

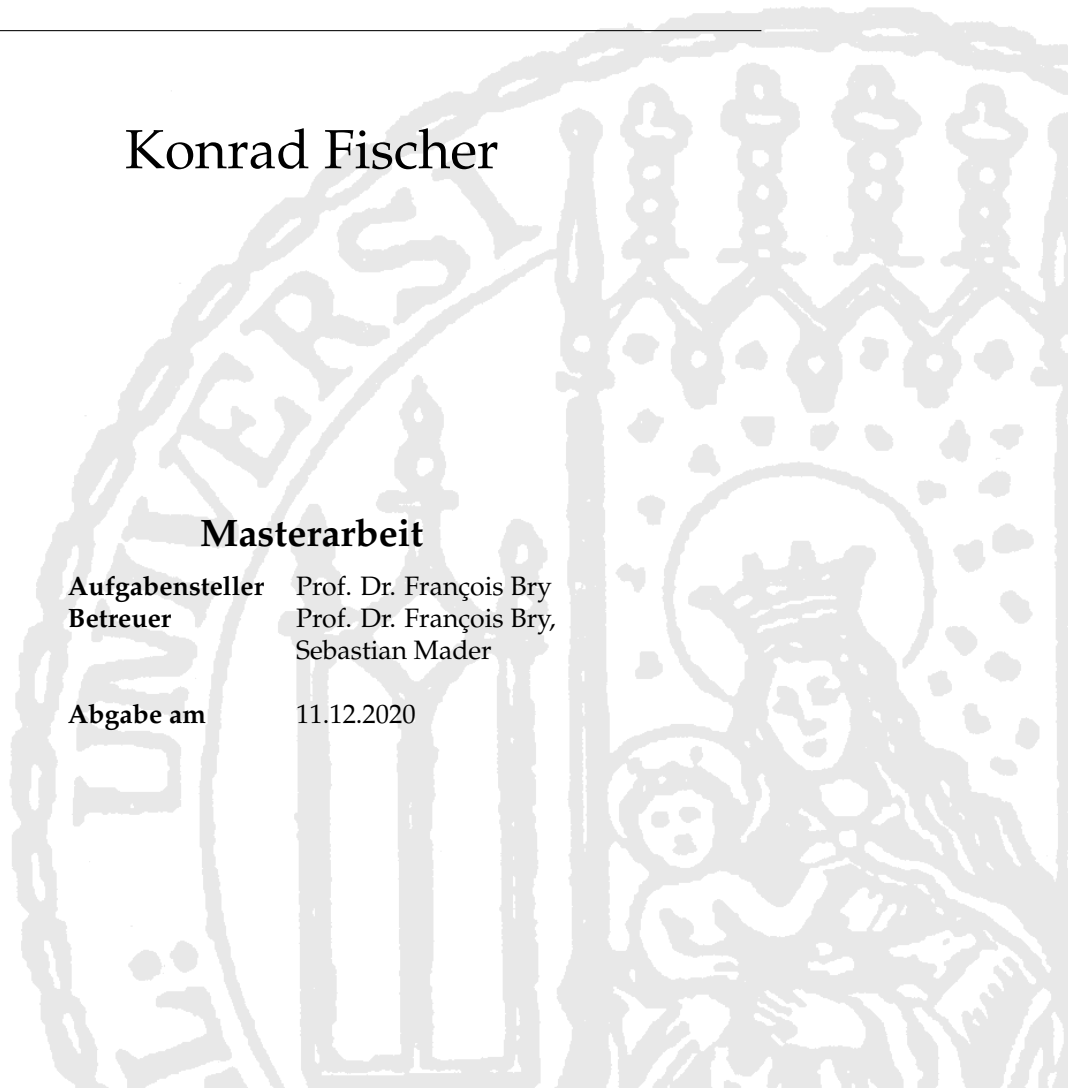
INSTITUT FÜR INFORMATIK
der Ludwig-Maximilians-Universität München

A COLLABORATIVE ANNOTATION EDITOR FOR BACKSTAGE 2

Konrad Fischer

Masterarbeit

Aufgabensteller	Prof. Dr. François Bry
Betreuer	Prof. Dr. François Bry, Sebastian Mader
Abgabe am	11.12.2020



Erklärung

Hiermit versichere ich, dass ich die vorliegende Arbeit selbstständig verfasst habe und keine anderen als die angegebenen Hilfsmittel verwendet habe.

München, den 11.12.2020

Konrad Fischer

Annotations are an important way for learners to interact with digital documents, with other learners, with lecturers, or with tutors. Mader [30, 31] developed the digital learning and teaching system *Backstage 2*, which includes a component for collaborative annotation enabling such interactions. This thesis describes a new annotation system that improves on the one integrated in *Backstage 2*.

As a first step towards this goal, related work on annotations, annotation systems, computer-supported collaborative learning, and cognitive load is summarized. Then data gathered from the usage of *Backstage 2*'s annotation component at *LMU Munich* from 2017 to 2020 is analyzed and design goals for the new system are formulated based on the results of the data analysis. The design goals include that the new system should make it easy to adapt it to different use cases and easy to switch between customizations of the system, that it should foster users' communication awareness, and motivate users to interact with the system by voting on annotations or creating comments regarding them.

The newly built system is then described on a conceptual level with a focus on how it fulfills the goals defined based on the data analysis. More specifically, the functionalities of the system, its design, its architecture, and potential extensions are explained. Regarding the system's functionalities, the different configurations that its interface can be used in are explained and regarding the system's architecture, its modularity is shown in detail and examples of how the system could be adapted to various use cases are given.

Afterwards, the systems functionalities, e.g., the process of creating an annotation, are described from the point of view of a user.

Lastly, the results of a data analysis and the description of the implemented system are summarized and an outlook on potential future improvements, extensions, and use of the system is given.

Zusammenfassung

Annotationen sind für Lernende eine wichtige Art der Interaktion mit digitalen Dokumenten, anderen Lernenden, Dozenten oder Tutoren. Mader [30, 31] entwickelte das digitale Lern- und Lehr-System *Backstage 2*, welches eine Komponente für kollaboratives Annotieren enthält, die solche Interaktionen ermöglicht. Diese Arbeit beschreibt ein neues Annotations-System, welches das in *Backstage 2* enthaltene verbessert.

Als erster Schritt zu diesem Ziel werden bestehende Arbeiten zu Annotationen, Annotations-Systemen, Computer-Supported Collaborative Learning und Cognitive Load zusammengefasst. Anschließend werden Daten aus der Nutzung der Annotations-Komponente von *Backstage 2* an der *LMU München* von 2017 bis 2020 analysiert und Design-Ziele für das neue System werden basierend auf dem Ergebnis der Datenanalyse formuliert. Die Design-Ziele beinhalten, dass das neue System es einfach machen soll, es an verschiedene Einsatzzwecke anzupassen und zwischen verschiedenen Anpassungen des Systems zu wechseln, dass es das Bewusstsein der Nutzer für Kommunikation im System erhöhen soll und Nutzer dazu motivieren soll, durch Abstimmungen über Annotationen und das Erstellen von Kommentaren zu Annotationen mit dem System zu interagieren.

Das neu entwickelte System wird anschließend auf einer konzeptuellen Ebene beschrieben, wobei ein Fokus darauf gelegt wird, wie es die Ziele, die basierend auf der Datenanalyse definiert wurden, erfüllt. Genauer werden die Funktionen des Systems, sein Design, seine Architektur und potentielle Erweiterungen beschrieben. Bezüglich der Funktionen des Systems werden die verschiedenen Konfigurationen, in denen die Oberfläche genutzt werden kann, erklärt und bezüglich der Architektur des Systems wird seine Modularität im Detail gezeigt und Beispiele dafür, wie das System an verschiedene Einsatzzwecke angepasst werden kann, werden gegeben.

Anschließend werden die Funktionen des Systems, z.B. der Prozess für das Erstellen einer Annotation, aus der Perspektive eines Nutzers beschrieben.

Abschließend werden die Ergebnisse einer Datenanalyse und die Beschreibung des implementierten Systems zusammengefasst und ein Ausblick auf potentielle zukünftige Verbesserungen, Erweiterungen und Nutzung des Systems wird gegeben.

Acknowledgments

I would like to thank my supervisor Sebastian Mader for his guidance, patience, and the opportunity to use the data he gathered for his dissertation for my thesis. He provided constant support, feedback, and help when things did not go according to plan. I will miss our insightful discussions.

I would also like to thank Prof. Dr. François Bry for the opportunity to pursue this topic and all the helpful feedback on this thesis throughout the process of its creation.

Thanks to my parents and my sister who supported me throughout the creation of this thesis and throughout my studies in general.

Last but not least, I would like to thank K. for proof-reading this thesis and always being there for me.

Contents

1	Introduction	1
2	Related Work	3
2.1	Definition of Annotations	3
2.1.1	Annotations on the Web	4
2.1.2	Categorizations	4
2.2	Advantages of Annotations	4
2.3	Annotation Systems	5
2.4	Collaborative Assessment	9
2.5	Computer-Supported Collaborative Learning	10
2.5.1	Definitions	10
2.5.2	Theoretical Foundations	11
2.5.3	History	12
2.5.4	Studies	12
2.5.5	Advantages and Challenges	14
2.5.6	Cooperative Learning	15
2.5.7	Additional Resources	15
2.6	Classification of Annotation Systems	15
2.7	Cognitive Load	16
2.7.1	Underlying Principles	16
2.7.2	Types of Cognitive Load	17
2.7.3	Cognitive Load Effects	18
2.7.4	Limitations of Short Term Memory	19
2.7.5	Learning	19
3	Analysis of Previous Annotation Behavior	21
3.1	Backstage 2	21
3.2	Methods	22
3.2.1	Data Gathered	22
3.2.2	Filtering	22
3.2.3	Measures	23
3.2.4	Analysis	24
3.3	Results	24
3.3.1	Overview	24
3.3.2	Annotation Creation Time	25
3.3.3	Annotation Purposes	25
3.3.4	Questions Answered in Courses	25

3.3.5	Questions Answered Over Time Since Start of Lecture Session	25
3.3.6	Aggregation by Page	27
3.3.6.1	Number of Annotations	27
3.3.6.2	Total Text Length	27
3.3.6.3	Annotation Creation Time	29
3.3.7	Individual Annotations	29
3.3.7.1	Comments	29
3.3.7.2	Votes	31
3.3.7.3	Text Length	31
3.4	Interpretation	32
3.4.1	Annotation Purposes	32
3.4.2	Questions Answered	33
3.4.3	Questions Answered Over Time Since Start of a Lecture Session	33
3.4.4	Aggregation by Page	33
3.4.4.1	Number of Annotations	33
3.4.4.2	Text Length	34
3.4.4.3	Annotation Creation Time	34
3.4.5	Individual Annotations	34
3.4.5.1	Comments	34
3.4.5.2	Votes	35
3.4.5.3	Text Length	35
3.5	Threats to Validity	35
4	Concept	37
4.1	Framework	37
4.1.1	Overview Mode	37
4.1.2	List Mode	40
4.1.3	Detail Mode	41
4.1.4	Conclusion	44
4.2	Configurability	44
4.2.1	Annotations	44
4.2.2	Documents	45
4.2.3	Content Editors	45
4.2.4	Filtering	46
4.2.5	Sorting	46
4.2.6	Situational Annotation Filter	46
4.3	Configuration Example: Annotations on Videos	47
4.4	Exchanging Customizations	47
4.5	Interactions with Annotations	48
4.5.1	Comments	49
4.5.2	Votes	49
4.6	Communication Awareness	50
4.6.1	Marking New Annotations	50
4.6.2	Notifications	51
4.6.3	Activity Stream	51
4.6.4	Communication Awareness Regarding Questions	52
4.6.4.1	Notifications	52
4.6.4.2	Visually Highlighting Questions and Showing Their Status	52
4.7	General Extensions	54
4.7.1	Avatars	54
4.7.2	Roles	54
4.7.3	Moderation Tools	54

<i>CONTENTS</i>	xi
4.7.4 Visibility Levels	54
4.8 Client-Server Communication	55
5 Implementation	57
5.1 Overview Mode	57
5.2 List Mode	59
5.3 Detail Mode	63
5.4 Annotation Creation	63
5.5 Document Viewers	66
6 Conclusion	69
Bibliography	71
A Appendix	77
A.1 Attribution for Icons Shown in Mockups	77
A.2 Attribution for Icons Shown in Screenshots	77

CHAPTER 1

Introduction

The current COVID-19 pandemic represents an enormous challenge for societies all around the world. Contact restrictions have been implemented in many countries to curb the spread of the virus. These restrictions have had a substantial effect on the education sector, as they, in many cases, prohibit classroom teaching and other in person meetings among students and between students and lecturers. The modalities of teaching at universities therefore had to undergo fundamental changes in the last few months. Alternatives to classroom teaching and in-person learning in the form of online learning and online teaching formats became more important. This also led to digital documents and videos gaining in importance, as they both represent important components of online learning. Examples of digital documents used in this context are slide decks and digital textbooks, which students use for studying, and assignments, which students submit in digital form. Videos are used in the context of online learning, e.g., in the form of video lectures.

When working with digital documents, annotations are a vital tool of interaction with the documents themselves, with other learners, and with lecturers or tutors. Students can use annotations to aid them in understanding the contents of scientific papers, slide decks, or text books by attaching explanatory notes in the form of annotations to specific parts of these documents. Furthermore, annotations are an important tool when creating a digital document, e.g., a report, as a group of students, as group members can use them to give others feedback on their work. Apart from that, annotations can be utilized to exchange feedback between different authors or groups of authors in peer review processes. Annotations can also be employed to communicate with lecturers or tutors by marking specific parts of scripts, lecture slides, or assignments and attaching a question to the marking using an annotation.

Mader [30, 31] developed the digital learning and teaching system *Backstage 2* that has been in use at *LMU Munich* from 2017 to 2020 and includes a component for collaborative annotations that is capable of enabling the listed interactions with digital documents. The goal of this thesis is to build a collaborative annotation system that improves on *Backstage 2*'s system by providing an adaptable architecture that prepares the system for future extensions, incorporating learnings from four years of usage of the system, and utilizing modern technologies.

Apart from interactions with digital documents, learners could also use annotations to interact with videos, e.g., by attaching explanatory notes to passages of video lectures or even to specific parts of the video frame at a specific time in a video lecture.

In Chapter 2 of this thesis, a survey of related work on, among other things, annotations, annotation systems, computer-supported collaborative learning, and cognitive load is presented, followed in Chapter 3 by an analysis of data gathered from the existing annotation system in *Backstage 2*, which is used to motivate the design goals for the new system. After that, the concepts behind the implemented system and how it fulfills the design goals are explained in Chapter 4 together with potential extensions of the system. The system's implementation is then showcased in Chapter 5. Chapter 6 concludes the thesis by summarizing the results of a data analysis and the description of the implemented system and providing an outlook on potential future improvements, extensions, and use of the system.

CHAPTER 2

Related Work

According to Wolfe [74], annotating is a process that goes back to medieval academics utilizing the margins of documents to store and distribute information, discussing the meaning of the document, and explaining the best ways of studying it. These annotations gained more and more value as a document was passed from one scholar to the next, the author explains, and often contained more information than the primary source. Due to their high value, annotations were often even duplicated along with the content of the document, Wolfe notes.

2.1 Definition of Annotations

On a general level Lin and Lai [28] define annotations as “marks readers make on reading materials” [28, p. 264] with Hwang et al. [21] using an almost identical definition. Ovsianikov et al. [45] are more specific and define an annotation as a “datum created and added by a third party to the original document, which can be a written note, a symbol, a drawing or a multimedia clip” [45, p. 340]. Campbell [4] focuses on the action of annotating which they define as “providing commentary on information objects created at other times and usually by other people” [4, p. 1].

In the context of this thesis an extended version of Lin and Lai’s [28] definition of annotations will be used: An annotation is a mark a person leaves on media they view. This extended definition, unlike Lin and Lai’s [28], also allows for annotations on types of media apart from written text, e.g., on videos, and for the author of an artifact to create annotations, which Ovsianikov et al.’s [45] definition does not cover.

The extensiveness of the idea of annotating and annotations is mentioned by multiple authors. Marshall [36] notes that an annotation is a wide concept and has been interpreted in different ways: “as link making, as path building, as commentary, as marking in or around existing text, as a decentering of authority, as a record of reading and interpretation, or as community memory” [36, p. 40]. Similarly, Wolfe [74] also describes that annotations appear in various forms: They can be “formal”, “informal”, “published”, or “unpublished” [74, p. 473, 476] and can be created for only their creator themselves or to be viewed by others.

2.1.1 Annotations on the Web

In the more specific context of annotations on web pages, Sannomiya et al. [57] note that annotations contain “personal opinions, comments, or impressions on web pages” [57, p. 40] and describe “a facet of the contents of web pages” [57, p. 40]. More specifically, in the context of the *Semantic Web*, Euzenat [10] defines an annotation as the meaning of the document expressed in a formal language that is connected to the document and Kahan et al. [24] as metadata that links additional information to documents. The *Semantic Web* is an evolved version of the *World Wide Web* where documents, e.g., web sites, are enhanced with information that makes them easier to semantically process for computers [3]. Within the description of *Amaya* [51], a tool to view, create, and edit web sites that supports creating annotations collaboratively, the *World Wide Web Consortium (W3C)* [71] defines annotations as “comments, notes[, or] remarks that can be attached to any Web document or a selected part of the document” [51].

2.1.2 Categorizations

Several authors propose categorizations for annotations. Wolfe [74] generally states that annotations can be intended either only for the author or for other people reading the same text. They mention that the typically intended audiences of annotations from a reader’s perspective are the reader themselves, the author of the annotated text, or other people reading the text. Furthermore, an author might create annotations for their readers, Wolfe adds. According to Marshall [36], the dimensions that annotations can be categorized along are their formality, publicness, scope, lifespan, explicitness, level of engagement with the text, and target audience. In another paper Marshall [35] identifies six roles of annotations, namely [35, p. 136f]

- “procedural signals” [35, p. 136f] that represent a plan for future actions,
- “placemarkings and aids to memory” [35, p. 136f] which mark passages as important,
- “in situ locations for problem-working” [35, p. 136f] where the author worked on an issue near the passage that it is related to,
- “a record of interpretive activity” [35, p. 136f],
- “a visible trace of the reader’s attention” [35, p. 136f],
- and “incidental reflections of the material circumstances” [35, p. 136f] that represent the situation the author was in when they read a text.

2.2 Advantages of Annotations

Annotations are not only helpful to convey information to other readers of a text but also aid their author in engaging with a text while reading it and in construction of knowledge [20]. A survey on the usage of annotations on paper among academics found that the most common uses of annotations were sharing commentary on documents with others, aiding the reader’s thought process, recalling information, and storing information that is of use for understanding a document’s content [45].

While annotations on paper already offer numerous benefits to both their authors and readers, they are even more effective and offer more functionalities when used in digital form [45]. One of the problems that occur when working with analogue annotations on paper is that gathering feedback from multiple reviewers of a text is difficult because the amount of feedback that can be left is limited by the space on a page [74]. An additional

challenge is the sending and receiving of feedback on physical documents, one solution to which is for participants to physically meet [74]. Another solution consists in participants sending the document to each other via mail each time they want to exchange opinions. This makes reviewing a text difficult on an organizational level as meetings have to be scheduled to exchange documents, or potentially expensive and lengthy, due to postage costs and delays if mail is used, which might be the only viable option for authors separated by long geographical distances. A way of using analogue annotations on paper while avoiding the disadvantages of mailing documents is scanning annotated documents and sending them to the recipient digitally, e.g., via email. However, this is associated with additional time investment. In general, exchanging feedback using analogue annotations often requires using additional means of communication apart from the one the document is transferred with [20].

When an online annotation software is used, the process of creating annotations and sharing them with other users is simplified [74, 20]. Furthermore, such an application can help to intensify the interaction between the involved parties if it allows for discussions related to the annotations [74]. Another advantage of digital annotation systems is that they often offer ways to search documents and to alert users to events that occurred in relation to their documents [20].

