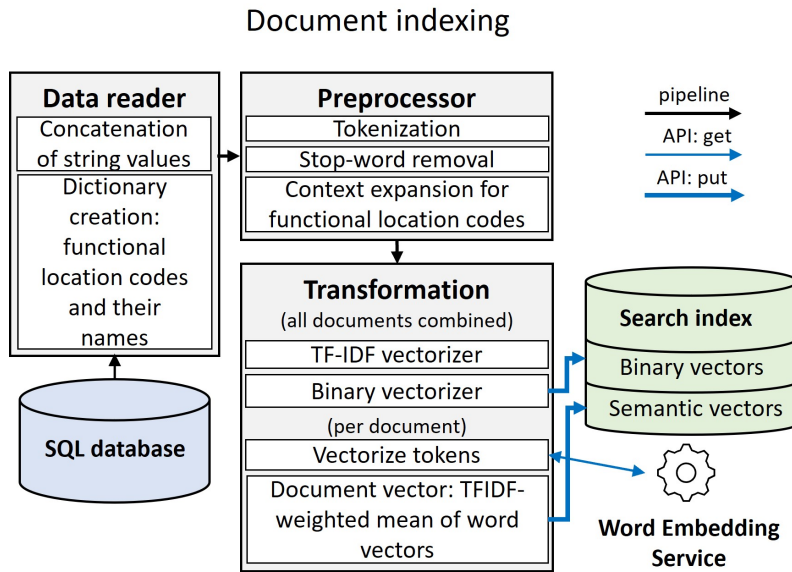


Fig. 9 Indexing text data from a SQL database is a process that includes data reading, preprocessing, data transformation through vectorization of words and documents, and saving the vectors into a search index.

4.2.2 Development - Backend & Frontend

Indexing

Indexing starts with accessing an SQL database and reading values from string columns (Figure 9). We concatenate all text fields of each record into one text document. Preprocessing performs *context expansion* to add semantic descriptors to some domain terminology. We transform documents into multiple vector spaces: (1) vectorization into boolean and term frequency-inverse document frequency (TF-IDF) representations that require full-text collection, (2) per-document vectorization into semantic representation, which employs word embedding service to represent words into fastText word vectors and computes TF-IDF weighted-mean word vectors to represent documents. Finally, we save boolean and semantic document representations into a search index. Document retrieval uses the same word embedding service and represents a query as an average of word vectors. First, candidate retrieval retrieves documents that contain terms to be mandatory present in the documents (depending on if a query contains terms in quotes). Second, it filters the documents by semantic similarity (Section 4.2.2). Lastly, we rank the top M most semantically similar documents by combined semantic similarity on the document and term levels.

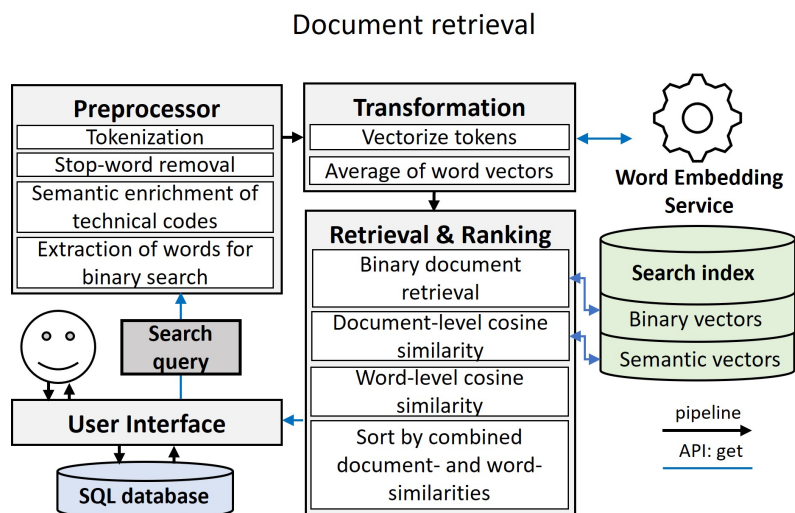


Fig. 10 Record retrieval preprocesses a search query, transforms it into a vector representation, retrieves documents based on binary retrieval and semantic similarity, and ranks the documents based on a combined document and term similarity.

Preprocessing

Preprocessing focuses on (1) providing additional context for some domain-specific terms, i.e., dictionary-based context expansion, and (2) normalizing text for semantic search that employs static word embeddings, e.g., lowercasing, stopword removal, and lemmatization.

To enable context expansion, we create a dictionary about functional locations, i.e., production machinery and equipment installed in a plant. The dictionary contains information about long IDs, their short equivalents, and descriptions. We created this dictionary based on a graph structure of the machinery from the logging system.

We use the domain dictionary to extract descriptions of a mentioned machinery, e.g., in Figure 2, a description of “R105.12” is “reactor.” The dictionary serves to “translate” machinery codes. Descriptions of the codes are added to the text before their occurrence. This approach is similar to automatic query expansion when query formulation is improved with additional information (Carpineto and Romano, 2012). Such context expansion enables retrieving records by containing either specific codes of functional locations or their descriptions.

Representation of the domain-specific language

Our semantic search is based on fastText word vectors, i.e., a skip-gram model for distributed representation on the char n-gram level (Mikolov et al, 2018; Bojanowski et al, 2017; Grave et al, 2018). Most of the word embedding models, such as word2vec (Mikolov et al, 2013) and GloVe Pennington et al (2014), contain only vectors for the words that were in a training corpus. Out-of-vocabulary (OOV) words pose a big problem when word embeddings are applied to specific domains, e.g., biology or chemistry (Galea et al, 2018). Some models propose training word embeddings from

scratch on a domain-specific corpus (cf. (Efstathiou et al, 2018)) or using subword information to learn the representation of the words’ parts to be able to represent later OOV words of this domain (cf. (Zhang et al, 2019)). But we used a pre-trained fastText model because the vectorization on the n-gram level was sufficient to represent German compound terms of this domain³. We use the TF-IDF-weighted average of the word vectors to highlight the impact of the semantically most important tokens. Weighting word embeddings with IDF has successfully been applied for IR tasks (De Boom et al, 2016; Galke et al, 2017; Schmidt, 2019; Jalilifard et al, 2021; Arroyo-Fernández et al, 2019; Agarwal et al, 2019).

We did not use contextualized word embeddings, e.g., language models, because word representation may suffer from incomplete grammar or syntax of the domain short texts in our dataset. Therefore, we used simpler static word embeddings, which can more robustly represent the writing style of hand-written notes and do not require word sense disambiguation within the same narrow domain. Later prototype improvements will be based on the domain adaptation of the dense passage retrieval based on sentence transformers (Reimers and Gurevych, 2019; Wang et al, 2022; Guo and Yu, 2022; Thakur et al, 2021; Wang et al, 2021).

Retrieval and ranking

Document retrieval consists of two parts: semantic search for regular words and exact match for any digit-containing terms (or on demand). While semantic search accommodates synonyms and modification of word forms, exact match enforces specific criteria. Such an approach defines how specific search results should be. For example, a user can search for temperature deviation in general or an inventory code of a functional location.

From the candidate documents, we retrieve documents that are the most semantically related to a query. Candidate retrieval computes document similarity as cosine similarity between a vector representation of a query and vector representations of all candidate documents. A linear search is not time-efficient, but in the current prototype phase, higher accuracy of the approach is more important than execution time. Search in the domain of daily operation is similar to Web search, i.e., only a few of the most relevant records are important for the users. Therefore, we retain the top $K = 200$ most similar documents and proceed with ranking.

Our ranking is based on a harmonic mean of document and term similarities. Term similarity prioritizes documents containing the query’s terms as-is and ranks the documents containing synonyms higher than the query’s terms. Term similarity is computed as a mean of the maximal similarities of the query’s terms to the document’s terms (Figure 11):

³For example, the vocabulary of a trained fastText model does not contain the word “Temperaturschwankungen” (i.e., “temperature fluctuations”). However, a fastText model can still represent this word so that it is semantically close to the words “Temperatur” and “Schwankungen.” Moreover, a shortening such as “Temp.Schwank.” will be vectorized and be semantically similar to “Temperaturschwankungen.”

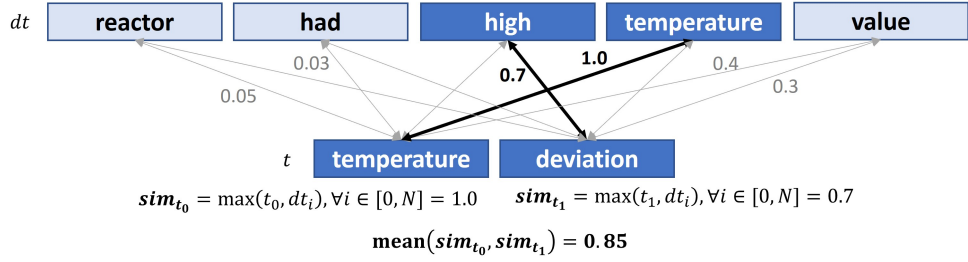


Fig. 11 Term-level similarity identifies the most semantically similar terms between a document and query.

$$TS_j = \frac{\sum_{qt \in QT} \max_{dt \in D_j} \text{cossim}(qt, dt)}{|QT|} \quad (1)$$

where QT is query terms, $|QT|$ is a number of query terms, qt is a vector representation of a query term, D is a document, and dt is a document term. A harmonic mean balances general semantic relevance on the document level and precise similarity on a term level that allows synonyms or shortenings of the terms to gain a high rank. If an exception situation occurs that a query contains only terms for an exact match, then the documents are ranked by their timestamp.

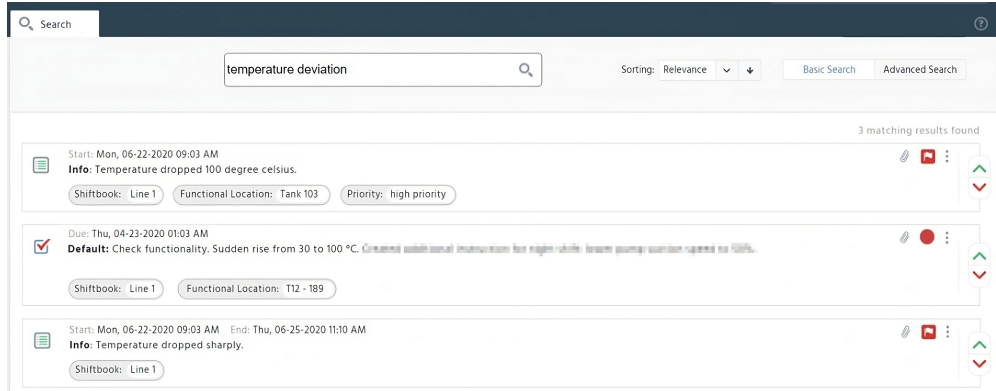


Fig. 12 User interface of the prototype resembles the proposed concept (Figure 8). The prototype contains functionality for the relevance feedback for each record but does not contain previously proposed visual elements that indicate computed relevance. This is made to prevent bias in the user assessment during user studies.

User interface

The user interface of our prototype resembles a click-dummy prototype proposed for the concept evaluation. Figure 12 depicts an interface of the prototype that contains the text of the record, enables sorting by relevance and timestamp, and contains functionality to collect user feedback for the prototype evaluation. Additionally, the prototype provides a detailed view by clicking on a record of interest. For the first version of the prototype, we excluded the functionality of highlighting and visualization of the relevance score to prevent judgment bias during the user-based evaluation of the prototype accuracy (see Section 4.3.2).

4.3 Evaluation

Our user-driven prototype evaluation aims to assess how users utilize the prototype in their work routine, i.e., how the effectiveness of their daily tasks changes after the unmoderated interaction with the prototype, and how accurately the prototype performs. We recruited a focus group of seventeen domain experts and advanced users of the logging system from one plant for the prototype evaluation with a diverse set of roles to ensure opinion variance in the user studies (Table 1). Some study participants were previously recruited for the exploratory generative UX research (see Section 4.1.2). The users were asked to perform two types of interaction: (1) freely use the system as they see fit during their work tasks (*user-driven prototype evaluation*), (2) perform a set of tasks to assess the relevance of search results (*accuracy evaluation*).

Table 1 Demographics of the user study participants for the prototype evaluation. All participants are domain experts and advanced users of the logging system.

Role	Invited to the prototype evaluation	Performed relevance assessment & Questionnaires	Participated in interviews
Shift leader	10	7	4
Second plant manager	2	2	1
Production assistant	2	2	1
Production & Maintenance	2	2	1
Laboratory manager	1	1	1
Total	17	14	8

The section is structured as follows. Section 4.3.2 reports the relevance assessment and evaluation according to the information retrieval (IR) methodology. Section 4.3.1 summarizes questionnaires and user interviews that measure differences in the user experience before and after interacting with the prototype.