Grbac et al. [16] conducted a focus group study with students from various subjects to evaluate a prototype for an annotation system. The following paragraph refers to Grbac et al.'s findings unless stated otherwise. The participants in the study reported that they preferred learning with and annotating on paper compared to other media. Participants used additional material to augment their primary analogue documents used for studying. They did so by adding the content of the additional material to the main document instead of establishing some kind of link. The inability to create such a link was mentioned as an issue by the participants as well as by Hoff et al. [20]. The most important issue that the students identified with their studying habits was that most of the paper they used was either thrown away, because there was not enough room to store it and doing so would have been effortful, or stored but never retrieved as there was no way to easily find it. They mentioned that they would have liked to have access to documents they discarded in a number of cases. Contrary to their analogue documents, most students in the study stored their digital documents as no room was required to do so and they were conveniently searchable.

2.3 Annotation Systems

Extensive research has been conducted related to annotations. Examples include the development and evaluation of annotation systems and systematic reviews of existing applications.

Hwang et al. [21] created a web application for collaborative annotating and evaluated their tool through a user study among freshman college students. The following paragraph summarizes their findings. The system allows the students to create annotations in the form of "highlighting, underlining, textual annotation, as well as online voice recording" [21, p. 5]. In the study it was used in three different modes that differed in the visibility of the annotations the students created: only visible to themselves, a group of peers, or all their classmates. The results of the study revealed that a majority of the students felt that their reading of digital documents improved through the annotations and that the usage of the application led to higher levels of enjoyment, interest, and better learning outcomes. Students also reported that they felt like they interacted with the meaning of the documents in a better way when using the system. More than half of the participants believed that the mode of the system that allowed them to share annotations with a group boosted

their motivation to learn and two thirds believed that the mode that made annotations visible among all students did so. When the authors evaluated the scores of the participants on tests conducted during the course, they found that participants that had access to the annotation system achieved significantly higher scores in all three scenarios (creating annotations only visible for themselves, visible for their group, or visible for all peers). Also, among the students in the group visibility scenario those that created more annotations got better scores. The magnitude of this effect was smaller when looking at the scenario where annotations were visible for all students, however. The authors found no significant increase in scores on the final examination when comparing the experiment and the control group. They explain this with the motivation to study for the final exam having more influence on the students' performance on that exam than the usage of the annotation tool.

Su et al. [61] developed the collaborative annotation system *PAMS 2.0* and evaluated it within a user study among freshman computer science students. Their work is summed up in this paragraph. *PAMS 2.0* allows for creating and distributing annotations on documents and discussing them among the students. According to the authors, the collaborative use of annotations allows learners to structure documents, add marks to easier locate specific information later, make team members focus on individual parts of the document, take in the knowledge of team members and therefore build up knowledge together, and, in the case of teachers or expert students, contribute to the knowledge pool. The results of Su et al.'s user study revealed that the system improved the students' learning and aided the distribution of annotations among them. Specifically, students believed that the usage of an online annotation application caused an improvement of their reading of digital documents. Furthermore, the majority of participants reported that the annotation software caused an improvement in their performance, enjoyment, and interest when participating in collaborative learning. The students also generally had a positive opinion on the annotation system and were open to using it again. According to the authors, the *PAMS 2.0* learning system improved the students' learning performance in some of the scenarios they tested but not in all. Furthermore, the authors found indications of a positive correlation between the number of annotations a user created and their learning performance which validates Hwang et al.'s [21] results in this regard.

Lebow and Lick [27] introduced *HyLighter*, another application that allows users to collaboratively create annotations on texts by highlighting sections, discussing them, and with the capability to "aggregate or merge annotations [...] from multiple readers and generate composite displays" [27, p. 1]. The authors' system further includes the possibility to compare the annotations of one user to either specific other users or all users in a specialized view. Lebow and Lick claim that *HyLighter*, among other positive effects, enhances users' ability to think critically, their reading capabilities, their motivation and self-control regarding studying, and abilities crucial for working with source material.

Razon et al. [53] evaluated *HyLighter* using two studies among undergraduate and graduate students. Their evaluation is outlined in this paragraph. When evaluating the results of the study among undergraduate students, the authors failed to find a significant improvement in scores on quizzes during the course, positive emotions, or scores on summative assessments. However, Razon et al. note that, on a descriptive level, there were positive effects on these variables. Furthermore, the students using *HyLighter* reported more motivation to learn and "higher levels of excitement, optimism, happiness, motivation to read, and motivation to read more" [53, p. 354]. The authors conclude that this first study shows that *HyLighter* helps to improve students' learning. The second study among graduate students also did not show a significant positive influence of the annotation software. On a descriptive level, however, the authors found that the usage of the *HyLighter* system led to "improved learning comprehension" [53, p. 357]. Furthermore, the system did not improve students' self-efficacy, anxiety levels related to tests, or their striving towards their mastery goal. "Mastery goals are aimed at attaining a standard of competence defined by self-

improvement or skill development” [50, p. 37]. Self-efficacy is “people’s domain-specific perceptions of their ability to perform the actions necessary to achieve desired outcomes” [11, p. 314].

Lin and Lai [28] developed the *CAFAS* system that allows students to view annotations created by other students explaining the solutions to formative assessment tests or to create such annotations themselves. They then evaluated that system within a user study among students of various business-related subjects. Their findings are summarized in the following paragraph. The mean scores of the students using Lin and Lai’s system on an exam after the first half of the course and on an exam at the end of the course were significantly higher compared to the scores of the students that did not use the *CAFAS* system. Furthermore, the participants in the experiment group spent significantly more time reexamining the learning material before the first and second exams. The authors interpret this as evidence that the experiment group was able to extract a significant amount of knowledge from the collaborative annotations. The results of a questionnaire and interviews revealed that the system motivated students to devote more of their time to working with the application and that they perceived the system as helpful in delivering insights into the material. Some participants voiced concerns regarding the added burden of annotation creation when learning, however.

Yeh and Lo [76] developed the *Online Annotator for EFL Writing* system to provide feedback to and investigate the errors of second language students. EFL hereby stands for English as a Foreign Language. The system allows teachers to correct students’ writing and is able to divide the marked errors into classes, the authors explain. Yeh and Lo conducted a study among Applied Foreign Languages freshman students studying EFL where feedback on their English writing was provided to the experiment group using the digital annotation system and to the control group using paper. The results of the study, as presented by the authors, revealed that the students using the digital annotation system were significantly better at correctly noticing errors in written texts. The authors see this as confirmation that using their application has a positive influence on the capabilities of EFL learners to rectify errors. They explain this influence with their system offering a number of functions that were not accessible when using paper to give feedback and which jointly established “a more interactive environment for EFL error correction and corrective feedback” [76, p. 890]. Yeh and Lo mention that a component that allows students to view annotations created by their peers would be a good addition to their system.

Nokelainen et al. [44] created the *EDUCOSM* tool for shared annotations and evaluated it in a study. Both their tool and their findings are described in the following paragraph. The system includes, among others, features that allow for creating, searching, and viewing shared annotations on documents as well as discussing them. Among other tasks, participants of the study had to create annotations on a document and debate the document with another student. Results of the study showed all participants reported that the application improved their learning technique, had a positive impact on their study routine, and that they would endorse the usage of the application in other classes. Furthermore, participants rated annotations they created themselves as more helpful than those created by others. Apart from that, the authors interpret from their results that students that are motivated to work with the system and can describe their actions to others are the ones that achieve the best learning outcomes and create the best annotations.

Grbac et al. [16] developed a prototype of a hybrid solution that allows the user to digitally interact with annotations created on paper. Their contribution is described in this paragraph. The user needs a mobile phone, which is mounted to a stand, and a laptop to use the system. The phone is used to acquire video footage of the document while the laptop is used for handling that footage. The system offers two views to the user, both shown via the laptop’s screen: In one view the user is shown the document they are currently interacting with and annotations either created by themselves or created by

other users. The second view shows a virtual version of the document reflecting the real document placed under the mobile phone's camera. In this view the user can use their pen to interact with the virtual document and annotations, either created by the user themselves or others, that are displayed together with the document. Annotations can be created on the physical document and then the authors' system is able to convert them into digital annotations based on images of the physical document. Other interactions that the system supports are searching for text, which is overlaid by annotations, on the Internet and adding links to web pages to the document. The authors asked members of a focus group study to provide feedback on their prototype. The participants did not have any issues with making annotations they had created available to others, but some doubted whether their annotations would be useful to other students. Furthermore, the students reported that they did not use others' annotations as they were but merged them with their own. Participants in general had a positive opinion of the prototype the authors created and did not see issues related to privacy or interference with others. One concern was, however, that the prototype was too inconvenient to use to bring it to courses, but that disadvantage was seen as less severe when studying for exams. When asked which features were missing from the prototype, participants wished for the ability to link to a specific part of a web site or a video.

Glover et al. [13] developed a web-based annotation application that augments web pages by enabling students to add annotations and remarks to them or to change their content. These additions and changes can then be made accessible to other students, according to the authors. They further explain that the system allows the user to either view only the annotations and changes they created themselves, those created by other users, or the original page.

Hoff et al. [20] created a scheme to categorize annotation software, identified a number of problems with annotation systems available at the time of writing, and proposed a new annotation model to solve them. The first issue the authors mention is that a user typically needs to use different annotation software depending on the document format and media type they want to annotate. This represents a disadvantage compared to annotating documents on paper as the same writing utensils can be used to write on printouts of various media, e.g., images or text, Hoff et al. point out. According to the authors, however, the majority of annotation applications do neither work with different file formats nor different media. Apart from that, the authors criticize that most annotation software does not support the recording of connections between different documents despite working with multiple documents being an integral part of the study routine of many students. They note that most web-based annotation applications allow for the usage of hyperlinks to create a connection to another document but that these links are not powerful enough as they only allow for references to certain kinds of documents, only support linking to the entire resource, and only create a reference pointing in one direction. Another issue the authors see with most annotation systems available at the time is that they do not offer extensive enough ways to share annotations among collaborators. They specifically criticize the absence of group-based access restrictions in systems that only differentiate between annotations visible to the author and those visible to everyone, while other systems do not even include that differentiation. Furthermore, the authors identify inadequate representation of annotations' contexts as another issue. They state that the context an annotation exists in is often required for the information stored in that annotation to be usable and criticize that most annotation software loses some or all of that context. To solve the identified problems, the authors then propose an annotation model.

Furthermore, a number of annotation systems and frameworks [57, 10, 4] have been developed for the *Semantic Web*.

One example of a *Semantic Web* framework is *Annotea* that Kahan et al. [24] describe. The following paragraph refers to their concept. Within the concept of the *Semantic Web*,

Annotea's role is to allow users to connect metadata to resources, e.g., web pages. *Annotea* sees annotations as a form of metadata and is based on the resource description framework (RDF) as a data storage, which is a "flexible framework for describing properties of any Web resources" [24, p. 593]. In its role as a framework, *Annotea* is not built for a particular client but only specifies how annotations are attached to documents and how they can be retrieved from the RDF data storage. Some of *Annotea*'s noteworthy features are that it allows for an optional differentiation between private and public annotations, that it enables users to define their own categories which annotations can be assigned to, and that it does not assume the existence of a central annotation server but instead is able to interact with multiple servers that the users can post their annotations to and download and view others' annotations from. After a user has sent an annotation to a server, it then becomes available to all users with access to that server. One of the features that the authors would like to integrate into the *Annotea* framework in the future are ways for the users to debate annotations.

Kahan et al. [24] also present an implementation of their framework based on *Amaya* [51], which is a tool created by the W3C to build, change, and view web sites that also contains a component that allows users to collaboratively work on annotations. Kahan et al.'s implementation makes use of *Annotea*'s support for a distinction between private and public annotations. Therefore, after an annotation is added by a user it initially only exists on their machine, the authors explain, and the user can then decide to publish the annotation to an annotation server. Apart from the features specified in the *Annotea* framework, Kahan et al.'s application also provides the possibility to filter annotations by different criteria.

Another example of a *Semantic Web* system is Yang et al.'s [75] application that they built to showcase the metadata models they developed which allow them to "formally describe content and annotation" [75, p. 79]. The authors list creating annotations, attaching them to documents, examining them, debating them in real time, and clustering them as features of their system. It processes the data entered by the users to build metadata models that characterize the annotations in a formal way and is then able to infer ontological connections between these models, according to the authors. Using the models, the system can then, among other features, offer "query-by-annotation for semantic search" [75, p. 80], Yang et al. explain. In the context of the *Semantic Web*, ontologies are "metadata schemas, providing a controlled vocabulary of concepts, each with explicitly defined and machine-processable semantics" [32, p. 72], according to Maedche and Staab.

In addition to the mentioned annotation systems, some authors [15, 2] explored digital ink annotations.

2.4 Collaborative Assessment

Macdonald [29] conducted a case study to explore the opinions of college students and teachers on classes that assessed students online in a collaborative way. They found that the number of participants in collaborative online exercises is higher if they are connected to some kind of grading and that the presence of such a grading component might cause an increase in the extensiveness and excellence of discussions. According to the author, working collaboratively, apart from offering interaction with others, can be beneficial for students regarding their advancement and building of expertise. Macdonald further reports that the students' opinions on collaborative grading were less positive when it implicated their success depending on peers.

2.5 Computer-Supported Collaborative Learning

When used for educational purposes, collaborative annotation applications can be seen as an instance of a computer supported collaborative learning (CSCL) system. CSCL will be defined and examined further in the following section.

2.5.1 Definitions

According to Alavi [1], collaborative learning in general “involves social (interpersonal) processes by which a small group of students work together (i.e., cooperate and work as a team) to complete an academic problem-solving task designed to promote learning (i.e., get actively involved and participate in problem solving)” [1, p. 161].

Gokhale’s [14] definition of collaborative learning is more focused on the composition of the group: “[A]n instruction method in which students at various performance levels work together in small groups toward a common goal” [14, p. 22]. According to the author, this leads to the learners having a responsibility for the learning of their peers as well as for their own, thereby supporting each other when succeeding. Gokhale sees improving students’ abilities with regards to critical thinking using collaborative learning as an important objective of education in technology-related subjects.

Dillenbourg et al. [9] define collaborative learning as “a variety of educational practices in which interactions among peers constitute the most important factor in learning, although without excluding other factors such as the learning material and interactions with teachers” [9, p. 3] and clarify that the term ‘computer supported’ “not only [refers] to connecting remote students but also to using technologies to shape face-to-face interaction” [9, p. 3]. They elaborate that for them a CSCL application is not only a way to aid interaction of students situated far apart physically but can also be used in traditional and remote education for guiding verbal exchanges and for real-time evaluation, display and recording of these exchanges. The goal the authors define for CSCL applications is to provide an environment that fosters effective exchanges between members of a group. They further point out that collaboration in itself does not lead to better learning but communication within the groups does, which Dillenbourg and Fischer [8] reassert.

Stahl et al. [59] more specifically define CSCL as “collaborative learning that is facilitated or mediated by computers and networked devices” [59, p. 479], adding that it can happen “*synchronously*, with learners interacting with each other in real time (e.g., a chat room), or *asynchronously*, with individual contributions stretched out over time (e.g., an e-mail exchange)” [59, p. 479, emphasis in original]. According to Stahl et al., the characteristic feature of CSCL is the usage of electronic devices communicating over a network for learning, not that the learners are situated far apart physically. Collaboration in general is mainly understood as “a process of shared meaning construction” [59, p. 487] that results from interactions, according to the authors. Therefore, in the CSCL view, as the authors describe it, learning happens when multiple parties work out meaning together, not when they study on their own. The aim of CSCL, as reported by Stahl et al., is, based on this understanding of learning, to create tools that support this way of constructing meaning.

The term of CSCL has to be distinguished from other educational concepts, Stahl et al. [59] state. According to the authors, one of these is e-learning, which they define as “the organization of instruction across computer networks” [59, p. 480] and describe as often aiming at getting learning material to large groups of students without the instructor having to continuously interact much with the learners. Stahl et al. criticize that distribution of material does not equal good teaching and that e-learning is more demanding regarding the time the instructor needs to spend than traditional approaches. CSCL goes further than e-learning by laying a focus on the learners collaborating with each other instead of them just consuming educational content on their own, the authors explain. This communica-

tion between students is where learning occurs and the use of computers and computer networks to facilitate it is of great importance for the CSCL approach, Stahl et al. point out.

Furthermore, Stahl et al. [59] highlight the differences between collaborative and cooperative learning. According to them, in a collaborative context the learners build knowledge together in a social way, whereas in a cooperative context the learners study individually first and then merge the outcomes of their work into a group artifact. Therefore, the authors elaborate, the type of learning activity differs between collaborative and cooperative learning, where the learners, despite being part of a group, learn individually. Collaborative learning can also contain components that involve individual studying but, Stahl et al. emphasize, always includes more than that.

For this thesis, Stahl et al.'s [59] definition of CSCL combined with Dillenbourg et al.'s [9] definition of collaborative learning will be used as Stahl et al. are rather vague in that regard. The definitions are the following: CSCL is "collaborative learning that is facilitated or mediated by computers and networked devices" [59, p. 479]. Collaborative learning takes place if interactions between learners are the most important aspect of the learning that happens during a certain activity, although other aspects like "the learning material and interactions with teachers" [9, p. 3] can also play a role [9].

2.5.2 Theoretical Foundations

One of the theoretical roots of CSCL is Vygotsky's description of the connection between students' learning and collaborative learning that occurs within a group [59]. Hereafter, a short synopsis of parts of Vygotsky's works [68, 69, 70] based on Stahl et al.'s [59] representation of [68], Gokhale's [14] of [68], and Warschauer's [72] of [68, 69, 70] will be given. Stahl et al., Gokhale, and Warschauer explain that Vygotsky sees a difference between the potential improvement that learners can achieve when working alone compared to the one they can reach when working collaboratively, with the latter being larger [59, 14, 72]. As a measure of this difference, Vygotsky, as summarized by Stahl et al. and Warschauer, defines the zone of proximal development [59, 72]. Vygotsky, as presented by Warschauer, emphasizes that collaborative learning is important in allowing students to progress through their zone of proximal development [72]. This collaborative learning can take place between a student and an instructor or between peers [72]. One of the reasons for this effect consists in different experiences and knowledge in a group improving learning, according to Vygotsky as summarized by Gokhale [14]. More generally speaking, the importance of collaborative learning stems from Vygotsky's view, as described by Warschauer, that complex human abilities evolve from social interaction between humans utilizing language [72]. Regarding the construction of studies for CSCL systems, this leads to tests of the performance of individuals when working alone not allowing for inference about learning that occurred during collaborative work [59].

Alavi [1] lists three properties of processes that allow for effective learning [1, p. 161]:

- "Active learning and construction of knowledge" [1, p. 161] as students can learn most effectively when obtaining, creating, changing, and organizing knowledge
- "Cooperation and teamwork in learning" [1, p. 161] as interactions with group members present opportunities to receive feedback, which can be used to refine and improve one's ideas, as well as to receive social backing
- "Learning via problem solving" [1, p. 161] as the mental models that students build are put to the test and improved in such situations

The author further explains that the concept of collaborative learning, according to their definition introduced in the previous section, is built on these three properties and im-

proves learning as it offers learners opportunities to use, validate, harden, and enhance their mental models using interactions with peers when working on problems.

2.5.3 History

According to Stahl et al. [59] and Dillenbourg et al. [9] the development of CSCL started in the 1990s as a response to applications that were aiming to improve students' learning, but only in isolation instead of utilizing collaborative learning. Stahl et al. [59] mention the ENFI Project at Gallaudet University, the Computer Supported Intentional Learning Environment (CSILE) Project at the University of Toronto, and the Fifth Dimension (5thD) Project at the University of California San Diego as the three initial projects that shaped the development of CSCL [59, p. 482]. These projects all aimed at helping learners with creating more meaningful and valuable content, the authors note. To do so, Stahl et al. explain, they all utilized computers in combination with innovative ways for students to interact with each other on a social level within the context of learning. Dillenbourg et al. [9] see CSCL's second development phase from 1995 to 2005 as during this time research on CSCL became more established as "it acquired its own conference cycle, book series, society and journal" [9, p. 4]. At the date of publication in 2009, Dillenbourg et al. predicted that the third development phase of CSCL from 2005 on would consist in CSCL fading as a separate topic with collaboration being introduced into more environments.