4.3.1 User-driven prototype evaluation

UX evaluation estimated how user utility and effectiveness change before and after using the prototype, e.g. if users found the information more efficiently and effectively. The users were asked to complete a questionnaire and were interviewed before the unmoderated interaction with the prototype (i.e., control point). After several

Table 2 Summary of the questions from the user interviews before and after using the prototype about the tool preference to search for the information. Although the interviewees continued using list view for specific tasks, semantic search became a preferred method of finding information of interest.

	List view	Search functionality
Before prototype	<ul style="list-style-type: none"> - The search is not efficient enough - Got used to the list view (x2) - Need to perform comparison of multiple entries and/or specific fields (x2) - Don't know any other way - The fastest solution 	<p>(Keyword search)</p> <ul style="list-style-type: none"> - Use a search to filter out a lot of non-relevant entries - A habit of googling - Know terms for searching - Determine redundancy in the records - Find solutions
After prototype	<ul style="list-style-type: none"> - Search for specific information, e.g., data of old analysis - Faster for some tasks 	<p>(Semantic search)</p> <ul style="list-style-type: none"> - Faster and easier, especially when looking for the older entries - Easy to use - Look for specific solutions, problems, or completed tasks - Find the results faster (x2)

weeks of free use of the prototype, we repeated the questionnaires and interviews to reflect on how their pattern of acquiring information has changed. Additionally, through interviews, we collected free-form feedback on using the prototype and ideas for improvements from a user perspective. Below, we report quantitative analysis on the questionnaires and qualitative analysis of the interviews.

Questionnaires: Control point and reflection

The questionnaires aimed to compare two “snapshots” of the user experience before and after the prototype use. We wanted to see how the study participants obtained information by using the prototype compared to old methods of the logging system when only the default search had been available.

Figure 13 reports the aggregated results of the questionnaires. For each question in Figure 13, the answers are sorted from the worst to the best option. After using the prototype, we observed a shift in the answers towards higher values for all eight questions. We summarized the improvements made by the prototype as follows. (Q1) Before the participants accessed the prototype, they read through a list of logs over the past 24-72 hours instead of using the simple search functionality. (Q2-Q3) We observed that the frequency of search functionality also increased, probably due to the speed-up and better search results of finding answers to their questions. (Q4) The functionality of the context expansion facilitated more efficient search and, consequently, frequent searching for information about them. (Q5) The ultimate goal of the project was to increase the effective lifespan of the logged information (see Figure 7). The chart depicts that after the introduction of the semantic search, the study participants showed more interest in the historical data over one month. (Q6) Performing daily operations became more informed due to the increased number of cases when search results helped answer the questions. (Q7) The study participants indicated they needed to rephrase their queries more rarely than before to obtain expected results. (Q8)

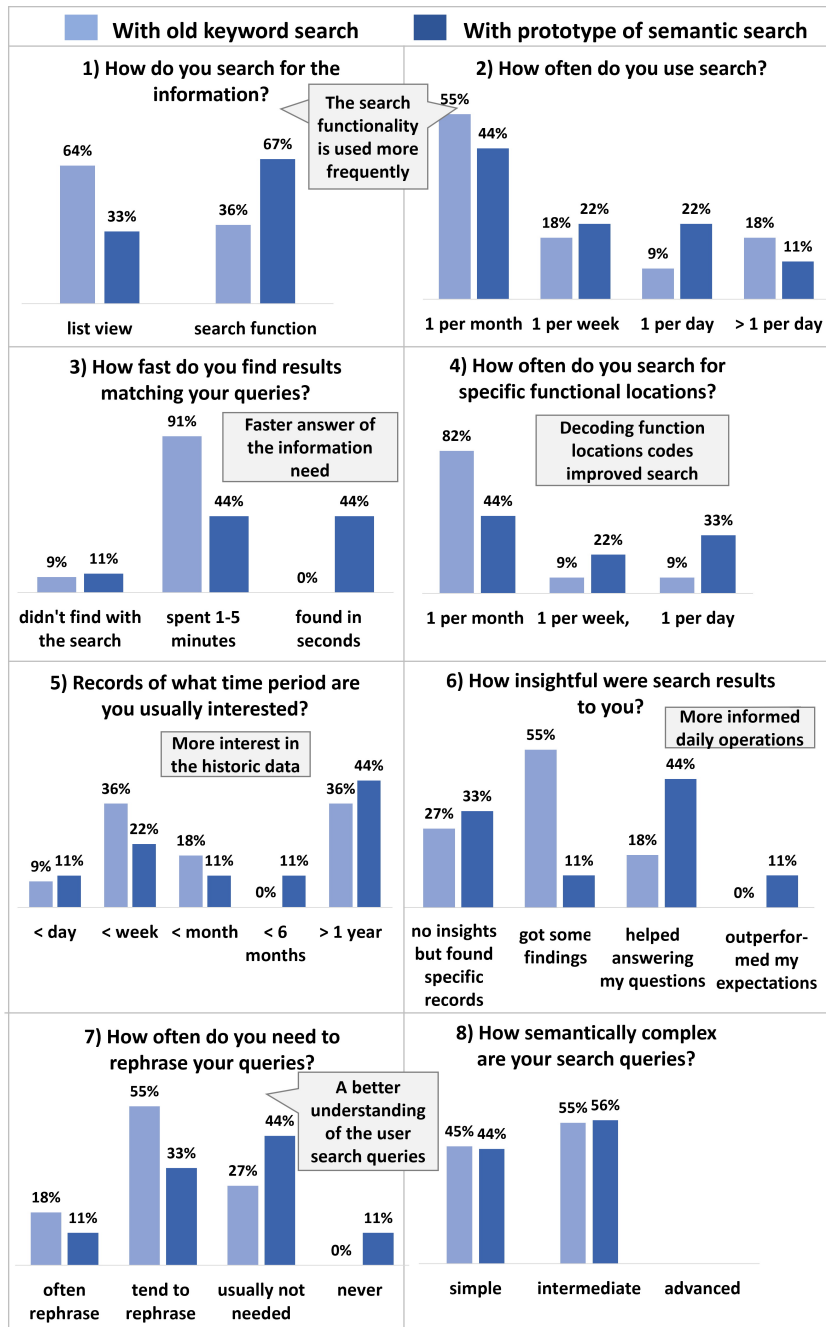


Fig. 13 Comparison of the user experience in searching for information in the system. The prototype showed a positive effectiveness and user utility trend using historical data to assist their daily work.

Although we saw a marginal increase in the complexity of the search queries, the quality of the search results improved (see other questions). We expect the numbers to be more pronounced if operators use semantic search in a production system and create more semantically complex queries.

User interviews: Control point and reflection

User interviews with open-ended questions helped in reflecting on the developed prototype. We collected feedback from interacting with the prototype and ideas for improvements in the next version:

1. How was the experience with the semantic search? How helpful was it? Did it increase your performance?
2. How did you feel about using the prototype? Excited? Curious? Disappointed?
3. Did you search with your search queries? Were there any search queries for which you would have liked to have found more? Was there a search query for which you were particularly impressed by the results or for which you found interesting information?
4. Imagine the prototype is deployed into the productive system. Which functionality would make a difference in using the semantic search regularly? What could be improved?

Table 2 reports summaries of the eight interviews and compares finding information by reading records in a list view to using search functionality, i.e., keyword vs. semantic search. Before the prototype evaluation, most interviewees preferred using a list view as it was a fast method that complemented the inefficient old keyword search. Despite low functionality, the old search was still used to find records by keywords. According to some interviewees, using any search engine felt natural due to regular interaction with the Google search engine. After the pilot, on the one hand, a list view was still preferable for specific tasks, i.e., data comparison. On the other hand, semantic search enabled more effective and efficient access to historical data. According to some interviewees, semantic search was a helpful tool to find solutions to similar problems accounted for in the past.

4.3.2 Accuracy evaluation

The accuracy evaluation focused on the precision-oriented IR evaluation and estimated the relevance of the N top-ranked results. Although such “shallow” evaluation could be incomplete, it is consistent if users are interested only in top-ranked documents in the domain application (Liu and Özsu, 2009). Therefore, the relevance assessment enabled the estimation of the prototype’s accuracy for daily use.

Collecting a test dataset

For the user evaluation, we prepared a list of 100 query topics. We obtained ideas for the query topics during the concept evaluation phase (see Section 4.1.2) and created the queries based on the plant’s logging system records. By following the methodology of Clough and Sanderson (2013), we collected diverse queries, phrases of which occurred in the system database at least once. The queries differ by length of query

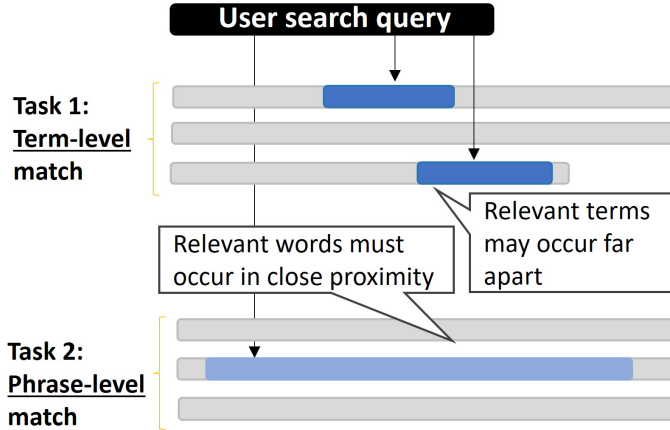


Fig. 14 Study participants needed to perform a set of tasks in the prototype. The tasks included assessing document relevance on two levels: term and phrase match. The difference between the tasks is in the proximity of the relevant terms in a document.

topics, i.e., a query contains 1-4 search terms. Moreover, the queries contain diverse domain terms such as machinery IDs, processes, statuses, and repair steps.

Controlled task execution

Each of the fourteen participants received 20 queries for evaluation and needed to assess the relevance of the search results on two levels: term and phrase (Figure 14). We distributed a list of query topics among the participants so that at least two participants evaluated one query topic. The same queries needed to be executed twice to assess the relevance of each of the 20 top-ranked results on the two levels. For a term level, a document was to be labeled as relevant if all query terms occurred in a document as-is or as a synonym independently from their word order or location within the text. For a phrase level, a document was relevant if a query phrase or its meaning occurred in close proximity, i.e., in a narrow context. For example, for a query topic “problem with a pump,” the documents that described a problem without explicitly mentioning the word “problem,” e.g., “a leakage in a pump,” were to be annotated as relevant.

We calculated inter-rater agreement between the assessors who had evaluated document relevance for the same queries to ensure the experiment’s validity. The assessors evaluated approximately half of the distributed queries, i.e., 44 queries for the first task and 16 queries for the second task were assessed at least twice. We calculated Cohen’s kappa for each task (McHugh, 2012) to measure inter-rater agreement. On the term level, Table 3 reports $\kappa_{term} = 0.856$, a near-perfect agreement between the study participants. On the phrase level, $\kappa_{phrase} = 0.673$, i.e., the study participants reached a substantial agreement. We assume that lower inter-rater agreement on the phrase-level document relevance depends on two factors: (1) an insufficient number of assessed queries for the second task yielded less reliable kappa computation, (2) assessing phrase-level relevance remains a harder task due to ambiguity in the task

description, i.e., what is close proximity. For example, a relevant term of a functional location’s code was at the beginning of a record, whereas the remaining relevant terms occurred later in the document. Therefore, a setup could be unclear on how to assess.

Table 3 Inter-rater agreement of Cohen’s kappa for two tasks: term and phrase level similarity. The level of agreement for both tasks is substantial.

Document relevance level	# queries	κ
term level	44	.856
phrase level	16	.673

To create a test dataset, we combined the assessors’ relevance votes for the query topics assessed more than two times. For each query, we took a list of search results. We summed up votes across assessors and tasks and assigned the results as document relevance scores. For example, if two assessors indicated that a document was relevant on both term and phrase levels, then a document got a relevance score of 4.

Computing accuracy metrics

We evaluated the prototype and the baselines against the documents collected during the expert assessment of the prototype. We acknowledge that such evaluation may cause certain biases favoring the prototype over baselines. But for the precision evaluation of the prototype, we consider such accuracy approximation sufficient to show how user utility improved by using the prototype. In our evaluation setup of a few user study participants, A/B testing would have been inappropriate and could have caused statistically irrelevant results (Gilotte et al, 2018).

After implementing the prototype, we provided two baselines to evaluate how user utility changed: a keyword search (i.e., exact term match) and BM25. The keyword search is a default SQL-based search from the logging system with which the domain users are familiar. BM25 represents a standard baseline used in IR. We implemented a more flexible version of the keyword search for the experiments. A document was retrieved if at least one term from a query was found. The documents were ranked by the decreasing number of overlapping terms and then by a timestamp of the records. For BM25, we normalized German text records and accommodated local domain setup with lemmatizing and removing stopwords. We retrieved documents if cosine similarity was higher than 0.15 and ranked them using it and a document’s timestamp. We implemented an optional context expansion in both baselines to evaluate its impact on handling the domain-specific language (see Section 4.2.2).