2.5.4 Studies

Gokhale [14] conducted a study among college students with the aim of comparing the effectiveness of collaborative learning and individual learning in improving the ability of students to think critically and to solve drill-and-practice problems. The following paragraph refers to Gokhale's work. They define individual learning as "[a]n instruction method in which students work individually at their own level and rate toward an academic goal" [14, p. 23]. The results of their study revealed that participants that studied using collaborative learning achieved significantly higher test scores on questions requiring critical thinking compared to participants that used individual learning. There was no difference in performance between the two groups on drill-and-practice questions, however. Additionally, the participants reported that working in groups energized their thinking, made it easier for them to comprehend the learning material, and made them feel less fear of solving problems as the responsibility was distributed among the group members, according to the author. Some participants complained about the time investment necessary to help group members understand the learning material, however. Gokhale sees the interaction with peers within the groups, which the students solved the tasks in within the collaborative scenario, as one of the reasons for the performance improvements when using collaborative learning as it allows participants to profit from each other's knowledge, background, and abilities. Furthermore, they believe that the need of explaining and defending one's opinions against group members, which was introduced with the group work component, had a positive effect on students' learning. The author sees an informal environment as a catalyst of discussions and exchanges and notes that their results concur with learning theories that address collaborative learning. Gokhale concludes that collaborative learning supports students in thinking critically as it encourages discussing and explaining thoughts and examining those of others. Based on the results of their study, they see collaborative learning as having a particularly positive effect on critical thinking of students.

Alavi [1] conducted a study investigating whether the use of computer-based group decision support systems (GDSSs) has a positive effect on the impact of collaborative learning on students. This paragraph summarizes their publication. The author describes a GDSS as a tool that aims to improve the work process of groups that are working together at the

same location by limiting the disadvantages of group work, compared to group members working on their own, and amplifying the advantages. The main fields that GDSSs help groups in are processes and tasks, according to Alavi. The results of the author's study show that participants that were using a GDSS reported more improvement in abilities, being more interested in learning, better learning, and in general a better learning experience both in their groups and in class than participants that did not use such a system. Evaluation of participants' comments validated the author's claim that interacting with different viewpoints improves learning through creation and improvement of mental models. Furthermore, participants that worked with a GDSS achieved significantly higher final scores in the course, although there was no significant change in scores on a test conducted halfway through the course.

Capdeferro and Romero [5] conducted a study among master students on the annoyances that students experience in connection with CSCL and their origins. The following paragraph describes the study's results. The authors found that students rarely experienced annoyances connected to CSCL and that their emergence influences the students' feelings and studying between highly and moderately. Capdeferro and Romero did not find a significant influence of annoyances related to CSCL on the participants' content with their master programmes. The causes of annoyance most frequently reported by participants were (in descending order) [5, p. 31f]:

- Inequality regarding dedication, labor, and responsibility among group members when looking at specific tasks and generally looking at CSCL
- Differences in objectives among peers
- Communication issues
- Members missing features that are important for building social connections
- Mismatch in value of the artifacts contributed
- Unbalanced distribution of work within the group
- Grading of individual group members
- Teacher position

Based on their findings, Capdeferro and Romero present recommendations for education to avoid the mentioned issues. They mainly recommend that students should be informed about the educational methods, specifically the usage of CSCL, and that they should be trained in working in teams and social interactions for better collaboration. Additionally, the students should be able to estimate the effort a course requires, adjust their expectations accordingly, and act responsibly. Apart from that, the institution should ensure that the assessment of students is consistent. The teacher should be an active part of the collaboration within the course and should know when to interfere in the CSCL process and how much.

Pérez-Sanagustín et al. [47] conducted a case study on using CSCL and mobile technologies to help new students with getting to know their university. The authors created three stages of activities for the students: one where information on parts of the campus was provided on-site using "Near Field Communication/Radio Frequency Identification (NFC/RFID) technologies" [47, p. 2] and mobile phones or via other media, a second one where students collaboratively created a presentation on their information gathered in the first step utilizing computers, and an evaluation of the whole process [47]. At this point it has to be noted that the phones used in the study differ from the smartphones common at the time of writing in that they have significantly less capabilities. The on-site information distribution in the study was executed using tags with information about specific

facilities on them, including videos, images, and audio, which could be viewed using the mobile phones given to the students using NFC/RFID [47]. Apart from using this interactive approach in the first stage of the activities to get to know the campus, students could also choose to do research on the university's website or gather information from banners placed on campus and potentially ask passersby [47]. The two groups using these ways of acquiring information were together considered the control group, while the students that used the mobile phones represented the experimental group [47]. Results of the study revealed that members of the experimental group received marginally better scores for the presentations they created and were able to describe the campus in a more accurate and authentic way [47]. Despite some issues of technical nature and issues related to using the system utilizing mobile phones, participants favored the activity using the phones [47]. The authors conclude their interpretation of their results by claiming that the interactive experience they designed using mobile phones led to "significant and motivational learning benefits to students" [47, p. 12].

2.5.5 Advantages and Challenges

Kreijns et al. [26] note that one of the advantages that CSCL applications offer is that they allow learners to work together despite being located far apart geographically and to work on their assignments whenever they want. According to the authors, this leads to a transition from traditional learning groups to "asynchronous distributed learning groups (DLGs)" [26, p. 8]. Kreijns et al. also mention that advanced CSCL applications can be utilized to aid in the application of educational plans and patterns and therefore in relieving teachers to some extent. This can make education more cost-effective, the authors note. They further highlight the importance of social interaction for learning and claim that teachers have observed a positive connection between social communication and the results of learning and the satisfaction of learners.

In another publication, Kreijns et al. [25] identify two major issues with CSCL. The first one they list is that it is assumed that social communication will take place simply because the systems offer the opportunity for it. Apart from that, the authors also criticize that social communication is often limited to learning, and therefore, the significance of the social side of such exchanges is overlooked. Such interactions, despite them not being directly connected to the subject, are critical for building a community among the learners, according to Kreijns et al. They note more specifically that although these interactions can work in the context of assignments they will work more likely in other contexts. Therefore, the existence of contexts apart from the assignment context has a positive effect on the development of a community among the learners, Kreijns et al. claim. One of the concrete improvements the authors suggest to avoid the mentioned issues is integrating interactive elements into CSCL applications as no social communication is possible without them. They point out, however, that these elements allow for strategies that foster the forming of communities but do not trigger it themselves. Furthermore, they suggest to focus on groups of learners instead of individuals and on the students instead of the instructor and conclude with stressing that social exchanges are critical for the success of CSCL.

Dillenbourg and Fischer [8] support the first issue Kreijns et al. found as they also point out that constructive social exchanges do not happen by themselves but that the CSCL application must be built to cause them.

Dillenbourg et al. [9] note that CSCL applications can inhibit learners' motivation to interact with others, e.g., by leading to annoyances or doubts regarding the learners' capabilities.

2.5.6 Cooperative Learning

Although cooperative and collaborative learning are not equivalent, research on cooperative learning is also worth noting.

Johnson and Johnson [22] suggest that instructors should conceptualize their students' learning to, among other goals, prevent close-minded approaches to dealing with challenges. To achieve this, the authors recommend using the coherent learning approaches of "cooperative learning and academic controversy" [22, p. 51]. They define that "controversy exists when one student's ideas, information, conclusions, theories, and opinions are incompatible with those of another, and the two seek to reach an agreement" [22, p. 52]. Johnson and Johnson stress the educational advantages of controversies namely better learning performance, social connections of higher quality, and increased social skills. The authors see a cooperative environment as critical for controversy to have a positive effect on learning instead of a negative effect that can occur in a competitive environment.

Johnson et al. [23] conducted a study among eighth-grade students comparing cooperative computer-supported learning to individual studying. They found that computer-supported cooperative learning leads to [23, p. 668]

- improved skills regarding utilizing facts to solve problems,
- better comprehension of facts,
- more and more valuable accomplishments per day,
- and improved competence regarding using facts in examinations.

Additionally, the authors found that learners which participated in cooperative learning were better at working out problems compared to those that studied on their own or in a competitive environment. Apart from that, participants in the cooperative environment also solved the tasks more precisely and needed less time to solve them, Johnson et al. report.

2.5.7 Additional Resources

Dillenbourg and Fischer [8] give an overview over research regarding CSCL and show a number of concepts with importance to education beyond the use of CSCL in vocational education that they extracted them from. Magnisalis et al. [33] surveyed research regarding "design and impact of adaptive and intelligent systems for collaborative learning support (AICLS) systems" [33, p. 5]. Gillani and Eynon [12] investigated massively open online courses (MOOCs), specifically the patterns of communication within MOOCs, whether there was a connection between students taking part in exchanges in forums and their grades, and the demographics of the participants. Based on their results, they then present recommendations for the creation of MOOCs and their position within education. Hämäläinen and Vähäsantanen [18] discuss the concepts behind collaborative learning and creativity and ways of encouraging creativity and coordinating collaborative learning. The authors also show effects of the theory on practice. Strijbos [60] examined grading of CSCL, specifically problems with existing concepts, the subject of such assessments, and ideas for the future of grading in CSCL.

2.6 Classification of Annotation Systems

According to the definitions selected for this thesis from Dillenbourg et al. [9] and Stahl et al. [59] in section 2.5.1, computer-based annotation systems are an instance of CSCL if

the goal of the annotation system is to educate its users and if the interactions among the learners that use the system are the most significant part of the learning experience as all of the systems trivially fulfill the requirement of using networked computers.

Therefore, annotation systems that only serve purposes apart from education [24, 75], e.g., annotation systems that are used to add information to the *Semantic Web* like Kahan et al.'s [24] implementation of *Annotea* using *Amaya*, cannot be categorized as CSCL systems. Among the applications that serve an educational purpose, those that do not focus on interaction among the learners, like Yeh and Lo's [76] *Online Annotator for EFL Writing*, which allows teachers to create annotations that can then be viewed by students, can also not be classified as CSCL applications. Grbac et al.'s [16] system is an edge case here because although it allows for sharing of annotations and the authors examine the utility of this functionality in their study, the main feature of the system seems to be the management of annotations without sharing them. A number of systems that fulfill both criteria [61, 21, 27, 28, 44, 13], e.g. Lebow and Lick's *HyLighter* [27] that allows students to add annotations to web pages and view and comment on other's annotations, remain as CSCL systems. These systems add to the body of research on CSCL and strengthen the argument for its positive contribution to education.

As shown in this chapter, annotation systems have been used in a wide variety of contexts, many of them being of educational nature, and in many cases have been positively evaluated in these contexts. Therefore, this work aims to build a system that is configurable and versatile to support as many of the various applications of annotation systems as possible. Its peer review functionality should allow learners on a meta level to improve their skills in formulating criticism of others' work and in defending their position in a discussion. On the subject level, the system should help students to improve their domain knowledge by receiving feedback on their works. As the goal of the application is to improve students' learning through interactions with other students, it can be considered a CSCL application according to the definition used.

2.7 Cognitive Load

Sweller [63, 64, 62] has introduced and refined the concept of cognitive load that will be explained in this section.

2.7.1 Underlying Principles

Unless stated otherwise, the following section refers to Sweller [62].

Sweller presents a number of principles that their cognitive load theory is built upon. One of these is the information store principle, which states that long-term memory is of central importance for the human mental system as it keeps the large quantities of information that humans use for the majority of their activity. In this view presented by Sweller, someone that is perceived as highly skilled in a certain profession has gathered a large amount of knowledge regarding that profession in their long-term memory. Therefore, according to Sweller, "competence is domain specific" [62, p. 47].

Another principle, which is of importance in Sweller's theory, is the borrowing and reorganizing principle. It states that when humans acquire knowledge from others they merge it with the existing knowledge in their long-term memory and then keep the resulting restructured knowledge instead of just duplicating the input. Put another way, humans keep schemas of information in their memory instead of strict duplicates of the information, according to Chi et al. [6]. In a previous paper, Sweller defines a schema as a "cognitive construct that organizes the elements of information according to the manner with which they will be dealt" [63, p. 296]. The borrowing and reorganizing principle is of

great importance because, according to Sweller, humans get most of the knowledge in their long-term memory from others.

The next base principle of Sweller's theory is the randomness as genesis principle. They argue that the information that humans receive based on the borrowing and reorganizing principle is initially created when trying to find a solution to a new problem for which the long-term memory does not contain a blueprint solution. In such a situation the only possible way to move forward is to randomly create potential solutions and verify whether they are effective as all other procedures require some form of existing information on the problem.

Furthermore, Sweller explains the narrow limits of change principle, which states that the human mind should keep the amount of new information that it has to process at the same time low. They base this advice on research [40, 48] showing that the storage capacity in working memory is relatively low as is the durability of that storage. Sweller further points out that knowledge that the working memory cannot handle because of its size can then not be moved to long-term memory and therefore no learning can take place.

Apart from that, Sweller also introduces the environmental organizing and linking principle. It consists in the working memory choosing, based on external information, the parts of long-term memory that are directly important at the moment. The limitations regarding the working memory's ability to store large quantities of information and to store it over time do not exist anymore when it handles knowledge from long-term memory, according to Sweller. Therefore, long-term memory enables humans to engage in activities that would not be conceivable if they only had short-term memory at their disposal, the author adds. They further state that such information can be activated by external influences.

Sweller connects their principles in the following way: "The information created by the randomness as genesis and narrow limits of change principles, transmitted by the borrowing and reorganizing principle and stored by the information store principle, can be used to determine action that is appropriate to a particular environment" [62, p. 56f].

2.7.2 Types of Cognitive Load

Apart from the underlying principles, Sweller [62] also introduces different notions of cognitive load. One is intrinsic cognitive load which comes from complexity rooted in the features of the information itself, the author explains. "It can only be changed by changing what is learned or by changing the knowledge levels of learners" [62, p. 57]. The second type of cognitive load the author introduces is extraneous cognitive load which originates in the method that information is conveyed by and can be decreased by changing that method.

The third type of cognitive load, which Sweller [64] introduces in a previous paper, is germane cognitive load. Germane cognitive load "refers to the working memory resources that the learner devotes to dealing with the intrinsic cognitive load associated with the information" [64, p. 126]. This kind of cognitive load does not represent an origin of working memory load on its own but only through its dependence on intrinsic cognitive load, according to the author [64]. If extraneous cognitive load grows, the germane cognitive load declines, and therefore, the two are complementary [64]. The total cognitive load that a learner is subjected to is only made up of intrinsic and extraneous cognitive load [64].

Intrinsic and extraneous cognitive load are directly influenced by element interactivity [64]. Element interactivity, as defined by Sweller, describes the dependence of units of information on others [64]. Content with low element interactivity consists of units that can be processed without needing many other units while high element interactivity content is made up of units that require other units to be processable [64]. For a specific assignment both the intrinsic and extraneous cognitive load is based on the element interactivity [64]. In the case of intrinsic cognitive load, the element interactivity is a result of the features of

the information itself, resulting in a certain number of units having to be handled and in case of extraneous cognitive load of the number of units that have to be handled caused by the way the information is conveyed [62]. This element interactivity then decides the level of intrinsic or extraneous cognitive load [62].

Based on the notions of extraneous and intrinsic load, Sweller presents a number of effects that can be used to improve learning [62, 64]. While the full list of effects can be found in [62, 64], those with potential applicability to the annotation system that is the subject of this thesis will be introduced in the following section.

2.7.3 Cognitive Load Effects

The following subsection entirely refers to Sweller [62]. One of the cognitive load effects Sweller covers is the *isolated elements effect* which consists in certain educational methods potentially lowering the intrinsic cognitive load in specific situations. Such a situation exists if the learning material has a high level of element interactivity causing high intrinsic cognitive load, so high that there are too many components to fit into working memory. The learner is then unable to continue until they have processed enough of the material to be able to merge connected components into schemas, which can then be dealt with as if they were a single component, as stated by the environmental organizing and linking principle. To avoid this issue, Sweller suggests that it could be beneficial to show the learner the components of the material in isolation to allow them to be handled in separation, despite the learner not being able to grasp the whole material yet.

Furthermore, Sweller defines the *split-attention effect*. It is associated with the *worked example effect* which consists in potentially improved learning when a puzzle and the solution for that puzzle are examined by the learner instead of them finding a solution themselves. The positive impact of the *worked example effect* is diminished if the example requires the learner to split their attention between different components of the material, which often need to be found among irrelevant components first. This leads to an increase in extraneous cognitive load as more elements need to be processed. Therefore, the author believes physically consolidated worked examples to be preferable compared to split variants. They point to an array of research supporting this claim, a summary of which would have gone beyond the scope of this thesis. Sweller further highlights that the *split-attention effect* is of importance for all other ways of teaching beyond worked examples. For the *split-attention effect* to be applicable, the learner needs to be unable to understand the components of the material separately, they point out, and therefore a consolidation of the material has to be required.

Another effect Sweller introduces is the *redundancy effect*, which can appear when components that are not necessary are added to the learning material. Adding such elements leads to more connections between elements and therefore to element interactivity connected to extraneous cognitive load. The *redundancy effect* does not differ much from the *split-attention effect* but is still separate. While with the *redundancy effect* the components of the material can be understood on their own, the components that the *split-attention effect* deals with have to be merged before they can be comprehended.

Additionally, Sweller presents the *transient information effect* which is another effect related to extraneous cognitive load. According to the author, “[i]nformation is transient if it disappears with the passage of time” [62, p. 71]. The effect can emerge if learners have to recall material with high element interactivity, as they can no longer access it, and then merge it with material they are absorbing at the moment. These activities can together lead to an overburdening of the working memory.

Sweller also defines the *modality effect* which is associated with the *split-attention effect*. It can appear when the learner is presented with two components that they are not able to comprehend separately. In this case, the components can be conveyed in different forms,

e.g., one as a text and one as an audio clip, to improve learning. Such delivery in different formats should lead to a boost in effective working memory and therefore to a decline in cognitive load.

A special cognitive load effect is the *element interactivity effect* as it can limit the impact of the other effects. It states that in situations where the element interactivity that is connected to intrinsic cognitive load is low to begin with, other element interactivity that is connected to extraneous cognitive load might only have limited influence on learning. The reason that Sweller gives for this effect is that in such situations the learner might be able to handle the information they are being presented with without reaching the limits of their working memory. Therefore, other effects, including those previously introduced, might not occur if the element interactivity of the learning material causing intrinsic cognitive load is already relatively low.

2.7.4 Limitations of Short Term Memory

Miller [40] found, after reviewing a number of studies, that a limit exists on the capacity of short term memory of about seven chunks. Moreover, the capacity seems to be independent of the information that is contained in a chunk, according to the research surveyed by Miller at the time, and therefore the author suggests that the capacity can be increased by forming bigger chunks of information. Miller further notes that they are not certain how a chunk is defined exactly.

However, Sweller [63] notes that Miller's notion of a chunk could be used instead of their notion of a schema. As it seems that a schema can be considered a chunk, the two terms can be considered commutable.

2.7.5 Learning

Apart from Sweller's previously presented conception of cognition, which serves as a basis for the cognitive load effects they introduce, they [63] also presented a more fundamental view on learning in an older publication. A central part of that understanding of learning are schemas, which according to the author are the basis of most learned abilities of humans. They explain the schemas' central role with a large repertoire of them being one of the differences between experts and beginners in a field. Therefore, obtaining abilities and learning relies on building schemas as a elementary component of knowledge, according to the author. Apart from their general role in learning, Sweller also notes that schemas boost the capacity of the working memory as they allow for combining multiple components into a single component. Therefore, schemas are a tool to increase working memory capacity and also of importance for long-term storage in the author's view.

In the case of the application central to this thesis, ideally an annotation would be processed as a single chunk or schema and would therefore only occupy one space in short-term memory. This compression of the different information associated with an annotation into one element is unlikely to happen for new users but it could potentially take effect for somebody that has been using the platform for some time.

Another integral part of Sweller's [63] explanation of learning is the automation of abilities as one becomes more skilled within a certain domain leading to less mental effort being required to execute certain automated actions. Automation enables us, according to the author, to go around our working memory when executing actions. In combining the role of automation with their previously mentioned view on learning and schemas, Sweller claims that one of the main purposes of learning is to "store automated schemas in long-term memory" [63, p. 298].

Analysis of Previous Annotation Behavior

To identify the needs of users of annotation systems, an exploratory data analysis was conducted on usage data of the collaborative annotation system of *Backstage 2*. This chapter first gives an overview of *Backstage 2*, then explains the methods used to analyze the data, followed by the results of the analysis and their interpretation.

3.1 Backstage 2

Mader [31] created the teaching and learning system *Backstage 2* which builds on Pohl's [49] previous system *Backstage*. In the following, an overview of *Backstage 2* based on Mader's dissertation [31], which the remainder of this section refers to unless stated otherwise, is given.

The issue which the author aims to address with *Backstage 2* is low involvement of students in their courses and courses lacking a hands-on quality for the students. Mader uses different components, a collaborative annotation system and an audience response system, to develop teaching and learning formats to combat this problem. The audience response system will not be discussed further as it is not within the scope of this thesis; refer to [31] for details. In large lecture sessions the collaborative annotation system is used to establish a way for students to communicate, thereby creating a backchannel. A backchannel is "a secret, unofficial, or irregular means of communication" [38].

A core concept within *Backstage 2* are courses. They can model periodical physical lecture sessions but also extend beyond just that by offering "asynchronous learning activities" [31, p. 9] to the learners, which they can work on at a pace convenient to them without any physical component, or hybrid learning experiences using both online and physical components. The author stresses that *Backstage 2* is capable of providing many more different learning experiences as it is use-agnostic. The remainder of this chapter will not use *Backstage 2*'s terminology however, but refer to a series of lectures spread over an entire semester as a course and to one lecture within a course as a lecture session. Furthermore, a seminar course, consisting in a series of seminar sessions usually spread over an entire semester and including the students creating a written artifact, e.g. an essay, will be referred to as a seminar.

In smaller courses or seminars, the collaborative annotation system is used to add a collaborative component to peer review. Topping [66] defines peer assessment, which is

synonymous to the concept of peer review as it is used within this thesis, as “an arrangement in which individuals consider the amount, level, value, worth, quality, or success of the products or outcomes of learning of peers of similar status” [66, p. 250]. While a peer review normally is executed in sequential stages, *Backstage 2* allows all stages of the peer review process to be executed concurrently.

Backstage 2’s collaborative annotation system enables students to annotate any material of a course or seminar they are taking part in. Annotations created are instantly made available to other users. Apart from creating their own annotations, users can also interact with others’ annotations by “commenting or voting on them” [31, p. 20]. When voting, users can up- or downvote an annotation to indicate whether they approve of its content or not.

In conclusion, the purposes of *Backstage 2*’s collaborative annotation system are giving feedback on learners’ works in collaborative peer reviews and providing a backchannel in large courses, which allows learners to communicate during the lecture sessions and afterwards.

3.2 Methods

The following section describes the structure of the data and explains how the usage data from *Backstage 2* was filtered, aggregated, and analyzed.

3.2.1 Data Gathered

Data was gathered in multiple seminars and courses at *LMU Munich* from 2017 to 2020. The data includes the artifacts generated in these courses, including but not limited to documents, annotations, and comments left on annotations, and information regarding the users’ usage of the system, e.g., an approximation of how many times a specific annotation was read.

Specifically, the data contains logs and artifacts from nine seminars. In these seminars students were tasked with producing different kinds of documents, namely scientific papers and job applications. Collaborative peer review was utilized to provide feedback on the drafts of the students’ documents. They provided their feedback in the form of annotations they created on their peers’ documents in *Backstage 2*.

Apart from these seminars, the data also contains logs and artifacts from two courses held in 2018 and 2019. In these courses, *Backstage 2*’s annotation system was used to allow students to create annotations during the lecture sessions and to ask questions and discuss the material by creating annotations outside of the lecture sessions.

3.2.2 Filtering

Prior to aggregating the data, a number of filters were applied. For both seminars and courses annotations created by a lecturer were filtered out and are therefore not included in the analysis. Comments created by lecturers were kept, however they were distinguished from those created by students. Furthermore, documents of types other than PDF, e.g., Markdown documents, and the annotations created on those documents were excluded as they do not represent course or seminar material. One specific document from seminar #7, which was created as a playground for students to explore the annotation system’s capabilities, was also excluded from the further analysis. Apart from that, portions of the data were omitted from some plots in order to improve their readability. These data points are still included in the reported quartiles though. Furthermore, 11 annotations could not be assigned to a lecture session within a course due to technical reasons and therefore

could not be classified as having been created during or outside of a lecture session. These annotations were excluded from all measures relying on such a classification.

3.2.3 Measures

For both peer review seminars and courses the annotations were grouped in various ways to be able to analyze them within different contexts. On the most general level, the annotations were grouped by the individual seminar or course they were created in. Moreover, annotations were grouped by the document they refer to, like the slide deck of a lecture session or a seminar paper of a student for peer review. Lastly, annotations were also grouped by the individual page of the document they belong to. The annotations that were created in the two courses that were evaluated were also grouped by the specific lecture session they belong to. In most cases, there was only one document used per lecture session but in one lecture session two documents were used.

Each annotation left on the documents associated with a lecture session was classified as having been created during the lecture session or outside of it. This classification was done based on the times that the lecture sessions were held at, assuming a duration of three hours per lecture session.

Additionally, each question annotation was classified as answered or unanswered based on the comments left on the annotation. In this context, an annotation is considered a question if it was marked by a user as such by selecting the purpose “question”, and a comment was considered an answer if it was left on an annotation with the purpose “question” and it itself had the purpose “answer”. Users had to explicitly choose a purpose for each annotation and comment they created using *Backstage 2*.

Apart from the mentioned binary classifications of annotations, a number of numerical measures were computed for each annotation, namely

- the number of characters of the annotation’s textual content
- the number of votes (which includes upvotes and downvotes) on the annotation
- the number of comments on the annotation
- the time elapsed since the start of the lecture session the annotation refers to (only for annotations created in a course).

Other measures were calculated for the contexts the annotations appeared in, partially aggregating the individual annotations’ measures, namely

- the number of annotations per page
- the total number of characters of all annotations on a page
- the portion of annotations that were created during a lecture session per page
- the number of annotations that were created during lecture sessions across all lecture sessions
- the number of usages of each annotation purpose across all seminars and across all courses
- the number of question annotations that an answer was provided for across all courses.

The aggregation process was first conducted per seminar or course and then the data from all seminars and from all courses was merged into one dataset for all seminars and one dataset for all courses.

seminar	#annotations	#documents	#pages
seminar 1	451	13	145
seminar 2	219	14	52
seminar 3	787	21	232
seminar 4	468	18	194
seminar 5	243	14	38
seminar 6	78	5	51
seminar 7	424	5	67
seminar 8	602	20	209
seminar 9	412	12	130
all seminars	3684	122	1118

Table 3.1: Overview over the artifacts found in the seminars' data

course	#annotations	#documents	#pages
course 1	482	14	444
course 2	769	29	464
all courses	1251	43	908

Table 3.2: Overview over the artifacts found in the courses' data

3.2.4 Analysis

For some measures a Mann-Whitney U test was conducted to determine whether there was a significant difference between groups of data. McKnight and Najab [37] list a number of conditions that have to be met to be able to conduct a t-test, which they describe as being used to determine whether two groups of data are dissimilar. One of these conditions is that the variable being examined has to be normally distributed [37]. As neither of the measures that the significance of the difference between groups was to be determined for is normally distributed, the Mann-Whitney U test was used, which McKnight and Najab describe as a “nonparametric version of the parametric t-test” [37] and suggest to use if the prerequisites for a t-test are not met [37]. More specifically, McKnight and Najab summarize Mann and Whitney [34] and Wilcoxon [73] by explaining that the Mann-Whitney U test “tests for differences between two groups on a single, ordinal variable with no specific distribution” [37]. As the measures, which the difference between groups was to be determined for, qualify as ordinal variables with no definitive distribution, a Mann-Whitney U test was determined to be fitting in these cases.

3.3 Results

In this section the results of the conducted exploratory data analysis are presented.

3.3.1 Overview

Table 3.1 shows an overview of the evaluated seminars, including the number of annotations, documents, and pages per seminar and of all seminars combined. Note that the listed counts represent the state after the application of certain filters explained in Section 3.2.2.

Table 3.2 shows an overview of the evaluated courses, after application of the mentioned filters, with the same measures shown as for the seminars.

purpose	#annotations
language	1141
misc	923
form	583
remark	583
positive	211
structure	151

Table 3.3: Purposes of the annotations created in the nine seminars evaluated

3.3.2 Annotation Creation Time

Annotations from the two courses analyzed were classified as having been created during or outside of a lecture session. The results show that more than double the number of annotations were created during lecture sessions (865) than outside of them (375).

3.3.3 Annotation Purposes

When creating an annotation using *Backstage 2*, users had to specifically choose a purpose for it, e.g., “question” or “positive”. In the two evaluated courses, there were 800 annotations with the purpose “remark”, 270 with “question”, and 181 with “answer”. Therefore, students created nearly three times the number of annotations with the purpose “remark” than with “question”.

Table 3.3 shows the purposes of the annotations in the nine seminars evaluated. Most annotations students created in the seminars were of purpose “language”, followed by “misc”, equally many with the purpose “form” and “remark”, followed by “positive” and “structure”. There were annotations with the purposes “question” and “answer” in the seminars but these were a result of misconfigurations of the system and are therefore omitted from Table 3.3.

3.3.4 Questions Answered in Courses

In Figure 3.1 the number of question annotations created in the courses with and without an answer is shown. This includes answers given by the lecturers. In this case, there are 161 questions with answer(s) and 109 without. If answers by lecturers are filtered, 130 questions with answers and 140 without remain.

From the listed counts it is evident that activity by a lecturer led to 31 of the questions being answered that would otherwise have been left unanswered. That corresponds to $\approx 19\%$ of all questions answered (when counting answers by the lecturer). The counts also show that there is a considerable number of questions that were left unanswered. When counting answers by the lecturer, unanswered questions make up $\approx 40\%$ of all questions and when filtering answers by the lecturer $\approx 52\%$.

3.3.5 Questions Answered Over Time Since Start of Lecture Session

Figure 3.2 shows the answer status of questions asked in the courses evaluated over the time elapsed since the start of the lecture session to which the question refers to. Answers given by a lecturer were not counted in this instance. Six questions could not be assigned to a lecture session for technical reasons and were therefore excluded.

The relation between the time since the start of a lecture session and whether a question received an answer was examined using a Spearman’s correlation test which yielded a




Figure 3.1: Number of questions asked in the courses with and without answers




Figure 3.2: Answer status over time since the start of the lecture session across all courses

weak negative correlation ($r_s \approx -0.221$) [65]. That is, with every minute that passes after the beginning of a lecture session, it becomes slightly less likely that a question that is posted, referring to that lecture session, will receive an answer.

3.3.6 Aggregation by Page

One way the analyzed data was aggregated was by grouping the annotations by the page of a PDF document they were created on. Based on that, several measures were then calculated per page.

3.3.6.1 Number of Annotations

Figure 3.3 shows the number of annotations per page across all seminars with the first quartile of the displayed data being at one annotation per page, the median at two annotations per page, and the third quartile at five annotations per page.

Figure 3.4 shows the number of annotations per page across all courses. Note that two pages with more than 23 annotations per page are hidden for better readability (recall Section 3.2.2). Here, both the first quartile and the median are at zero annotations per page and the third quartile is at two annotations per page.

A one-sided Mann-Whitney U test with $\alpha = 0.05$ was conducted to determine whether students created significantly less annotations per page in courses than in seminars. The results confirmed a statistically significant difference ($p \approx 4.094 \cdot 10^{-57}$, $p \ll \alpha$).

3.3.6.2 Total Text Length

Figure 3.5 shows the total text length of all annotations on a page across all seminars. The first quartile is at 4.25, the median at 124, and the third quartile at 312.5 characters per page.




Figure 3.3: Number of annotations per page across all seminars




Figure 3.4: Number of annotations per page across all courses




Figure 3.5: Total text length of all annotations per page across all seminars

Note that three pages with more than 2399 characters per page are not displayed for better readability (recall Section 3.2.2).

Figure 3.6 shows an equivalent plot for all courses. Here, both the first quartile and the median are at zero characters and the third quartile is at 58.25 characters per page.

A one-sided Mann-Whitney U test with $\alpha = 0.05$ was conducted to determine whether the total text length of annotations on a page is significantly smaller in courses compared to seminars. The results confirmed a statistically significant difference ($p \approx 5.316 \cdot 10^{-71}$, $p \ll \alpha$).

3.3.6.3 Annotation Creation Time

Figure 3.7 shows for each page of a document from the courses the percentage of annotations created on that page that were created during a lecture session. The first quartile here is at 0%, the median at 90.91%, and the third quartile is at 100% of annotations per page that were created during a lecture session.

3.3.7 Individual Annotations

Apart from grouping annotations by the page they were created on, the annotations were also examined independent from the context they were posted in. The following section shows a number of measures that were computed per annotation.

3.3.7.1 Comments

When including comments created by lecturers, users created a total of 427 comments on 3684 annotations in seminars and 450 on 1251 annotations in courses. If comments created by lecturers are filtered, 294 comments were created in seminars and 349 in courses. For




Figure 3.6: Total text length of all annotations per page across all courses




Figure 3.7: Percentages of annotations created during a lecture session per page across all courses




Figure 3.8: Text length per annotation across all seminars

both the seminars and the courses, the first quartile, the median, and the third quartile are all at zero comments per annotation, both when counting comments left by lecturers and when excluding them.

3.3.7.2 Votes

Apart from leaving comments on annotations, *Backstage 2* also allows users to upvote or downvote annotations to express their opinion of them. In the evaluated seminars users cast a total of 798 votes on 3684 annotations and a total of 1535 votes on 1251 annotations in the courses (these include both up- and downvotes and votes cast by lecturers). For both the seminars and the courses, the first quartile and the median are at zero votes per annotation, but for the courses the third quartile is at one vote per annotation while it is at zero votes per annotation for the seminars evaluated.

3.3.7.3 Text Length

Figure 3.8 shows the text length of annotations in the evaluated seminars. The first quartile is at 15, the median at 44.5, and the third quartile at 96.25 characters per annotation. Note that seven annotations with 600 characters or more are hidden in Figure 3.8 for better readability.

Figure 3.9 shows an equivalent plot of the text length of annotations in the courses evaluated. Here, the first quartile is at seven, the median at 22, and the third quartile at 50 characters per annotation. Two annotations with 500 characters or more are hidden in Figure 3.9.

To determine whether the texts of annotations created in courses are significantly shorter than those of annotations created in seminars, a one-sided Mann-Whitney U test with




Figure 3.9: Text length per annotation across all courses

$\alpha = 0.05$ was conducted. A statistically significant difference was found ($p \approx 3.311 \cdot 10^{-54}$, $p \ll \alpha$).

Within annotations from the courses, the first quartile of the text length of those created during a lecture session is at six, the median at 18, and the third quartile at 41 characters per annotation. For annotations created in a course but outside of a lecture session the first quartile is at 17 characters per annotation, the median at 37, and the third quartile at 84.

A one-sided Mann-Whitney U test with $\alpha = 0.05$ was conducted to determine whether the text of annotations created during lecture sessions is significantly shorter than that of annotations created outside of lecture sessions. The results confirmed a statistically significant difference ($p \approx 5.973 \cdot 10^{-20}$, $p \ll \alpha$).

3.4 Interpretation

In the following section the results reported in the previous section are analyzed and interpreted.

3.4.1 Annotation Purposes

In the evaluated courses, the purpose “remark” was the one used most, while the purpose “misc” was the second most used one in the seminars (recall Section 3.3.3). Both of these purposes are relatively general and do not allow for much inference of the specific content of the annotation. Therefore, there seems to be potential for the introduction of more purposes with higher specificity. Besides, although there were sizeable differences between the existing purposes regarding the number of uses, all of them were used by a significant number of users. This suggests that there are not too many purposes already, supporting

the observation that there is potential for more purposes. Therefore, the architecture of the system that is being created in this thesis should make it as easy as possible to introduce new annotation purposes.

3.4.2 Questions Answered

The questions that would not have received an answer without a lecturer providing one make up a significant portion of the total questions (recall Section 3.3.4) but by far not the majority. Therefore, it seems that students are helping each other by answering their peers' questions, and thus, the system seems to be largely successful at encouraging student interaction. From Section 3.3.4 it is also evident that there is a relatively high portion of unanswered questions. This could be explained by users not noticing that there are unanswered questions when using the system or users noticing but not being able to answer the questions because of the system being overly complicated or unintuitive to use. However, the latter explanation is rather unlikely to be true because *Backstage 2*'s interface is already relatively simple and intuitive. Another explanation could be that the students were aware of the unanswered questions but were not motivated or unable, because of a lack of knowledge, to provide answers to the questions. While improvements to the system could potentially improve users' awareness of questions without answers and make it easier and more intuitive to reply to them, students' motivation and ability to answer a question depend on numerous other factors besides the system's design.

In summary, it seems that *Backstage 2* largely succeeds at motivating students to answer questions asked by their peers but there is room for improvement to decrease the reliance on the lecturers and increase the number of answered questions. Adding means for improving communication awareness might be a promising approach to achieve such improvements.

3.4.3 Questions Answered Over Time Since Start of a Lecture Session

Section 3.3.5 reports questions' answer status over the time that had passed since the start of the lecture session the question refers to. One of the possible reasons for the observed negative correlation could be that students are less likely to notice newly asked questions as time passes during a lecture session and afterwards. Therefore, decreasing attention over time could be a factor that causes the negative correlation. In consequence, improving communication awareness could be a valuable contribution to an enhanced annotation system.

3.4.4 Aggregation by Page

The following subsections contain the interpretation of data analysis results that were obtained from annotations grouped by the page of the document they were created on.

3.4.4.1 Number of Annotations

As reported in Section 3.3.6.1 the third quartile of the number of annotations per page is at five for the analyzed seminars and at two for the analyzed courses. Based on that, it seems reasonable to assume that if the system created within this thesis could display eight annotations per page while being easy and intuitive to use, the examined use cases in the context of university teaching should be covered. From Figure 3.3 it can also be inferred that in the seminars there are many pages with only few annotations on them and a small number of pages with large numbers of annotations on them. Because of this distribution and the mentioned quartiles, the targeted eight annotations that the system should be able

to display comfortably should not be a limiting factor for the vast majority of pages in seminars similar to the evaluated ones.

The low number of annotations per page in the courses compared to the seminars could be due to the creation of annotations being one of the students' main tasks in peer review seminars, while in the courses, where the annotation system was used as a backchannel, annotation creation was voluntary and not the main focus of the lecture sessions. The statistically significant difference between seminars and courses regarding the number of annotations per page can be interpreted as a sign of courses and seminars representing different use cases of an annotation system. A new annotation system should therefore be flexible to allow it to be adapted to the requirements of various use cases.

3.4.4.2 Text Length

The difference in the total text length of all annotations on a page between seminars and courses (see Section 3.3.6.2), which was found to be statistically significant, can be seen as evidence of the different requirements that the seminars and courses impose on an annotation system. This observation provides further evidence for the conclusion that a new annotation system should be easily adaptable to the requirements a specific use case imposes on it.