We evaluated the accuracy of the prototype at top N records with precision-oriented metrics, i.e., precision $P@N$, and metrics for ranking quality, i.e., mean average precision $mAP@N$ (Liu and Özsu, 2009), mean reciprocal rank MRR (Radev et al, 2002) and normalized discounted cumulative gain $nDCG@N$ (Järvelin and Kekäläinen, 2002). The metrics were calculated for the top 5 and 20 records in our experiments. The combined relevance scores in the test dataset were used to compute $nDCG@N$.

Table 4 Results of the relevance assessment. The proposed prototype of semantic search significantly outperforms all baselines.

Search method	Context expans.	Avg. # retr. docs	<i>MRR</i>	<i>P@5</i>	<i>mAP@5</i>	<i>nDCG@5</i>	<i>P@20</i>	<i>mAP@20</i>	<i>nDCG@20</i>
keyword search	-	3300	.153	.068	.125	.076	.063	.112	.096
	+	5944	.132	.180	.077	.152	.101	.066	.100
BM25	-	736	.582	.294	.548	.322	.178	.479	.291
	+	657	.572	.302	.523	.337	.182	.459	.291
prototype: semantic search	+	90	.888	.749	.880	.725	.668	.805	.772

Table 4 reports the proposed prototype’s evaluation results and baselines. Our prototype outperformed all baselines, but even BM25 showed significant performance improvement compared to the original search method in the logging system. Compared to the keyword search, BM25 showed 3-4 times higher values in all metrics, whereas the semantic search improved evaluation by 5-10 times.

Using context expansion improved precision $P@5$ and $P@20$ of both baselines. Figure 15 depicts the influence of context expansion. The ranking values measured by MRR , $mAP@5$, and $mAP@20$ that ignore the annotated relevance score of the documents have slightly dropped. On the contrary, $nDCG@5$ and $nDCG@20$ have increased. The results show that context expansion assisted in retrieving more candidate documents and ranking relevant documents higher.





Query: “high temperature in reactor”			
	Text	Exact search	Semantic search
w/o context expansion	AB1234 was too hot. CD5612 was clogged up and needed to be cleaned. AB1234 now continues with the recomm. temp.		 (hot, temp.)
with context expansion	reactor AB1234 was too hot. pipe CD5612 was clogged up and needed to be cleaned. reactor AB1234 now continues with recomm. temp.	 (reactor)	 (hot, temp., reactor)

Fig. 15 An example of the positive influence of context expansion on all search methods in the domain language of daily operations in the process industry.

5 Discussion

The paper reports a case study for which we devised and implemented a methodology of making generative UX research a part of a pipeline for prototyping NLP applications. The case study showed that learning directly from users or target audiences enabled the identification of their problems and needs, which consequently led to effective and efficient prototyping with positive user feedback in the evaluation phase. Therefore, instead of evaluating hypothetical ideas for an NLP application, the concept generation phase addressed these specific needs in a problem-solving way. Our

case study took approximately 30 weeks to perform a full cycle to develop and collect extensive feedback on the MVP (4 weeks for exploration, 16 weeks for prototyping, and 10 weeks for evaluation). This period is rather short in terms of full-scale project development but yields better products tailored to the needs of domain experts.

We expect the proposed methodology of combining UX and NLP methodologies to positively impact domain-specific NLP application development. Understanding and evaluating data quality and availability must come in a tight combination with understanding the domain users and their stories. A timeline depicted in Figure 1 can serve as a step-by-step guideline for the generative-driven prototyping of NLP domain applications. The pipeline describes specific tasks for UX researchers, data scientists, and project managers at each stage. Combining the steps in each type of UX research and the questions from the user studies (Section 4.1 and Section 4.3) expands the guideline and equips the UX practitioners with applied examples. Section 4.1.1, Section 4.1.2, and Section 4.2 exemplify data scientists how we performed exploratory data analysis, assessed and solved the challenges for a domain NLP prototype application.

We encourage NLP and UX researchers and practitioners to join in evaluating this methodology and reporting their results. Currently, we are testing the methodology in a second case study in which we prototype an NLP application for supporting domain users in logging information. There, we observe and explore how the domain users log information, which steps are hard for them, and which information they enter wrongly or incompletely. Although it is too early to present the results of the case study in full detail, they support our conclusion that using generative UX research for NLP applications strongly improves user acceptance.

Impact of generative UX research on the domain NLP applications

Generative UX research focuses on contextual inquiries to learn about the daily routines of domain users and how they use an existing product (if any). Observing and exploring which challenges, concerns, and needs they have enables opportunities for a new NLP application/feature. To facilitate immersed domain understanding, contextual inquiries include interviews with the domain users and observations in their environment. Generative research engages domain users from the first step of the development process and targets the identified pain points. This problem-solving methodology increases interest in the final product and trust in the final application. In our case study, we applied generative UX research to learn how the interviewees use the logging system in their daily operations in the process industry domain. The interviews highlighted two problems: (1) the logged text information has an uneven reporting style and suffers from partial incompleteness, and (2) users mostly used only recently logged information and used inefficient ways to access historical data. The identified problems establish a ground for concept generation.

Exploratory data analysis is essential for ML/NLP applications to understand data availability and quality. In generative research, the preliminary data analysis enables learning about the data constraints that must be considered during the concept generation phase. Moreover, the data analysis facilitates forming hypotheses about patterns of special usage of a domain system, which can be later tested during user interviews. Our exploratory data analysis revealed multiple challenges of the domain text data:

(1) semi-structured data records included structured attributes and unstructured text fields, (2) writing style in the records varied from large structured reports to notes quickly taken down with incomplete grammar and syntax, (3) domain language contained rare and domain special terminology, e.g., abbreviations, term shortenings, chemical formulas, numerical values, and codes of the machinery that act as domain terms.

Generated ideas and concepts from UX research help define the scale of a project. Selecting small ideas for a prototype phase speeds up evaluating the credibility of a bigger concept. Based on our UX research, we created an idea for a big-scale system for domain knowledge sharing, e.g., a solution recommendation system that is based on the experience of plant operation, collected and extracted from the historical records (see Figure 7). We selected a small part from this idea for prototyping and proposed a semantic search to provide effective access to the historical data. The prototype provided us with quick feedback from the domain users and suggested steps for further development.