3.4.4.3 Annotation Creation Time

The analysis of the annotation creation time of annotations from the courses in Section 3.3.6.3 shows that many pages are annotated almost exclusively either during or outside of lecture sessions, represented in Figure 3.7 by the leftmost and rightmost bars. The requirements of users during and outside of lecture sessions could differ, e.g., because of differences in the amount of time and mental resources the user has available to interact with the annotation system. During a lecture session these could be lower as the user has to focus on following the lecture. The observation that most pages are either almost exclusively edited during or outside of lecture sessions could suggest that a new annotation system could be most successful in supporting the users' work not by trying to fulfill all different requirements at the same time, but by adapting based on the current state of the environment that it is used in, e.g., whether there is an ongoing lecture session or not, to provide the best user experience possible at all times. Therefore, a new annotation system should not only be easily adaptable in general but also offer the possibility to easily switch between different customizations.

3.4.5 Individual Annotations

Below, the conclusions drawn from the analysis of annotations without the context they were created in are explained.

3.4.5.1 Comments

As Section 3.3.7.1 reports, users created a number of comments both in the evaluated seminars and the evaluated courses. When relating the number of comments to the number of annotations it seems, however, that a higher number of comments might be desirable because it might boost the users' feeling of being part of a community. Furthermore, through users interacting with the content of other users' annotations, a higher number of comments might lead to more answers to questions asked in annotations being posted. One improvement, which could be made to increase the rate of interaction between users via comments, could again be to increase communication awareness, i.e., in this case, making

users more aware of the comments other users write. This should then ideally lead to users feeling motivated to participate in ongoing discussions by posting comments themselves.

3.4.5.2 Votes

Students did use *Backstage 2*'s voting feature for annotations, as results in Section 3.3.7.2 show, but the number of votes cast was relatively low when taking the total number of annotations into consideration. In general, it would be desirable to motivate users to cast more votes on annotations, thereby bringing more social interactions to the system and providing more, albeit simple, feedback to other users. However, the goal is not to animate users to vote on all annotations but only on those that are relevant to them.

3.4.5.3 Text Length

The statistically significant difference between the text length of annotations in seminars and courses, as reported in Section 3.3.7.3, again suggests that the seminars and the courses represent two different use cases of the system. This observation adds to the evidence suggesting that a new annotation system should be adaptable to the different requirements different use cases impose on it. Similarly, the difference between the text length of annotations created during lecture sessions and those created outside of them, which was found to be statistically significant as well, also hints at the existence of two different use cases within a course. The argument that a new annotation system should offer the possibility to easily switch between different customizations, beyond being easily adaptable in general, is strengthened by this observation.

Outside of lecture sessions it would be desirable to motivate users to write longer annotations in some situations where short annotations do not suffice because of the complexity of the topic as longer annotations could be a sign of deeper involvement with the material in these situations. Besides, longer annotations could lead to more collaboration and interaction as they could attract more comments and votes because of them most likely containing more detailed feedback or new content. During lecture sessions longer annotations are not desirable, however, as the users' main time and mental resources should be spent on following the lecture and not on reading or writing long annotations.

Regarding the system being created within this thesis, one goal could be to be able to comfortably display annotations of up to 280 characters, the maximum length of a *Tweet* on *Twitter*¹ at the time of writing [67]. If that goal was to be reached when designing the new system, the vast majority of use cases both in the courses and seminars should be covered, according to the results in Section 3.3.7.3.

3.5 Threats to Validity

When classifying annotations from the evaluated courses as created within a lecture session or outside of one, each lecture session was assumed to have been three hours long, as scheduled. This is a threat to the validity of the analysis because lecture sessions might have ended earlier or taken longer than planned which might have led to annotations being classified incorrectly.

When analyzing the number of answered questions some of them might have been left unanswered not because of some property of the annotation system but because they were phrased badly, and therefore students were not able to understand them, or because they were too difficult to answer given the knowledge conveyed within the course. Missing

¹<https://twitter.com/>

categorizations of annotations as questions might also have influenced the number of annotations that were considered questions. The same applies for comments that might have not been marked as answers and therefore might have led to questions that were answered being considered unanswered. Hence, there most likely are more answers to questions asked in annotations than this analysis captures.

Furthermore, the number of comments per annotation does not capture whether these comments were of high quality and contributed to an insightful discussion. Therefore, the number of comments per annotation can only be considered a quantitative measure of interaction and not a qualitative one.

The following chapter explains the ideas behind the implementation of the annotation system built as part of this thesis and describes potentially beneficial extensions that could be added to the system.

4.1 Framework

The following section will explain the three different modes that the system can be used in: overview mode, list mode, and detail mode. These modes represent a framework around the various exchangeable components of the system, which can be used to customize it. Many of the features described have already been implemented in *Backstage 2* [31].

In explaining these modes, the following section will repeatedly refer to cognitive load effects, the impact of all of which could be, according to Sweller [62], limited by the *element interactivity effect*. This effect could occur due to low element interactivity connected to intrinsic cognitive load of the material shown and could limit the impact of other cognitive load effects, the author explains. Therefore, the benefit of some of the measures outlined below that aim to avoid negative cognitive load effects and facilitate positive ones could be limited. For annotations with simple content and few comments the occurrence of the *element interactivity effect* seems likely but it could be absent for annotations with complicated and long content and a large number of comments.

The system can be used in the document-centric overview mode, from which a user can get to the list mode by extending the sidebar. The list mode, among other things, allows for filtering and sorting annotations in the sidebar. From there the user can get to detail mode, which presents detailed information on one specific annotation, by clicking a link in an annotation's preview in the sidebar.

4.1.1 Overview Mode

When a user opens a document using the annotation system, the overview mode is the first mode they are shown. It allows the user to gain an overview of the document and its annotations without directing too much of their attention away from the document itself by showing only the document, the annotations' contexts, and small icons representing the annotations' purposes, as shown in the upper half of Figure 4.1. Note that all mockups

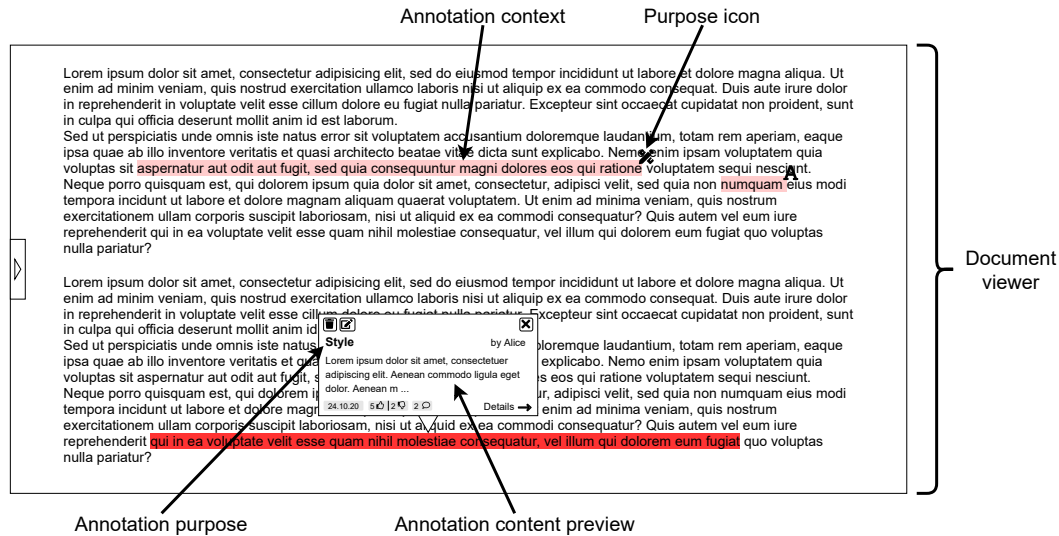


Figure 4.1: Mockup of the overview mode with open annotation preview popup (see Appendix A.1 for icons' sources)

referenced in this chapter show the implemented textual annotations on PDF documents. However, the system supports different documents and content types.

If the user clicks on an annotation's context in overview mode, the context is highlighted and a popup opens that shows a preview of the annotation's content, as shown in the lower half of Figure 4.1. In his master thesis Mader [30] defines an annotation's context as "the part of a medium an annotation is referring to, i.e., the part of a medium that is annotated" [30, p. 13]. These popups are the only source of information regarding annotations' contents in overview mode and therefore no *split attention effect* should occur. In the case of the textual annotations implemented within this thesis, the context preview consists in the first 100 characters of the annotation's content. Such a limit seems to be acceptable for textual annotations as, based on the analyzed courses and seminars, only the content of relatively few annotations will be truncated because of this limit. Apart from the preview of the annotation's content, the popup also contains a link that the user can click to open the detail view for that annotation. Furthermore, the popup contains metadata on the annotation including the username of the author of the annotation, its creation date, and its purpose. Some of this metadata is presented in summarized form, e.g., the number of comments created regarding an annotation is shown but not the comments themselves. Not showing metadata with an overly high level of detail that is not necessary to understand the meaning of the annotation should help to avoid the *redundancy effect*. Some of the metadata displayed also acts as triggers for actions related to the annotation. Clicking the number of upvotes, e.g., lets the user upvote the annotation themselves. Moreover, the popup contains buttons to delete the currently viewed annotation or edit its content if the current user is the author of the annotation.

To get from overview mode to list mode the user can expand the sidebar on the left of the document viewer. Both modes allow for the creation of annotations by, assuming the implemented textual annotations, first specifying the new annotation's context by selecting a text section or arbitrary point in the document viewer, then choosing a purpose (see Figure 4.2), and then entering the annotation's textual content (see Figure 4.3).

The overview mode should, assuming the implemented textual annotations on the PDF of a typical seminar paper, be able to display eight annotations at a time, in the form of

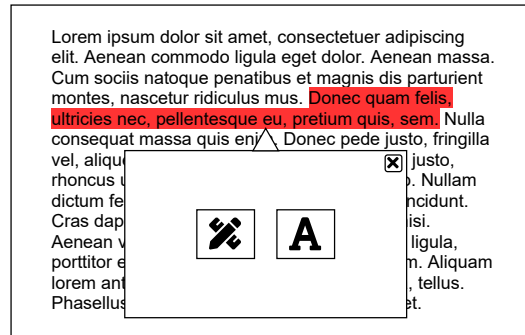


Figure 4.2: Mockup of purpose selection shown after selecting a new annotation's context (see Appendix A.1 for icons' sources)

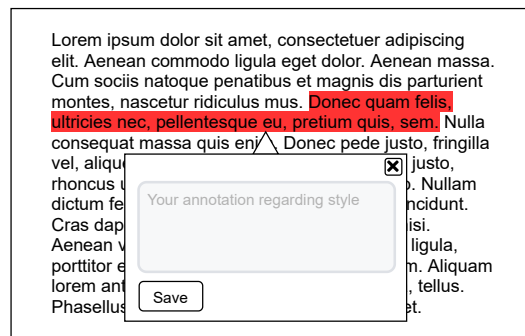


Figure 4.3: Mockup of content input using a simple textual content editor shown after selecting a new annotation's purpose (see Appendix A.1 for icons' sources)

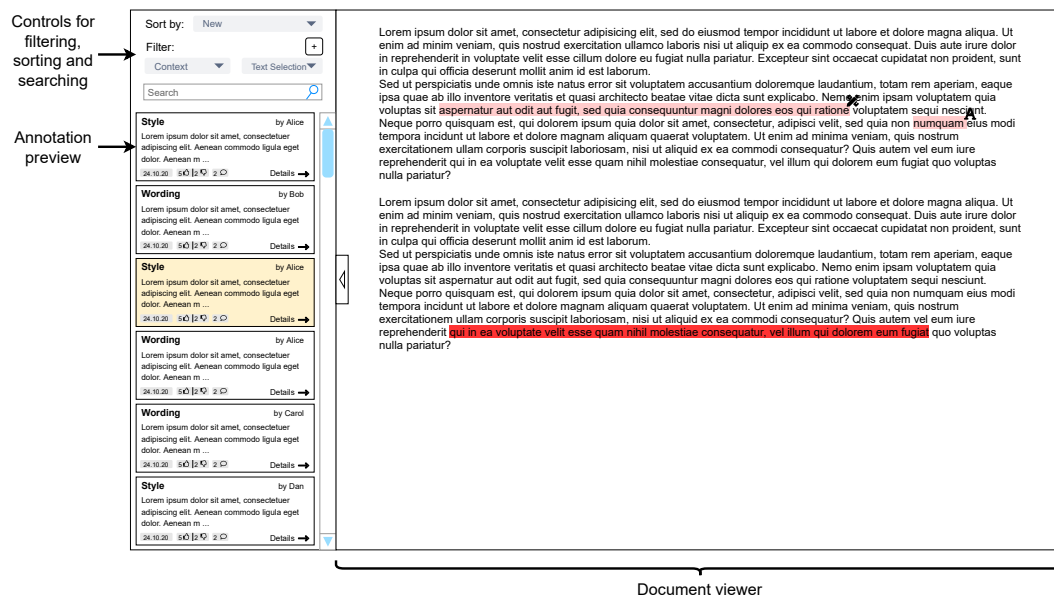


Figure 4.4: Mockup of the list mode (see Appendix A.1 for icons' sources)

their contexts in the document viewer, and therefore be able to cover most use cases based on the analyzed data.

The overview mode could be particularly well suited for reviewing an essay as part of the peer review process in a seminar or going through lecture slides outside of a lecture session, as it allows the user to focus on the document, create annotations with minimal effort and therefore minimal interruption of the reading flow, and view annotations without being shown too much information that would take long to process, which could also interrupt the reading flow. From the initial minimal display of an annotation's information, the user can then decide to view annotations they are interested in in detail by opening the detail mode from an annotation's popup.

4.1.2 List Mode

One way to activate the list mode, a mockup of which is shown in Figure 4.4, is to expand the sidebar while in overview mode. In list mode the document with the annotations' contexts is shown on the right of the screen while a list of annotation previews is shown in the sidebar on the left. The annotation previews shown in the sidebar are very similar to the previews shown in the popups in overview mode with the only difference being that those in the sidebar do not include a row of buttons above the purpose of the annotation. The annotation list in the sidebar is, assuming the implemented textual annotations, able to display the previews of eight annotations at a time and therefore should be able to cover most use cases based on the analyzed data. If a user clicks an annotation preview in the sidebar, the matching annotation's context is highlighted in the document viewer and vice versa, thereby allowing the user to connect an annotation's content to its context.

Above the list of annotation previews in the sidebar controls, which allow the user to filter, sort, or search the annotations (see Figure 4.4), are shown. For sorting annotations, users can choose from different sorting criteria, e.g., creation date, and an either ascending or descending sorting direction. For filtering annotations, users can add any number of filters, each consisting of a filter criterion, e.g., annotation purpose, and a value for that criterion. If multiple filters are used, only annotations that fulfill all filter criteria are shown.

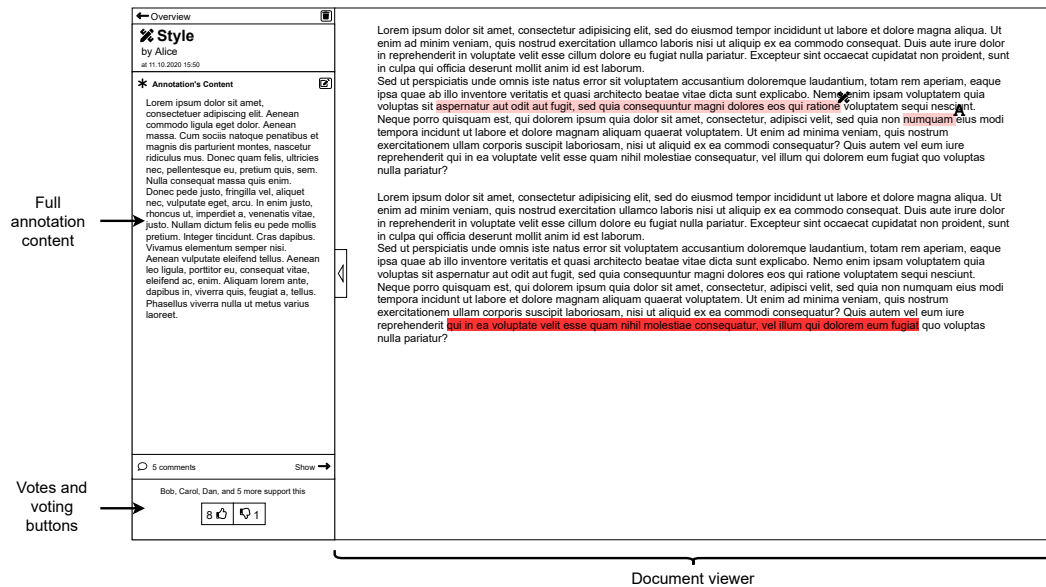


Figure 4.5: Mockup of the detail mode (see Appendix A.1 for icons' sources)

Using the search field in the sidebar, users can search the usernames of the authors of annotations and the textual representations of annotations' contents, which in the case of the implemented textual annotations are identical to the contents but might be different for other kinds of annotations.

The list mode allows the user to see all annotations belonging to the currently visible part of the document, e.g., a page, at a glance using the sidebar. On the one hand this might make it harder for the user to focus on the document itself, due to more information being permanently shown in addition to the document. Additionally, it might make it more difficult for them to process annotations in complex documents, as the connection between an annotation's content and its context, expressed using a highlighted content preview in the sidebar and a highlighted context in the document viewer, is not as clear as in overview mode. On the other hand it could make list mode particularly well suited for situations in which the focus is more on the communication taking place through annotations than on the document itself. A lecture session might be such a situation, as lecture slides have, compared to, e.g., essays, less content and communication is highly important during a lecture session, as students might want to react to remarks or questions posted by their fellow students as quickly as possible. A quick response to another student's annotation might be desirable because the memory of the lecture session's topic is still fresh during the lecture session and the lecturer might, if they read students' annotations during the lecture, address questions or remarks in a timely manner during the lecture session.

In list mode annotations can be created the same way as in overview mode.

4.1.3 Detail Mode

Detail mode is activated when a user clicks a link in an annotation preview in overview or list mode. In this mode the sidebar is extended and shows detailed information on one specific annotation while the document is still visible on the right of the sidebar, as shown in Figure 4.5. The context of the annotation that is shown in detail is highlighted in the document view.

In detail mode, basic metadata is shown in the sidebar: the annotation's purpose, the

username of its author, and the date it was created at. Furthermore, the sidebar shows the number of upvotes and downvotes the annotation has received and allows the user to cast their own vote (see Figure 4.5). It also shows the usernames of the users that upvoted the annotation: The usernames of the first three users are always shown while the full list of usernames is only shown if the user hovers their cursor over the first three usernames. Showing usernames instead of their votes just being included in the total count of upvotes should lead to increased social presence as these users should be perceived as more of a real person by the user of the system this way. Gunawardena and Zittle [17], in summarizing Short et al. [58], define social presence as “the degree to which a person is perceived as a ‘real person’ in mediated communication” [17, p. 9] and Rovai [55, 56] suggests that social presence positively influences the sense of community that students experience when participating in online learning activities. Apart from voting-related functionalities, the sidebar also contains buttons that allows the user to delete the annotation or edit its content, if they are its author.

A possible extension of the current system could be to prioritize certain users, whose opinion might be of particular importance, when creating the list of users that upvoted an annotation, by putting their usernames at the front of the list. Such a prioritization could, e.g., be based on user’s roles like lecturer or reviewer of a document. It would, however, lead to the order that the usernames are shown in not representing the order that the users voted in anymore.

Another possible extension of the current system could be to calculate a rating, using the number of upvotes and downvotes for each annotation, and replace the number of upvotes and downvotes in detail mode as well as in the two other modes with it. Such a rating already exists in *Backstage 2* [31]. Mader [31] utilizes an algorithm suggested by Miller [39] which uses the “[l]ower bound of [the] Wilson score confidence interval for a Bernoulli parameter” [39] to calculate the rating score. Millers algorithm was at one point used by *Reddit*¹ [54].

In its main area the sidebar in detail mode either shows the annotation’s content (see Figure 4.5) or a list of comments (see Figure 4.6), offering the user the possibility to switch between the two. If the detail mode is opened using a link from the overview or list mode, the annotation’s content is shown by default. This should help to avoid the occurrence of the *redundancy effect* as an annotation’s content is more essential to understanding its meaning in most cases than the comments created regarding it.

If the content section is extended in the sidebar, the full content of the displayed annotation is shown (see Figure 4.5). Additionally, this view allows the user to edit the annotation’s content using a button if the user is the author of the annotation. Assuming the implemented textual annotations, this view is able to show a textual content of 280 characters and more and should therefore be able to cover most use cases based on the analyzed data.

If the comment section is extended in the sidebar, the comments referring to the currently viewed annotation, together with an editor for the user to create their own comments, are shown in the sidebar. If the user is the author of a comment, buttons to delete or edit the comment are shown in the comment’s header. The sidebar with the extended comment section is shown in Figure 4.6. Due to the limited available horizontal space, the comments cannot be nested, i.e., replies are not indented relative to the comment they are referring to. One could argue that more horizontal space could be acquired by removing the document view to the right of the sidebar, which would mean losing the annotations’ contexts in list mode, or making it considerably smaller, which, depending on the displayed document, could make understanding an annotation’s context harder. Hoff et al. [20] see the context as an essential part of an annotation. They further note that the whole

¹<https://www.reddit.com/>

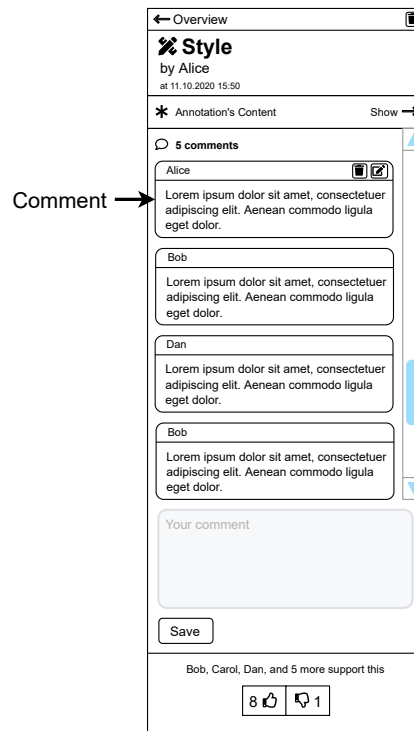


Figure 4.6: Mockup of the sidebar in detail mode with selected comment view (see Appendix A.1 for icons' sources)

annotation would not have any value without its context in most cases and list “loss of context [...] an annotation was created in” [20, p. 234] as one of the deficiencies they found in annotation systems they evaluated. Therefore, such changes do not seem beneficial.

As the user is shown an annotation's content without the related comments before being shown the comments in the sidebar in detail mode, the *isolated elements effect* should occur. The locations where the user is potentially shown an annotation's content before accessing its comments, out of which they pass at least one, are the content previews in overview mode and list mode and the full content view in detail mode.

When attempting to create an annotation in detail mode, the system will switch to list mode so that the user can immediately see their newly created annotation, including a preview of its content, in the annotation list in the sidebar, instead of being shown the details of the annotation they had previously selected in detail mode.

Although the document viewer is visible in all three modes, clicking an annotation's context in a mode apart from overview mode does not open the annotation preview popup in order to not duplicate information as the annotation's content and metadata are already shown in the sidebar in these modes. This is another measure that should help to avoid the occurrence of the *redundancy effect*.

In summary, the role of the detail mode is to allow the user to interact with one specific annotation while the other two modes are focused on allowing them to interact with all annotations belonging to one specific part of the document, e.g., a page. Part of the interaction with one annotation is allowing the user to communicate with other users via comments.

4.1.4 Conclusion

The three different modes the annotation system features are a first step towards fulfilling the different requirements of seminars and courses and usage within and outside of lecture sessions that the previous data analysis revealed. However, a flexible annotation system is still desirable as it might, through adaptations, allow for an experience even more tailored towards one of the use cases discovered during the previous data analysis or for adjusting the system to a totally different use case.

4.2 Configurability

This section introduces a number of key components of the annotation system and explains how these components can be exchanged and therefore how the annotation system can be adapted to different use cases.

4.2.1 Annotations

An annotation's representation in the implemented system is not monolithic but consists of a number of exchangeable parts: The content data and context data are the data representations of an annotation's content and context, while the content renderer and context renderer are their visual representations that depend on the data representations. The content renderer is used to display the content previews in overview mode (see Figure 4.1) and list mode (see Figure 4.4) and the full content in detail mode (see Figure 4.5). The context renderer is used in all modes to show the annotations' contexts on the document (see, e.g., Figure 4.1). For the implemented textual annotations on PDF and Markdown documents, the content data consists in an object holding text, the content renderer in a component displaying that text, the context data in either the corner points of a polygon or a single point and the length of the side of a square, and the context renderers in components rendering either a SVG polygon or a SVG rectangle. Apart from these parts, an annotation's representation also includes additional metadata in the form of the annotation's author, users that upvoted the annotation, the number of upvotes and downvotes, the time the annotation was created at, and comments created regarding the annotation. *Backstage 2* also makes a distinction between context data and content data and keeps similar metadata for each annotation as well [31].

The visual representation of an annotation's context or content can therefore be replaced without having to change its data representation by replacing the context renderer or content renderer. Therefore, the implemented SVG context renderers on PDF and Markdown documents could, e.g., be replaced by context renderers utilizing the *Canvas API*, a way "for drawing graphics via JavaScript and the HTML `<canvas>` element" [41]. Furthermore, the simple content renderers showing the implemented textual annotation contents could, e.g., be replaced by more advanced content renderers that not only show the content in textual form but also offer voice output.

As content-related functionality and context-related functionality are independent of each other, one can be changed without having to change the other, e.g., the current implementation of textual contents could be replaced by hand-drawn contents without having to change context data or context renderer. Conversely, the implemented textual content data and content renderer could be kept while replacing context data and context renderer, e.g., to allow users to select parts of an audio file as the context for their annotation. Mader [31] already mentioned hand-drawn images as a potential type of content and the ability to independently exchange components of the system relating to content and components of the system relating to context. Mader further implemented contexts compatible with au-

dio and video documents within their master thesis [30], but these components were never added to *Backstage 2* [31].

4.2.2 Documents

Apart from parts of the annotation's representation, the types of documents shown in the implemented system are also exchangeable in the form of different document viewers that can be integrated into the system. A document viewer is responsible for showing the document itself and the contexts of annotations created regarding the document, as well as creating the context data and context renderers of new annotations. It is part of all three modes.

Arbitrary documents combined with arbitrary annotation contexts can be integrated into the system as long as they abide by a set of basic rules. These include providing a mechanism for opening and closing the popups in overview mode, highlighting the context of the currently selected annotation in some way, and respecting the boundaries of the space allocated for the document viewer.

The implemented system offers an API for a document viewer to interact with other parts of the system, which includes functionalities like highlighting an annotation's preview in the sidebar (see Figure 4.4). Furthermore, it allows the document viewer to add custom controls to the top toolbar of the application (not shown in the mockups).

Within this thesis, a document viewer for PDFs and a document viewer for Markdown files were implemented. Both use libraries to render the content of the document and then superimpose a SVG layer, which is used to show annotations' contexts in the form of SVG polygons and SVG rectangles. The PDF document viewer utilizes the annotation system's API to add buttons to, e.g., change the page of the PDF document to the top toolbar. Other kinds of document viewers, e.g., an image viewer that could allow users to mark parts of the image by drawing polygons, which could then also be shown in the form of SVG polygons, can easily be integrated into the system. *Backstage 2* already includes the possibility to annotate images by drawing polygons [31]. The SVG annotation layer used for the PDF and Markdown document viewers was built to be reusable by other document viewers apart from the implemented ones, but other types of annotation presentation techniques can also be integrated into the system. An example of such a technique would be *XPath* expressions which can be used to find elements in a HTML document [43]. They therefore could be used to address passages of text the user has selected, e.g., in a rendered Markdown document, and highlight them instead of using SVG shapes, as the provided implementation does.

4.2.3 Content Editors

When the user has selected a context for a new annotation they are creating, they are presented with a popup (see Figure 4.2) where they can choose the purpose of their new annotation. Based on that choice the matching content editor is then shown (see Figure 4.3). It is also shown when the user clicks the edit button in the sidebar in detail mode (see Figure 4.5) or in the annotation preview popup in overview mode (see Figure 4.1). The content editor generates content data and a content renderer and therefore, as the user has already chosen the annotation's context, completes the annotation. The provided implementation only includes simple content editors for textual annotations but other types could be added to the system, e.g., a voice recorder to create annotations in the form of audio messages, a canvas to create hand-drawn annotations, or an upload form to create annotations consisting of images. These content editors could then be combined with any type of document and any type of context. Mader [31] already mentioned *Backstage 2*'s two-step content creation process, which consists of first selecting an annotation's purpose and then specifying

its content. He further listed graphics and audio as potential types of content that an annotation system could be extended with.

In the current implementation, a content editor is always associated with one specific purpose, as it is supposed to produce contents of that purpose. All content editors and their associations with a purpose are registered at a central registry. Adding a new purpose only requires adding the new purpose to the list of existing purposes and then registering a content editor that produces contents with this purpose at the central registry. Therefore, adding new purposes to the system is easy, which is a goal identified during data analysis.

4.2.4 Filtering

The system allows for the addition of arbitrary filters that the user can apply to constrain the annotations shown in the sidebar and the document viewer in list mode. These filters that the user can control from the sidebar in list mode (see Figure 4.4) can use any property of the annotations' representations, including their content data, context data, and all metadata. The current implementation allows the user to filter annotations based on their purpose or context type. Potentially, the user could be given the possibility to filter annotations, e.g., by the user that created them or by whether they were created within or outside of a lecture session, assuming the times of lecture sessions are known. As multiple filters can be used simultaneously, the addition of one new filter enables a multitude of new filter combinations. *Backstage 2* already allows for filtering annotations by their purpose [31].

4.2.5 Sorting

Similar to filters, arbitrary sorting mechanisms can be added to the system. A user can change the currently applied sorting criterion and the sorting direction in the sidebar in list mode (see Figure 4.4) above the filter selection. The selected sorting criterion influences the order the annotations in the sidebar are displayed in. The current implementation enables the user to sort annotations by their creation date, purpose, author's username (alphabetically), number of upvotes, and y-coordinate of their context on its page of the document. Same as filters, sorting mechanisms can also access all properties of the annotations' representations. Sorting criteria that could potentially be added to the system include the number of comments on an annotation and, assuming textual annotations, the length of the annotation's content. *Backstage 2* already allows for sorting annotations "either by rating or creation date and either ascending or descending" [31, p. 24].

4.2.6 Situational Annotation Filter

The implemented annotation system allows the document viewer to determine which annotations should be shown by setting and updating a situational annotation filter. This filter not only affects the annotations' contexts shown in the document viewer itself, but also the annotation previews in the sidebar in list mode (see Figure 4.4). The situational annotation filter is not applied on search results, however, which therefore can contain annotations from the whole document. In the executed implementation of textual annotations on PDF documents the situational annotation filter is used to only show annotations that relate to the currently shown page of the PDF. Due to the filter making no assumptions on the document's properties, it can be used for any type of document or can be disabled for document viewers where all annotations should always be visible, e.g., the implemented Markdown viewer. An example of a potential usage of the situational annotation filter are annotations on a map, where it could be used to only show annotations in the sidebar whose context is currently within the section of the map that the user is viewing in the document viewer.

4.3 Configuration Example: Annotations on Videos

One example of a scenario that was not implemented within this thesis but is viable using the implemented system are textual annotations attached to a region of a certain part of a video. Mader [30] already mentioned this concept of “spatio-temporal context” [30, p. 44] for videos. A user would pause the video, specify their annotation’s context by drawing a polygon with their mouse over a region they want to annotate, and then specify their annotation’s content by entering the text they want to annotate that region with. Such a capability to annotate videos might be particularly beneficial for learners during the current COVID-19 pandemic as universities are increasingly utilizing video lectures. Contexts for use with videos were developed by Mader within his master thesis [30] but were never added to *Backstage 2* [31].

To realize this scenario using the implemented system the following parts can be used: The document viewer would consist in a video player with an SVG layer on top of it. The user would use the player’s controls, which can be added to the top toolbar using the API the system provides to document viewers, to play and pause the video. When the user has found a part of the video they want to annotate, they would pause the video and draw a polygon on the SVG layer. Therefore, the new annotation’s context data would consist in a data representation of the polygon and the time in the video that the annotation refers to. The annotation’s context renderer can then use that context data to produce a SVG polygon. After the user has selected the annotation’s context in the form of the polygon, the annotation creation popup, which all document viewers share, would open and the user would first choose a purpose (analogous to Figure 4.2) and then enter their textual annotation using a simple content editor (analogous to Figure 4.3). The annotation would then appear on the video player in the form of a polygon. For viewing annotations, the document viewer, which in this case would contain the video player, would periodically, e.g., every five seconds, update the situational annotation filter so that only annotations within a sliding window of about 30 seconds, i.e., annotations whose time in the video approximately falls into the previous or next 15 seconds of video playback, would be shown. This would lead to only these annotations being displayed in the sidebar in list mode and in the document viewer in all modes. This usage of the situational annotation filter would be possible because the time in the video that an annotation refers to would have been saved in its context data before and the situational annotation filter can access all properties of annotations’ representations, including their context data. A sliding window of about 30 seconds seems like a reasonable default for educational videos, for which ten minutes or less seems to be an ideal length. This is supported by Harrison [19] who found in a study among online learners that “almost three-fourths of the participants indicated that they preferred the videos [to] be less than ten minutes [long]” [19, p. 185] and Ozan and Ozarslan [46] who found in a study among university students that 58% of videos that were watched to the end were less than ten minutes long. Assuming a video length of about ten minutes, a sliding window of about 30 seconds should allow the user to gain an overview over annotations in the current part of the video while not overloading the interface with annotations from distant parts of the video. To view an annotation’s content, the user would click the polygon superimposed on the video player and the annotation preview popup containing the content renderer (which would look similar to the popup in Figure 4.1) would open.

4.4 Exchanging Customizations

The annotation system built within this thesis fulfills the goal of allowing for adaptation to different usage scenarios. It also fulfills the goal of making it possible to switch between adaptations relatively easily as many of its components can be replaced without requiring

substantial effort. Content editors, filters, and sorting mechanisms are registered at central registries and the document viewer is, in encapsulated form, passed to the root component of the application as a parameter. Using the mentioned points of attachment of certain components, the system can be transformed to fit a different use case with relative ease.

One example of such a transformation would be offering simpler content editors during lecture sessions, to allow for faster creation of annotations that requires less mental resources, but then switching to more sophisticated content editors after the lecture session has ended, to allow for more detailed annotations. A property of text editors used outside of lecture sessions could be that they have larger text fields to nudge users towards writing longer annotations, which was identified as a goal for the implemented annotation system during data analysis. The transition of the system between the two states could be realized by having the server switch to different instructions for registering content editors at their central registry upon the end of the lecture based on the current time. Users would then be presented with different content editors upon their next reload of the web page. A potential extension of the system could consist in rolling out changes in real time while avoiding loss of data that might occur if a user's browser was forced to reload the web page while they were creating an annotation.

As certain parts of the system that generate data, namely the content editors and the document viewers, can be exchanged relatively easily, the data generated by these parts, namely the contents and contexts, can also be changed relatively easily. This aspect adds to the flexibility of the implemented system.

4.5 Interactions with Annotations

One conclusion from the data analysis is that more interactions of users with annotations, specifically in the form of votes and comments, would be desirable. This section describes measures that could help increase the number of interactions.

As one of these, the implemented capability to sort annotations by the number of upvotes they have received might provide a source of motivation for users to examine popular annotations and then potentially express their opinion of these annotations by voting on them or by creating comments.

Another measure that might help increase the number of interactions with annotations is the display of the names of the users who upvoted an annotation, which was added to the system as a measure to increase social presence. As previously mentioned, Rovai [55, 56] sees social presence as a positive factor of influence on the sense of community that online learners feel. An increased sense of community could then make them more motivated to participate in their learning community by commenting or voting on annotations or replying to comments.

A small extension that could be added to the system to potentially increase the number of comments created and votes cast is the possibility to sort annotations by their number of comments. This might lead to users using that sort capability to identify annotations where a large number of comments indicate an ongoing discussion and then potentially joining that discussion by either creating comments themselves or voting on the annotations. Adding such a sorting capability would require little effort, as explained in Section 4.2.5.

Analogous to adding the ability to sort annotations by their number of comments, the ability to sort annotations by their rating (the addition of a rating was proposed in Section 4.1.3), could also be added to the system.

As mentioned in Section 4.2.5, such a sorting mechanism could be added to the system relatively easily. It already exists in *Backstage 2* [31]. Users might utilize it to identify either popular annotations (with a high rating) or unpopular annotations (with a low rating)

and then express their opinion regarding them either by voting on them or by posting a comment. Miller [39], who proposed a rating algorithm that was discussed in Section 4.1.3, notes that “it may be more useful in a ‘top rated’ list to display those items with the highest number of positive ratings *per page view*” [39] than using the formula they proposed. Miller’s suggestion could be implemented by adding a system for counting the number of views an annotation receives and then offering the capability to sort by the number of upvotes per times an annotation was viewed, while still showing the previously explained rating in annotation previews and in detail mode. Such a duality of rating mechanisms might cause confusion among the users, however. Therefore, it could be beneficial to only add either a rating or the capability to sort by the number of upvotes per view to the system. Further research, e.g., in the form of a user study, would be needed to determine which of the two would benefit users more. *Backstage 2* already includes a mechanism for counting the number of times an annotation was read [31].

Another way to get more users to express their opinion in the annotation system could be to add the possibility to vote on comments. This feature already exists in *Backstage 2* [31]. Although it does not increase the number of votes cast on the annotations themselves, it might lead to more engagement of users with the annotations’ contents, which is often needed to understand the comments. The users might then form an opinion on that content, which they might want to express by voting on the annotation or creating a comment referring to it. Another reason why this extension of the system might drive users to write more comments is that they might find it motivating to receive additional feedback on their own comments beyond feedback in the form of other comments.

4.5.1 Comments

The annotation system contains ways for the user to jump to the comment creation view of the sidebar in detail mode from each mode: In both overview and list mode links in the annotation previews allow the user to do so and in detail mode clicking the link to open the comment creation view does. These numerous opportunities to jump to the comment editor from a large portion of the views of the application might help to increase the number of comments as the effort it takes for the user to create a comment is kept relatively low. Therefore, users that tend to not create comments because it requires too much effort might be more inclined to do so. The previous implementation of the annotation system in *Backstage 2* only offered the possibility to start the process of creating a comment from a sidebar, which provided similar functionalities to the current system’s sidebar in list mode [31].

4.5.2 Votes

Analogous to adding links to create a comment to all modes of the system, buttons to cast votes on annotations were also added to all modes: In both overview and list mode they are included in the annotation preview and in detail mode they are a permanent part of the sidebar and therefore always visible, no matter whether the comment or content view is selected or not. As with adding the links to the comment creation view, this measure aims at lowering the effort required to vote on an annotation as much as possible in the hope of getting users that refrained from voting on annotations because of the required effort to cast votes. Analogous to comment creation, the previous implementation of the annotation system in *Backstage 2* only offered the possibility to vote on annotations in a sidebar similar to this system’s sidebar in list mode [31].

4.6 Communication Awareness

This section introduces measures designed to improve users' communication awareness regarding annotations in general and specifically regarding question annotations, which, based on the analyzed data, was identified as a goal for the annotation system built within this thesis.