Prototyping for UX+NLP research: quick and effective

Prototyping for NLP applications should aim at the UX principle “fail fast, fail often.” It means that a primary goal is to evaluate prototype credibility, which leads to a prototype containing only basic or simplified functionality sufficient for estimating the user utility. Moreover, industrial or domain NLP applications may not contain top-tier research because real-world applications need to deal with more constraints than academic research. Simpler methods do not always mean less effective but more robust and suitable to a specific domain or low-resource language. In our case study, we used a combination of two solutions to handle the domain-specific language in the prototype: (1) fastText word embeddings that reliably represent domain compound German terms and (2) context expansion that adds descriptions to the mentioned machinery codes. Although these solutions are not from the cutting-edge research of language models and their domain adaptation, the prototype enabled reliable evaluation of the user utility and usefulness of the proposed NLP application.

Prototype evaluation: effectiveness of the daily task completion

Prototype usefulness and user satisfaction in domain-specific NLP applications are often bound to the personal effectiveness of completing the tasks for which the prototype was intended. Depending on the application, effectiveness can be replaced or extended with efficiency. Since the effectiveness of completing tasks and user utility often depend on the accuracy of an NLP application, the evaluation needs to address a combination of prototype accuracy evaluation when the system is in use and user satisfaction in assisting in their tasks. Quantitative and qualitative analysis of accuracy metrics, questionnaires, and interviews can yield the most complete arguments to decide about the prototype’s success or failure and points for its improvement.

In our case study, user utility in the process industry depends on the effective and efficient acting on occurring problems, which yields safety and quality of the technological processes. The evaluation methodology employs domain assessors to estimate the accuracy of the prototype and its impact on their daily routine. First, we perform

a precision-oriented evaluation of the prototype and ask domain users from a focus group to assess the top search results of a given list of queries. Second, we analyze user utility with questionnaires and personal interviews. Analysis of the questionnaire difference before and after the prototype evaluation showed a positive trend in improving user experience working with historical logged data.

6 Conclusion

The paper reports on a case study of developing a domain-specific NLP application prototype in the process industry. To develop a prototype, we proposed a methodology for integrating generative UX research into the data-driven methodology of developing NLP applications by reinforcing domain understanding through data analysis and user interviews. Generative UX research focuses on domain users and employs them at the initial stages of prototype development, i.e., ideation and concept evaluation and the last stage of the prototype evaluation. The core idea of generative research is that involving domain users increases user utility, interest, and trust in the final application since it targets their needs.

The case study reports a full prototype development cycle of a domain-specific NLP application for daily operations in the process industry. In the generative UX research, we learn about a complicated domain of the German language, varying writing styles, partial data incompleteness in the logging system, and the short lifetime of the text records. We propose a semantic search to enable effective access to the historical knowledge of plant operations and motivate users to improve data completeness by logging more information. The user-oriented prototype evaluation focuses on assessing the accuracy and user utility when interacting with the prototype with a set of provided tasks or openly as they see fit. The prototype evaluation demonstrates more informed, effective, and efficient task completion in daily operations.

The paper shows that generative research for NLP application development combines the strengths of both UX and NLP research. The methodology is a promising direction for broad use in the development of NLP applications. The UX+NLP combination reveals opportunities and constraints of the data- and user-driven nature through exploratory data analysis and contextual inquiries, which enables a more informed and effective development process. We encourage UX+NLP teams to work in close collaboration to test the proposed methodology and report their results.

Acknowledgements

This Project is supported by the Federal Ministry for Economic Affairs and Climate Action (BMWK) on the basis of a decision by the German Bundestag and by the Lower Saxony Ministry of Science and Culture and the VW Foundation.

References

Agarwal L, Thakral K, Bhatt G, et al (2019) Authorship clustering using tf-idf weighted word-embeddings. In: Proceedings of the 11th Annual Meeting of the

- Forum for Information Retrieval Evaluation. Association for Computing Machinery, New York, NY, USA, FIRE '19, p 24–29, <https://doi.org/10.1145/3368567.3368572>, URL <https://doi.org/10.1145/3368567.3368572>
- Agner L, Neczyk B, Renzi A (2020) Recommendation systems and machine learning: Mapping the user experience. In: Design, User Experience, and Usability. Design for Contemporary Interactive Environments: 9th International Conference, DUXU 2020, Held as Part of the 22nd HCI International Conference, HCII 2020, Copenhagen, Denmark, July 19–24, 2020, Proceedings, Part II. Springer-Verlag, Berlin, Heidelberg, p 3–17, https://doi.org/10.1007/978-3-030-49760-6_1, URL https://doi.org/10.1007/978-3-030-49760-6_1
- Amershi S, Cakmak M, Knox WB, et al (2014) Power to the people: The role of humans in interactive machine learning. *AI Magazine* 35(4):105–120. <https://doi.org/10.1609/aimag.v35i4.2513>, URL <https://ojs.aaai.org/aimagazine/index.php/aimagazine/article/view/2513>
- Anderson N (2022) Generative research: A complete guide to running a successful study. URL <https://dscout.com/people-nerds/generative-research-complete-guide>
- Arroyo-Fernández I, Méndez-Cruz CF, Sierra G, et al (2019) Unsupervised sentence representations as word information series: Revisiting tf-idf. *Computer Speech & Language* 56:107–129. <https://doi.org/https://doi.org/10.1016/j.csl.2019.01.005>, URL <https://www.sciencedirect.com/science/article/pii/S0885230817302887>
- Beyer H, Holtzblatt K (1997) Contextual Design: Defining Customer-Centered Systems. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA
- Bojanowski P, Grave E, Joulin A, et al (2017) Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics* 5:135–146. https://doi.org/10.1162/tacl_a_00051, URL <https://aclanthology.org/Q17-1010>
- Boukhelifa N, Bezerianos A, Lutton E (2018) Evaluation of Interactive Machine Learning Systems. In: Zhou J, Chen F (eds) *Human and Machine Learning Visible, Explainable, Trustworthy and Transparent*. Springer, p 341–360, URL <https://hal.inria.fr/hal-01845018/file/evaluation-interactive-machine.pdf>
- Brown T, et al (2008) Design thinking. *Harvard business review* 86(6):84
- Browne JT (2019) Wizard of oz prototyping for machine learning experiences. In: *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, CHI EA '19, p 1–6, <https://doi.org/10.1145/3290607.3312877>, URL <https://doi.org/10.1145/3290607.3312877>