4.6.1 Marking New Annotations

One way to increase users' communication awareness regarding newly created annotations could be to visually highlight such annotations. In the current implementation of textual annotations on PDF documents this could, e.g., be done by adding a border to the context of new annotations. Such a border could also be added to the preview of recently created annotations in the sidebar in list mode. *Backstage 2* already contains a similar feature where unread annotations are highlighted by showing them in bold face [31].

It could be beneficial to add a temporal dimension to the emphasis of a newly created annotation by, e.g., letting the border become smaller over time until it disappears completely. The temporal progression of such a visual change would have to be adaptable, as different use cases of the annotation system will have different definitions of what a 'new' annotation is. In a lecture session, e.g., an annotation that was created half an hour ago could already be considered relatively old, as users most likely check for new annotations frequently as the lecture session progresses, while the same annotation could be considered relatively recent in a seminar, where students will most likely work on peer reviews from home and only check for new annotations irregularly.

A simpler way to add a temporal dimension to the emphasis of recently posted annotations would be to base the level of emphasis on the position of an annotation in the list of the most recent annotations. In the case of the previously described emphasis using a border around an annotation's context this could mean displaying the most recent annotation with the thickest border, the second most recent one with the second thickest border and so on.

It might be beneficial to allow users to choose between the two described approaches as some users might find the movement, which the changing accents from the first approach can create as time passes, distracting. To determine whether one of the described approaches is better suited, further research, potentially in the form of a user study comparing them, would be needed.

Apart from allowing users to choose between different versions of temporal progression of the emphasis, it might also be beneficial to let them customize how noticeable the emphasis should be in general. An example of a situation where a user might wish for more noticeable emphasis than usually is a lecturer that uses the annotation system to allow users to ask questions during a lecture session. The lecturer might only be able to divert a small portion of their focus towards the annotation system and therefore might benefit from more noticeable emphasis than preferred by other users. For the same reason, students that want to follow the lecture session but simultaneously be aware of new annotations might also benefit from more noticeable emphasis. Furthermore, this situation also serves as an example for a possible application of the simpler version of temporal progression of the emphasis as the lecturer will most likely only check the annotation system irregularly and therefore the first proposed version of temporal progression might be of little use to them as the emphasis might already have disappeared when they check the system.

Certain annotations could be excluded from being highlighted as recently created. Examples of such annotations are those created by the user themselves and questions that

have already been answered. This would require a representation of the answered/unanswered status of an annotation in the system, which will be discussed in Section 4.6.4.2.

4.6.2 Notifications

Users could be notified regarding important events connected to annotations, e.g., the creation of a comment on an annotation they created, which might lead to an increase in their communication awareness. The term 'notification' in this case refers to a popup being shown on the user's screen for a fixed amount of time to alert them to an event and then disappearing again. The need for notifications is supported by Hoff et al. [20] who note that annotation systems should provide a sufficient way to notify users of events. An example they give is a user who would like to be notified of answers posted to a question they asked. Such a user would benefit from being notified by the system of such answers instead of having to check the system periodically, the authors explain. They further note that notifications should be customizable to fit a user's preferences regarding which events they want to be notified of. This suggestion seems reasonable as it could help to avoid users feeling annoyed because the system is sending them too many notifications or feeling that the system has not notified them of an event they considered important. Lecturers represent a group of users that might particularly benefit from customization options regarding notifications as they might want to disable them entirely during lecture sessions, where they might want to use the annotation system for feedback provided by users but might not want to be interrupted in their lecture. In his description of a notification feature that could be added to *Backstage 2*, Mader [31] already mentioned comments posted on users' annotations and annotations created on users' documents as examples of events that users could subscribe to.

One technology that could be used to implement notifications are *Web Notifications* which make it possible for web pages to send notifications to the user that are displayed the same way notifications from desktop applications are displayed [42]. These notifications can be shown even if the web page is not focused and are widely supported among modern browsers [42, 7].

4.6.3 Activity Stream

Notifications offer the advantage of immediate awareness of events that occurred connected to annotations but also have the disadvantage that they do not offer a permanent record of events. Therefore, a view showing a history of events related to a user's annotations, documents, and potentially comments, if actions like voting are implemented, might be a beneficial addition to the system. The user might want to access such a view, e.g., if they missed a notification or want to get an overview of the events that recently occurred. The activity stream would therefore act as a valuable complement to the notification component previously described, as it preserves events in a permanent way and offers an alternative way to access them for users that do not want to receive notifications or cannot receive them for technical reasons. As a result, an activity stream component might help to increase users' communication awareness.

One group of users that could benefit from a permanent history of events are lecturers who might not be able to follow events as they occur during a lecture session but might occasionally want to react to students' recently posted comments and remarks. The activity stream would allow them to deal with all events in chronological order. Functionalities that could be added to make the activity stream easier to use for lecturers include the ability to filter events, e.g., to only view annotation creation events but not comment posting events, and the ability to add a marker to an event to remember that events up to this one have already been dealt with.

Apart from serving as a permanent record of notifications, the activity stream could also be used to communicate events that are not deemed to be as important as the events conveyed via notifications and could therefore be a second communication channel with the user besides notifications.

A similar concept to the activity stream proposed here was already mentioned by Gruschke, as summarized by Mader [31], in the form of a summary of events the user missed in regards to their documents or comments.

4.6.4 Communication Awareness Regarding Questions

One goal for a new annotation system identified during data analysis is to improve users' communication awareness regarding questions asked in the form of annotations. The following section proposes extensions of the current system aiming at doing so. The potential extensions of the system mentioned below might also increase the number of comments created as comments are the medium of communication for providing solutions to questions and the extensions ultimately aim at more questions being answered.

4.6.4.1 Notifications

A possible measure to increase the number of answered questions could be to use the notification system introduced in Section 4.6.2 to call attention to questions that have been unanswered for some time. A crucial question in this context is which users should be notified. It seems reasonable to only notify users in the same course or seminar, as they are relatively likely to have the necessary knowledge to answer questions asked regarding the course or seminar they are attending.

One could argue that among users in the same course or seminar, users that have already answered a large number of questions should be notified as they are likely to be willing and able to answer more questions. A disadvantage of this approach is, however, that relatively inactive users will not be reached, although they might benefit greatly from the positive feeling of successfully answering someone's question. This positive experience might then lead to them being more active on the platform. Based on this assumption, another approach could be to notify users that have not answered any questions yet to try to enable them to have the positive experience of helping someone and to distribute work evenly among the users. A disadvantage when applying this idea could be that these users might not be interested in helping other users at all or might not be able to and therefore will not contribute towards answering more questions. A solution intelligently balancing both of these approaches might yield the best results.

In conclusion, the question of which users should be notified in order to get more questions answered seems to require more research, potentially in the form of a user study where different approaches are compared.

4.6.4.2 Visually Highlighting Questions and Showing Their Status

To improve users' communication awareness regarding questions, it could be helpful to visually mark questions and show their status (answered/unanswered). This could also improve users' awareness of questions without answers, which was identified as a goal that the annotation system should fulfill based on the analyzed data. In a limited form this could be realized for the implemented textual annotations without changes to the static parts of the current system. *Backstage 2* already allows users to assign the purpose of 'question' or the purpose of 'answer' to their annotations or comments [31].

To highlight question annotations, the annotations' content and icon could be customized to stand out. However, currently only icons from a fixed icon set can be selected; using ar-

bitrary icons would require some small changes to the system. An annotation's content is shown in full in detail mode and as a preview in list mode (in the sidebar) and in overview mode (in the popups).

Giving question annotations a special context, however, would require changes to the current system, as the user currently has to choose an annotation's context before selecting its purpose. Therefore, the context renderer of an annotation is already set before its purpose is determined and, as the context renderer in the current implementation cannot access the content data, of which the purpose is part of, the appearance of an annotation's context currently cannot be changed based on its purpose. Furthermore, the appearance of the annotation previews shown in the sidebar in list mode and the popups in overview mode cannot be changed based on an annotation's content currently as well. Changes to the system would therefore be necessary to support accentuated contexts for question annotations. Such changes would have to be carefully considered as they go against the architectural goal of minimizing dependencies between the content and context of an annotation to be able to replace them independently from each other (see Section 4.2.1).

The status of a question annotation could be integrated into its content. To achieve this, question annotations would only have to use a different content renderer than other annotations. Apart from in the body of an annotation's content, the status of a question could also be shown using an annotation's purpose, which could be changed based on its status. The purpose is displayed in the header of the sidebar in detail mode and in the annotation previews in list mode and overview mode. Furthermore, a question annotation's icon could be changed based on its status.

To change the status of a question in the current system, a user would have to edit the annotation's content in the sidebar. The content editor that is shown here could contain a way to select the annotation's status, e.g., through a checkbox. Although this way to mark a question as answered seems relatively complicated, it does not require the user to navigate to a different part of the application as the comments are also only accessible in detail mode. These comments could contain the answer to the question asked in the annotation. However, a potential extension of the current system could consist in allowing content editors to add UI controls to a fixed position in the sidebar in detail mode, similar to the document controls that document viewers can add to the top bar. This would make it possible to show a button to the user that would allow them to mark a question as answered without them having to edit the question annotation's content while keeping the application's flexible architecture. A disadvantage of this approach is that it does not allow the user to mark the comment that solved their problem as, e.g., *Stackoverflow*² does. To add such a functionality, more fundamental changes to the system would be required as an annotation's content data would have to be able to influence the way a comment is displayed.

Assuming a question's status is a property of its annotation's content or part of its annotation's purpose, adding a filter for answered or unanswered questions would only require little effort as arbitrary filters, which can use all properties of annotations including their content and purpose, can be added to the system easily. Similarly, sorting of question annotations by their status could also be added to the system easily.

In conclusion, the current architecture of the system allows for limited highlighting of question annotations, compared to other annotations, and for a limited implementation of question annotations' statuses. Possible extensions of the system could make more extensive implementations of these features possible.

²<https://stackoverflow.com/>

4.7 General Extensions

This section introduces a number of extensions that could improve the annotation system but are not aimed at improving communication awareness or interactions with annotations.

4.7.1 Avatars

A small addition that could be made to the system are user avatars. They could be shown next to an annotation's purpose in all three modes and could lead to increased social presence as they add an additional element, besides the username, that a user can use to distinguish themselves from others.

4.7.2 Roles

As already hinted at in the description of the detail mode, roles could be added to the system. They could, e.g., be document-specific or course-/seminar-specific roles. Examples of document-specific roles are a user that a document was assigned to for peer review and the author of a document, whereas examples of course-/seminar-specific roles are a lecturer and a tutor. A user's role could be displayed in the form of an icon next to their username. Adding such a marking could help users to judge the importance and quality of annotations, e.g., a user could trust an answer given to a question by a tutor more than if it was posted by a 'normal' user. Adding the ability to filter annotations by the role of their author, e.g., only showing annotations created by a lecturer, to the filtering system of the application should only require little effort as arbitrary filters can be added and these filters can use all metadata of annotations, which includes the author of an annotation. *Backstage*, which *Backstage 2* is based on, already marks posts of lecturers and tutors as such by showing a mark next to the author's username [49]. However, *Backstage 2* does not include roles [31].

4.7.3 Moderation Tools

A potential addition to the system that utilizes roles are moderation tools. They could allow users with certain roles, e.g., lecturers or tutors, to execute certain actions not only on their own annotations but also on the annotations of other users. Examples of such actions are deleting an annotation or editing its content. Applications of moderation tools include removing annotations with abusive contents, removing redundant information that leads to a cluttered document, e.g., questions that have already been asked and answered, and removing solutions to homework tasks that users posted on purpose or by accident by asking overly specific questions. To make spotting problematic annotations easier for users with moderation rights, a feature could be added that allows users without those rights to report annotations. Reported annotations could then show up in a list that is available to users with moderation rights.

4.7.4 Visibility Levels

Another potential extension of the system implemented within this thesis are visibility levels for annotations. *Backstage 2* currently allows users to choose between making an annotation visible for all other users or only for the user themselves [31]. Furthermore, Mader [31] discusses adding the possibility to only make annotations visible for a group of users, friends of the user, or an arbitrary group that the user defined themselves. All of these visibility levels would also make beneficial additions to the system implemented within this thesis.

Most of the implementation of this feature would have to be done on the server side of the annotation system as annotations that the user should not be able to see should not be sent to the client. As the current implementation only includes a frontend component that displays the annotation it gets from a source, e.g., a server component, most of the implementation of visibility levels would have to be done outside of the scope of this thesis. However, UI elements for setting and changing the visibility level of an annotation are currently not part of the client component and would have to be added. The visibility level of an annotation could be stored in its context data.

Apart from the visibility levels that Mader [31] already mentioned, an additional setting could be to make annotations visible only for the author of the document (a document-specific role) and the author of the annotation. This visibility level could be used if a user finds a potentially embarrassing mistake in another user's document during peer review and wants to alert the author of the document to their mistake without exposing it to others.

Another potentially useful visibility setting could be to make annotations visible for all users but lecturers (a global role). This would represent a weakened version of the previously explained setting and could allow users to point out mistakes committed by the author of a document during peer review without fearing that they could negatively influence the grade of the author of the document. An issue that arises with this visibility setting is that lecturers might want to check whether a user has participated in the peer review of a document, i.e., whether they have created annotations on a document in seminars where such participation is mandatory. To allow for this kind of check while also offering the mentioned visibility level, a summary of how many annotations each user has created on a document could be added to the system. The calculation of such a summary would have to be implemented on the server side as annotations that should be hidden from lecturers should not be sent to their clients.

4.8 Client-Server Communication

The implemented annotation system is purely concerned with the creation and display of annotations on the client of a user. A server component is not part of the provided implementation, but the client component provides some structures to interact with a potential server component: All annotations are stored in serializable form in a central data storage. This storage is queried to obtain the annotations that are shown in the document viewer, the sidebar, and the annotation preview popups. To change the annotations shown in the system, e.g., to add annotations received by the server, only this central data storage would have to be updated.

Data would also have to flow in the other direction, from the client to the server. Here, the implemented system offers a central access point where components that could be added to the system to communicate with a server component can register themselves to be notified of certain actions. If such an action is executed by the user, these components are notified of the type of the action and receive a copy of the annotation in the state before the action was executed (if possible) and after the action was executed (if possible). Currently the following actions are supported:

- Creating an annotation,
- upvoting an annotation,
- downvoting an annotation,
- adding a comment,
- deleting a comment,

- updating a comment,
- deleting an annotation,
- and setting the content data and content renderer of an annotation (currently executed after its content was edited by the user).

CHAPTER 5

Implementation

The following chapter will give an overview of the implemented system by explaining and showing its different components, which have been described on a conceptual level in the previous chapter, from the perspective of a user of the application. The application was implemented in *TypeScript*¹ using, among other libraries, *React*², *Redux Toolkit*³, and *Bootstrap*⁴.

5.1 Overview Mode

In overview mode, the mode that is active when the user opens the application, the sidebar is collapsed and only the document viewer is visible (see Figure 5.1). In this mode the user can both create annotations, as will be explained in more detail in Section 5.4, and view annotations. For the implemented context renderers and document viewers, an annotation's context is shown either in the form of a polygon superimposed on a section of text or a rectangle. Both new document viewers (see Section 4.2.2) and new context renderers (see Section 4.2.1) can be added to the system, however.

To view an annotation's content the user can click on its context, which highlights the context and opens a preview popup for that annotation (see Figure 5.2). The popup shows metadata on the annotation, namely its purpose, the username of its author, its creation date, the number of up- and downvotes it has received, and the number of comments created regarding it. Some of the metadata is only presented in summarized form in overview mode, namely the creation date, where only the date but not the time is shown, the comments, where only the number of comments but not the comments themselves are shown, and the upvotes, where only the number of upvotes but not the usernames of the users that cast those votes are shown. Some of the displays of metadata also act as triggers for actions: By clicking on the counter for up- and downvotes the user can cast their own vote on the currently viewed annotation and by clicking on the comment counter the user can jump to the comment editor in detail mode to create a comment regarding that annotation. Furthermore, the popup contains a preview of the annotation's content and, if the user is

¹<https://www.typescriptlang.org/>

²<https://reactjs.org/>

³<https://redux-toolkit.js.org/>

⁴<https://getbootstrap.com/>



Figure 5.1: Annotation system in overview mode showing polygon and rectangle contexts on a PDF document (see Appendix A.2 for icons' sources)

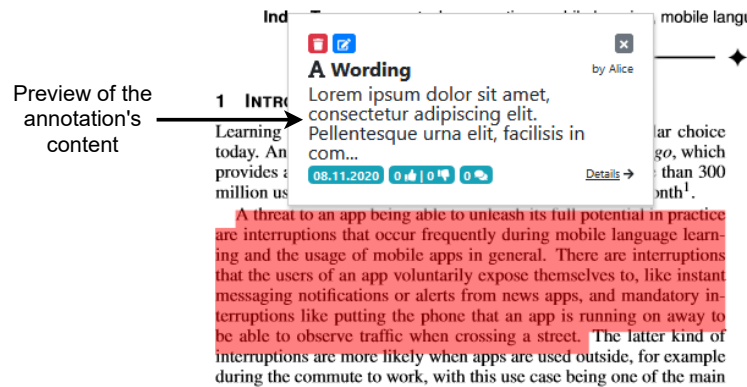


Figure 5.2: Annotation preview popup next to an annotated text section of a PDF document (see Appendix A.2 for icons' sources)

the author of the annotation, buttons for them to delete the annotation or edit its content (shown in the top of the popup). The latter will activate detail mode and show a content editor for the annotation there. The popup also contains a separate link to switch to detail mode in its bottom right corner. Only textual contents were implemented within this thesis but new types of contents can be added to the system by adding new content editors (see Section 5.4). From overview mode the user can get to list mode by extending the sidebar by clicking on the triangle button on the left of the document viewer (see Figure 5.1).

Both overview mode and list mode, which is explained further in the next section, only show annotations from the currently viewed part of the document using a document viewer-specific filter (see Section 4.2.6). For the implemented annotations on PDF documents this results in only annotations from the currently viewed page of the PDF being shown. For the implemented annotations on Markdown documents all annotations on a document are always shown as Markdown documents are not split into parts in the provided implementation.

5.2 List Mode

In list mode the sidebar is extended and shows a list of annotation previews (see Figure 5.3), which are very similar to the ones shown in overview mode. The only difference is that the annotation previews displayed in list mode do not include a row of buttons in the top, which in overview mode contains buttons to close the preview and to delete the annotation or edit its content if the user is the author of the annotation. Therefore, the annotation previews in list mode also contain metadata on the annotations and possibilities for the user to execute actions regarding the annotation, namely voting, switching to the comment creation view for the annotation, and viewing details on the annotation. If a user clicks on an annotation's preview in the sidebar, the annotation's context is highlighted in the document viewer and vice versa (see Figure 5.3).

Above the list of annotation previews, controls are shown to sort, filter, or search the annotations (see Figure 5.4). For sorting the annotation previews, the user can choose from multiple sorting criteria including the time an annotation was created at and the number of upvotes an annotation has received. Furthermore, they can choose between an ascending or descending sorting direction using the blue arrow button next to the dropdown for selecting the sorting criterion. Arbitrary sorting criteria can be added to the system (see Section 4.2.5).