- Cambo SA, Gergle D (2018) User-Centred Evaluation for Machine Learning, Springer International Publishing, Cham, pp 315–339. https://doi.org/10.1007/978-3-319-90403-0_16, URL https://doi.org/10.1007/978-3-319-90403-0_16
- Carmona K, Finley E, Li M (2018) The relationship between user experience and machine learning. Available at SSRN 3173932 <https://doi.org/http://dx.doi.org/10.2139/ssrn.3173932>, URL <https://ssrn.com/abstract=3173932>
- Carney M (2019) User Research for Machine Learning Systems: A Case Study Walkthrough — dscout.com. <https://dscout.com/people-nerds/user-research-for-machine-learning>
- Carney M (2020) What is Machine Learning + UX? URL <https://medium.com/ml-ux/what-is-ml-ux-71d5e6d6ce9>
- Carpineto C, Romano G (2012) A survey of automatic query expansion in information retrieval. *ACM Comput Surv* 44(1). <https://doi.org/10.1145/2071389.2071390>, URL <https://doi.org/10.1145/2071389.2071390>
- Clough P, Sanderson M (2013) Evaluating the performance of information retrieval systems using test collections. *Information Research* 18(2). URL <https://eprints.whiterose.ac.uk/78886/>, © Year The Author(s). This is an Open Access article distributed under the terms of the Creative Commons Attribution Licence (<http://creativecommons.org/licenses/by/3.0>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.
- De Boom C, Van Canneyt S, Demeester T, et al (2016) Representation learning for very short texts using weighted word embedding aggregation. *Pattern Recognition Letters* 80:150–156. <https://doi.org/https://doi.org/10.1016/j.patrec.2016.06.012>, URL <https://www.sciencedirect.com/science/article/pii/S0167865516301362>
- Dell’Era C, Magistretti S, Cautela C, et al (2020) Four kinds of design thinking: From ideating to making, engaging, and criticizing. *Creativity and Innovation Management* 29(2):324–344. <https://doi.org/https://doi.org/10.1111/caim.12353>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/caim.12353>, <https://onlinelibrary.wiley.com/doi/pdf/10.1111/caim.12353>
- Dove G, Halskov K, Forlizzi J, et al (2017) Ux design innovation: Challenges for working with machine learning as a design material. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, CHI ’17, p 278–288, <https://doi.org/10.1145/3025453.3025739>, URL <https://doi.org/10.1145/3025453.3025739>
- Dudley JJ, Kristensson PO (2018) A review of user interface design for interactive machine learning. *ACM Trans Interact Intell Syst* 8(2). <https://doi.org/10.1145/3185517>, URL <https://doi.org/10.1145/3185517>

- Efstathiou V, Chatzilenas C, Spinellis D (2018) Word embeddings for the software engineering domain. In: Proceedings of the 15th International Conference on Mining Software Repositories. Association for Computing Machinery, New York, NY, USA, MSR '18, p 38–41, <https://doi.org/10.1145/3196398.3196448>, URL <https://doi.org/10.1145/3196398.3196448>
- Evans E (2004) Domain-driven design: tackling complexity in the heart of software. Addison-Wesley, Boston
- Fronemann N, Peissner M (2014) User experience concept exploration: User needs as a source for innovation. In: Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational. Association for Computing Machinery, New York, NY, USA, NordiCHI '14, p 727–736, <https://doi.org/10.1145/2639189.2641203>, URL <https://doi.org/10.1145/2639189.2641203>
- Galea D, Laponogov I, Veselkov K (2018) Sub-word information in pre-trained biomedical word representations: evaluation and hyper-parameter optimization. In: Proceedings of the BioNLP 2018 workshop. Association for Computational Linguistics, Melbourne, Australia, pp 56–66, <https://doi.org/10.18653/v1/W18-2307>, URL <https://aclanthology.org/W18-2307>
- Galke L, Saleh A, Scherp A (2017) Word embeddings for practical information retrieval. In: Eibl M, Gaedke M (eds) INFORMATIK 2017. Gesellschaft für Informatik, Bonn, pp 2155–2167, https://doi.org/10.18420/in2017_215
- Gilotte A, Calauzènes C, Nedelec T, et al (2018) Offline a/b testing for recommender systems. In: Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining. Association for Computing Machinery, New York, NY, USA, WSDM '18, p 198–206, <https://doi.org/10.1145/3159652.3159687>, URL <https://doi.org/10.1145/3159652.3159687>
- Goodman E, Kuniavsky M, Moed A (2012) Observing the User Experience (Second Edition), second edition edn. Morgan Kaufmann, Boston, <https://doi.org/https://doi.org/10.1016/B978-0-12-384869-7.18001-5>, URL <https://www.sciencedirect.com/science/article/pii/B9780123848697180015>
- Grave E, Bojanowski P, Gupta P, et al (2018) Learning word vectors for 157 languages. In: Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018). European Language Resources Association (ELRA), Miyazaki, Japan, URL <https://aclanthology.org/L18-1550>
- Guo X, Yu H (2022) On the domain adaptation and generalization of pretrained language models: A survey. arXiv preprint arXiv:221103154
- Hanington BM (2007) Generative research in design education. International Association of Societies of Design Research2007: Emerging Trends in Design Research pp 12–15

- Hippchen B, Giessler P, Steinegger R, et al (2017) Designing microservice-based applications by using a domain-driven design approach. *International Journal on Advances in Software* 10(3 & 4):432–445
- Hippchen B, Schneider M, Giessler P, et al (2019) Systematic application of domain-driven design for a business-driven microservice architecture. *International Journal on Advances in Software* 12(3 & 4):343–355
- Hoffman RR, Shadbolt NR, Burton A, et al (1995) Eliciting knowledge from experts: A methodological analysis. *Organizational Behavior and Human Decision Processes* 62(2):129–158. <https://doi.org/https://doi.org/10.1006/obhd.1995.1039>, URL <https://www.sciencedirect.com/science/article/pii/S0749597885710394>
- Jalilifard A, Caridá VF, Mansano AF, et al (2021) Semantic sensitive tf-idf to determine word relevance in documents. In: Thampi SM, Gelenbe E, Atiquzzaman M, et al (eds) *Advances in Computing and Network Communications*. Springer Singapore, Singapore, pp 327–337
- Järvelin K, Kekäläinen J (2002) Cumulated gain-based evaluation of ir techniques. *ACM Trans Inf Syst* 20(4):422–446. <https://doi.org/10.1145/582415.582418>, URL <https://doi.org/10.1145/582415.582418>
- Kayacik C, Chen S, Noerly S, et al (2019) Identifying the intersections: User experience + research scientist collaboration in a generative machine learning interface. In: *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, CHI EA '19, p 1–8, <https://doi.org/10.1145/3290607.3299059>, URL <https://doi.org/10.1145/3290607.3299059>
- Khononov V (2021) *Learning Domain-driven Design: Aligning Software Architecture and Business Strategy*. O'Reilly Media, Incorporated
- Law ELC, van Schaik P (2010) Modelling user experience – An agenda for research and practice. *Interacting with Computers* 22(5):313–322. <https://doi.org/10.1016/j.intcom.2010.04.006>, URL <https://doi.org/10.1016/j.intcom.2010.04.006>, <https://academic.oup.com/iwc/article-pdf/22/5/313/1988154/iwc22-0313.pdf>
- Lindberg T, Meinel C, Wagner R (2011) *Design Thinking: A Fruitful Concept for IT Development?*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp 3–18. https://doi.org/10.1007/978-3-642-13757-0_1, URL https://doi.org/10.1007/978-3-642-13757-0_1
- Liu L, Özsu MT (eds) (2009) *Encyclopedia of Database Systems*. Springer Reference, Springer, New York, <https://doi.org/10.1007/978-0-387-39940-9>
- McHugh ML (2012) Interrater reliability: the kappa statistic. *Biochemia medica* 22(3):276–282

- Mikolov T, Sutskever I, Chen K, et al (2013) Distributed representations of words and phrases and their compositionality. In: Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2. Curran Associates Inc., Red Hook, NY, USA, NIPS'13, p 3111–3119
- Mikolov T, Grave E, Bojanowski P, et al (2018) Advances in pre-training distributed word representations. In: Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018). European Language Resources Association (ELRA), Miyazaki, Japan, URL <https://aclanthology.org/L18-1008>
- Millett S, Tune N (2015) Patterns, Principles, and Practices of Domain-Driven Design. Wiley
- Paley A, Urma RG, Lawrence ND (2022) Challenges in deploying machine learning: A survey of case studies. ACM Comput Surv 55(6). <https://doi.org/10.1145/3533378>, URL <https://doi.org/10.1145/3533378>
- Park S, Wang AY, Kawas B, et al (2021) Facilitating Knowledge Sharing from Domain Experts to Data Scientists for Building NLP Models, Association for Computing Machinery, New York, NY, USA, p 585–596. URL <https://doi.org/10.1145/3397481.3450637>
- Pennington J, Socher R, Manning C (2014) GloVe: Global vectors for word representation. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP). Association for Computational Linguistics, Doha, Qatar, pp 1532–1543, <https://doi.org/10.3115/v1/D14-1162>, URL <https://aclanthology.org/D14-1162>
- Radev DR, Qi H, Wu H, et al (2002) Evaluating web-based question answering systems. In: Proceedings of the Third International Conference on Language Resources and Evaluation (LREC'02). European Language Resources Association (ELRA), Las Palmas, Canary Islands - Spain, URL <http://www.lrec-conf.org/proceedings/lrec2002/pdf/301.pdf>
- Reimers N, Gurevych I (2019) Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP). Association for Computational Linguistics, Hong Kong, China, pp 3982–3992, <https://doi.org/10.18653/v1/D19-1410>, URL <https://aclanthology.org/D19-1410>
- Rohrer C (2014) When to use which user-experience research methods. Nielsen Norman Group 12:21
- Schmidt CW (2019) Improving a tf-idf weighted document vector embedding. 1902.09875

- Soini K, Pirinen A, Suominen J (2005) Workshops-collaborative arena for generative research. In: Proceedings of DPPI. Eindhoven University of Technology
- Thakur N, Reimers N, Rücklé A, et al (2021) Beir: A heterogenous benchmark for zero-shot evaluation of information retrieval models. In: Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS 2021) - Datasets and Benchmarks Track (Round 2), URL <https://arxiv.org/abs/2104.08663>
- Vernon V (2013) Implementing Domain-Driven Design. Addison-Wesley
- Visser FS, Stappers PJ, van der Lugt R, et al (2005) Contextmapping: experiences from practice. *CoDesign* 1(2):119–149. <https://doi.org/10.1080/15710880500135987>, URL <https://doi.org/10.1080/15710880500135987>
- Wang K, Reimers N, Gurevych I (2021) TSDAE: Using transformer-based sequential denoising auto-encoder for unsupervised sentence embedding learning. In: Findings of the Association for Computational Linguistics: EMNLP 2021. Association for Computational Linguistics, Punta Cana, Dominican Republic, pp 671–688, <https://doi.org/10.18653/v1/2021.findings-emnlp.59>, URL <https://aclanthology.org/2021.findings-emnlp.59>
- Wang K, Thakur N, Reimers N, et al (2022) GPL: Generative pseudo labeling for unsupervised domain adaptation of dense retrieval. In: Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Association for Computational Linguistics, Seattle, United States, pp 2345–2360, <https://doi.org/10.18653/v1/2022.naacl-main.168>, URL <https://aclanthology.org/2022.naacl-main.168>
- Yang Q (2017) The role of design in creating machine-learning-enhanced user experience. In: 2017 AAAI spring symposium series
- Yang Q (2018) Machine learning as a ux design material: how can we imagine beyond automation, recommenders, and reminders? In: AAAI Spring Symposia, pp 2–6
- Yang Q, Scuito A, Zimmerman J, et al (2018) Investigating how experienced ux designers effectively work with machine learning. In: Proceedings of the 2018 Designing Interactive Systems Conference. Association for Computing Machinery, New York, NY, USA, DIS '18, p 585–596, <https://doi.org/10.1145/3196709.3196730>, URL <https://doi.org/10.1145/3196709.3196730>
- Zhang Y, Chen Q, Yang Z, et al (2019) Biowordvec, improving biomedical word embeddings with subword information and mesh. *Scientific Data* 6(1):52. <https://doi.org/10.1038/s41597-019-0055-0>, URL <https://doi.org/10.1038/s41597-019-0055-0>

Citation for this Paper

[Click here for the Online Version](#)

```
@article{zhukova2024,  
  author={Anastasia Zhukova, Lukas von Sperl, Christian E. Matt, Bela Gipp},  
  title={Generative User-Experience Research for Developing Domain-specific Natural Language  
Processing Applications},  
  journal={Knowledge and Information Systems},  
  pages={36},  
  year={2024},  
  month={08},  
  doi={10.1007/s10115-024-02212-5},  
  note={This version of the article has been accepted for publication, after peer review  
(when applicable) but is not the Version of Record and does not reflect post-acceptance  
improvements, or any corrections. The Version of Record is available online at:  
http://dx.doi.org/10.1007/s10115-024-02212-5}  
}
```

Related Papers

1) Anastasia Zhukova, Felix Hamborg, Bela Gipp. ANEA: Automated (Named) Entity Annotation for German Domain-Specific Texts. <http://arxiv.org/abs/2112.06724v1>. 2021.