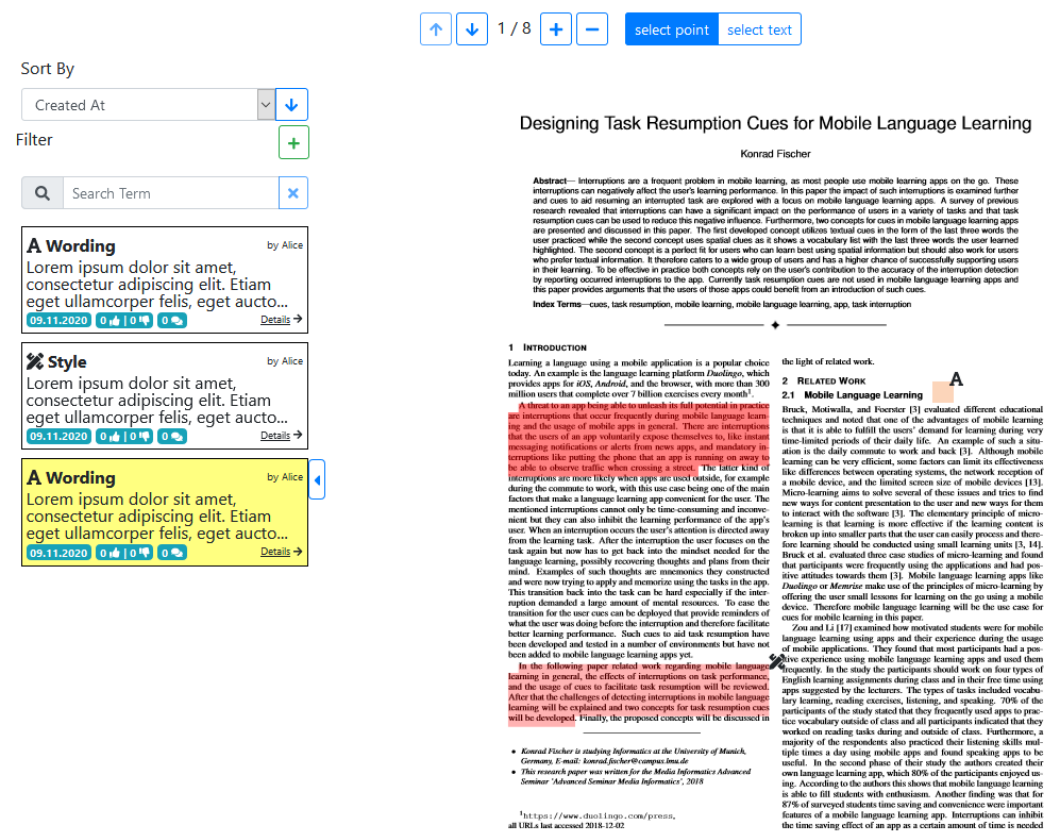


Figure 5.3: Annotation system in list mode showing annotations on a PDF document with one selected annotation whose preview is highlighted in the sidebar and whose context is highlighted in the document viewer (see Appendix A.2 for icons' sources)

The image shows a set of UI controls for managing annotations. At the top, under the heading 'Sort By', there is a dropdown menu currently set to 'Created At' and a blue button with a downward arrow. Below this, under the heading 'Filter', there is a green button with a plus sign. Underneath the filter heading, there are two dropdown menus: 'Purpose' and 'Wording', both currently showing a downward arrow. To the right of these dropdowns is a red button with a minus sign. At the bottom, there is a search bar with a magnifying glass icon on the left, the text 'Search Term' in the center, and a blue button with an 'x' icon on the right.

Figure 5.4: Controls for sorting, filtering, and searching annotations in list mode with annotations currently being sorted by their creation date in descending order and an active filter for annotations with the purpose ‘Wording’ (see Appendix A.2 for icons’ sources)

Users can also filter the annotations shown in the sidebar and the document viewer by adding one or multiple filters using the green plus button. Filters can be removed again using the minus button next to the dropdown that shows the current value of a filter’s filter criterion. If multiple filters are active, they are applied consecutively, i.e., an annotation has to pass all active filters before its context is shown in the document viewer and its preview is shown in the sidebar. The implemented system offers the purpose of an annotation and the type of its context (polygon or rectangle) as filter criteria, but arbitrary filter criteria can be added to the system (see Section 4.2.4).

Using the text field above the list of annotation previews in the sidebar the user can search the annotations created regarding the currently viewed document. Annotations that fit the search criteria are shown in the document viewer and the sidebar while all other annotations are hidden. The search results in the sidebar (see Figure 5.5) include annotations from the whole document, while only results from the currently viewed part of the document are shown in the document viewer. For the implemented annotations on PDF documents this results in the search results in the sidebar including annotations from all pages of the PDF. Apart from the annotation previews, indications of which part of the currently viewed document a search result belongs to are also shown in the sidebar. In the case of the implemented annotations on PDF documents the page of the PDF a search result is on is displayed below the annotation preview. The user can jump to the page of the PDF a search result is on by clicking on its annotation preview in the sidebar when using the implemented PDF viewer. If the user is using the implemented document viewer for Markdown documents, no indication of which part of the document a search result is from is given in the sidebar, as Markdown documents are not split into parts in the provided implementation. The search functionality considers the usernames of the authors of the annotations and the textual representations of annotations’ contents, which in the case of the implemented textual annotations are identical to the content of the annotations but might differ for other types of content. When searching annotations, filters can be added to further constrain the search results, but the sorting order of the results cannot be changed as they are sorted by their relevance regarding the search term. The search functionality is provided using the library *lunr.js*⁵.

⁵<https://lunrjs.com/>

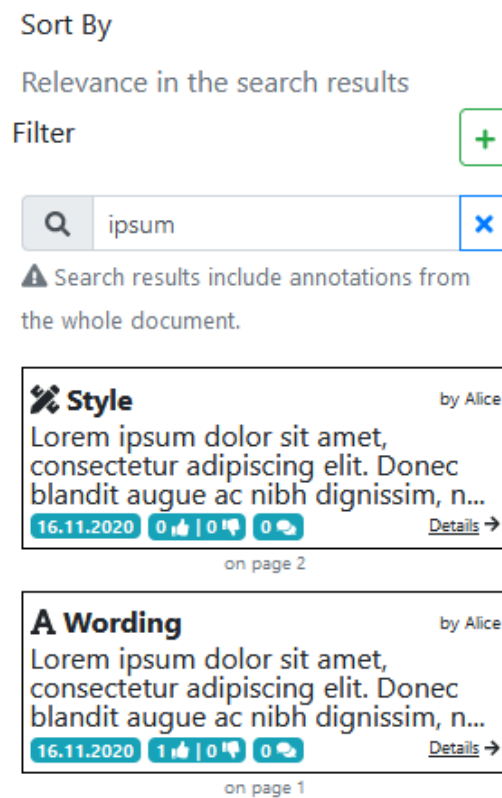


Figure 5.5: Sidebar in list mode showing search results from a PDF document (see Appendix A.2 for icons' sources)

5.3 Detail Mode

Users can get to detail mode by clicking on the link in the lower right corner or the comment counter in an annotation preview in overview mode (see Figure 5.2) or list mode (see Figure 5.3) or by clicking on the edit button above the annotation preview in overview mode. In detail mode, the document viewer is shown next to the sidebar, as it is in list mode, and the context of the annotation that is currently shown in detail is highlighted in the document viewer, as the context of the selected annotation is in list mode. The sidebar in detail mode (see Figure 5.6) shows, similar to the annotation previews in overview mode and list mode, metadata on the annotation with the difference being that some metadata that was shown in summarized form in the annotation previews in the other modes is shown in full in detail mode, including the full date that an annotation was created at. In the top the sidebar also offers a button for the user to delete the currently viewed annotation if they are its author.

The main area of the sidebar consists of two sections (content section and comment section, see Figure 5.6), only one of which is extended at a time. The content section is extended by default and shows the annotation's full content, while only a preview of the annotation's content is displayed in overview mode and list mode. If the user is the author of the currently viewed annotation, a button that allows them to edit the annotation's content is shown. If clicked, this button activates a content editor (see the next section) which is shown instead of the annotation's content.

The comment section of the sidebar shows, if extended, a list of comments created regarding the currently viewed annotation and an editor for the user to post their own comments (see Figure 5.7). If the user is the author of a comment, buttons to delete or edit the comment are shown in its header.

In the bottom the sidebar contains a section that allows the user to vote on the currently viewed annotation (see Figure 5.7). Above the voting buttons the usernames of users that upvoted the annotation are shown. While the usernames of the first three users that upvoted the annotation are always displayed, the remaining usernames are only shown when the user hovers their mouse over the usernames of the first three (see Figure 5.8).

5.4 Annotation Creation

In the implemented system the user can create annotations in overview mode and in list mode. The annotation creation process consists of three steps: First the user specifies the context of their annotation in the document viewer. The type of the annotation's context depends on the annotation mode they have selected in the top toolbar. The implemented document viewers for PDF and Markdown documents allow the user to create a context by either selecting a passage of text (see Figure 5.9) or by selecting a point in the PDF or Markdown document. However, arbitrary types of context can be added to the system by adding new document viewers (see Section 4.2.2).

After the user has selected the context for their annotation, they are presented with a popup where they have to choose a purpose, represented as icons that show the name of the purpose when the user hovers their mouse over an icon, for their new annotation (see Figure 5.9). Arbitrary purposes can be added to the system (see Section 4.2.3).

In the final step, the user then has to enter the content of their annotation using a content editor. The content editors implemented within this thesis are simple text fields that produce textual annotation contents (see Figure 5.10). However, different content editors, and therefore also different contents, can be added to the system (see Section 4.2.3). When the user is done with creating their annotation's content, they can then click the "Save" button and the annotation is complete. Its context will now show up in the document viewer

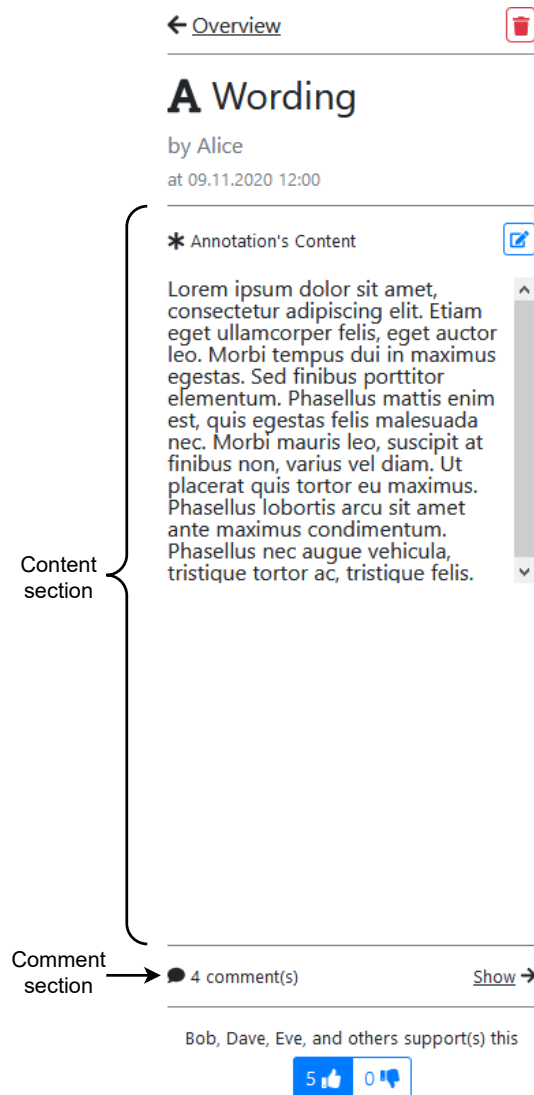


Figure 5.6: Sidebar in detail mode with the content section extended and the comment section collapsed (see Appendix A.2 for icons' sources)

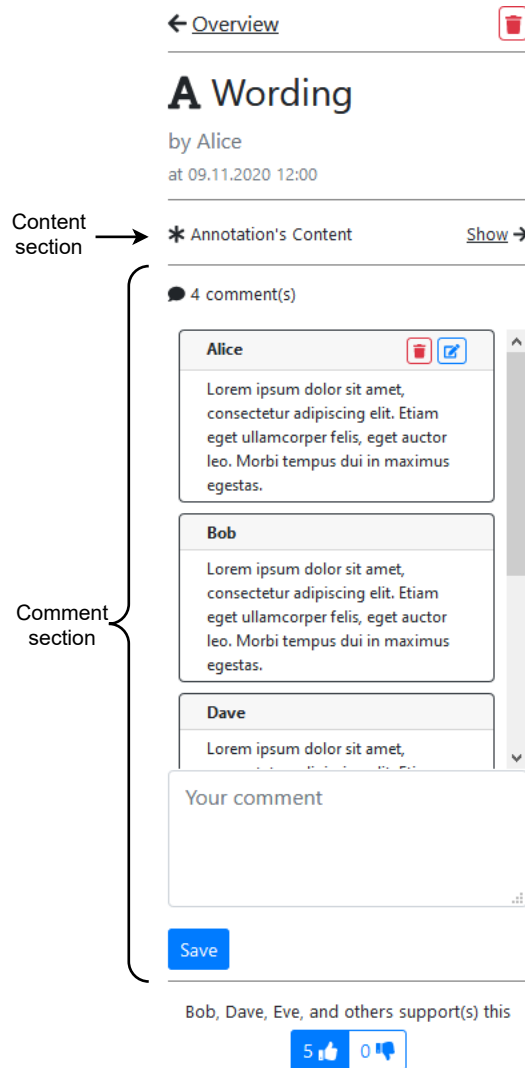


Figure 5.7: Sidebar in detail mode with the comment section extended and the content section collapsed (see Appendix A.2 for icons' sources)



Figure 5.8: Voting section of the sidebar in detail mode in the state that it is in when the user hovers their mouse over the usernames of the first three users that upvoted the annotation (see Appendix A.2 for icons' sources)

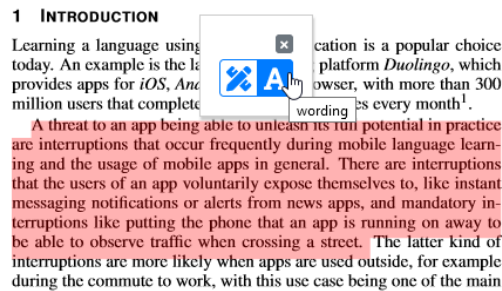


Figure 5.9: Text selection on a PDF document with the purpose selection popup (see Appendix A.2 for icons' sources)

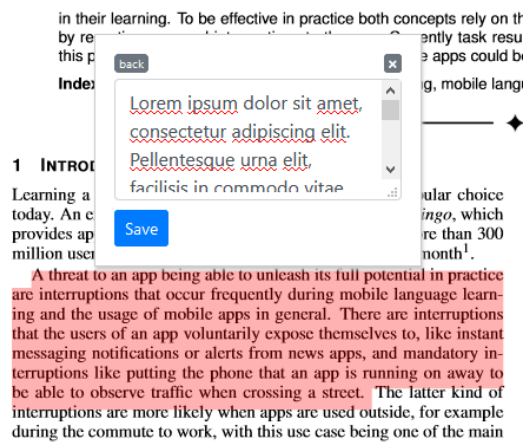


Figure 5.10: Textual content editor on a PDF document (see Appendix A.2 for icons' sources)

together with a small icon representing its purpose (see, e.g., Figure 5.1).

5.5 Document Viewers

As mentioned in the previous chapter, one of the main components of the system are document viewers. Document viewers are responsible for showing the document currently being annotated and allow the user to specify the context of a new annotation. Document viewers for PDF documents (using the *JavaScript* library *react-pdf*⁶ for rendering) and Markdown documents (using the *JavaScript* library *marked*⁷ for rendering) were implemented within this thesis. However, new document viewers can be added to the system (see Section 4.2.2).

Document viewers are not entirely self-contained units; there are some UI elements outside of the document viewers themselves that affect their behavior. One of these is the annotation mode selector, which is located in the top toolbar above the document viewer (see Figure 5.1 and Figure 5.11) and allows the user to choose from a range of annotation modes. In the case of the implemented PDF and Markdown viewers, the user can choose between selecting a section of text or a point as the context of their new annotation. The

⁶<https://projects.wojtekmaj.pl/react-pdf/>

⁷<https://marked.js.org/>




Figure 5.11: Top toolbar above the document viewer (version used with the implemented PDF document viewer) (see Appendix A.2 for icons' sources)

selection of annotation modes is variable, however, and can be adapted to arbitrary document viewers. The active document viewer is informed of the user's choice if they decide to change the annotation mode.

Furthermore, the annotation system offers a way for document viewers to insert their own controls into the top toolbar. The implemented PDF document viewer makes use of this possibility to show its controls for zooming the document and for changing the currently displayed page of the document (see Figure 5.11), while the implemented Markdown document viewer does not add any additional controls to the toolbar.

CHAPTER 6

Conclusion

Within this thesis, related work on annotations, annotation systems, computer-supported collaborative learning, and cognitive load, among other things, was reviewed, followed by an exploratory data analysis on usage data of an existing annotation system, the collaborative annotation system of *Backstage 2*. Based on that analysis, design goals for a new annotation system were formulated. These included:

- making adding new purposes for annotations to the system easy,
- making adapting the system to different use cases easy,
- making it easy to switch between different customizations of the system,
- fostering users' communication awareness,
- and motivating users to interact with the system by creating comments or voting on annotations.

The newly built system was then described on a conceptual level, explaining the ideas behind its architecture, the design of its user interface, and its functionalities. Here the focus was laid on the modular architecture of the system, which allows it to fulfill the architectural design goals defined based on a data analysis. Examples of how this architecture could be used to adapt the system to different use cases were given. Furthermore, potential extensions of the system were described, e.g., further measures to improve the communication awareness of users of the system. Apart from an overview of the system on a conceptual level, the system was also described from the perspective of a user of the application.

Apart from the mentioned concrete extensions of the system that could be added in the future, user studies could also be conducted to evaluate whether the system, compared to the annotation system integrated in *Backstage 2*, improves users' communication awareness and whether it increases interactions of users with the system. Furthermore, the suitability of the different modes the system can be used in for the scenarios they were designed for and users' satisfaction with the system in general could be evaluated in a user study. A UI/UX study could also be conducted to evaluate whether users are able to access and utilize all features of the system. Additionally, experts within the field of education could be asked for feedback regarding the system's functionalities, design, and structure. Apart

from that, the suitability of the implemented system for usage in education outside of the university context, e.g., at secondary schools, and for usage outside of the educational sector, e.g., for working on documents within companies, could be explored. To do so, existing research could be surveyed, experts could be consulted, and user studies could be conducted.

Moreover, the system's design could be improved so as to make it more visually appealing and accessible, and its architecture could be improved further in order to make adding customizations and switching between them easier. Features beyond the ones already mentioned could also be added, e.g., hotkeys to make using the system faster. Furthermore, usability of the annotations generated by the implemented system as training data for machine learning applications could be explored. An application utilizing such data could, e.g., aim at automatically highlighting sections of an essay that are important for understanding its content.

The current COVID-19 pandemic has forced many universities to switch from classroom to online teaching within a relatively short period of time. Rapanta et al. [52] predict, referring to the COVID-19 pandemic, that “[f]or higher education institutions around the world to be competitive (again), evidence of faculty preparedness in terms of professionalism is necessary [and] online teaching is an essential part of such professional preparedness” [52, p. 942]. As annotation systems are a vital component of online learning and online teaching, they will continue to be of high relevance for education after the end of the COVID-19 pandemic and therefore represent a promising area for further research.

Bibliography

- [1] Maryam Alavi. "Computer-Mediated Collaborative Learning: An Empirical Evaluation". In: *MIS Quarterly* 18.2 (1994), pp. 159–174.
- [2] David Barger and Tomer Moscovich. "Reflowing digital ink annotations". In: *CHI '03: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery (ACM), 2003, pp. 385–393.
- [3] Tim Berners-Lee, James Hendler, and Ora Lassila. "The Semantic Web". In: *Scientific American* 284.5 (2001), pp. 34–43.
- [4] D. Grant Campbell. "The Use of the Dublin Core in Web Annotation Programs". In: *Proceedings of the International Conference on Dublin Core and Metadata for e-Communities*. Firenze University Press, 2002, pp. 105–110.
- [5] Neus Capdeferro and Margarida Romero. "Are online learners frustrated with collaborative learning experiences?" In: *The International Review of Research in Open and Distributed Learning* 13.2 (2012), pp. 26–44.
- [6] Michelene T. H. Chi, Robert Glaser, and Ernest Rees. *Expertise in problem solving*. Learning Research and Development Center, University of Pittsburgh, 1981.
- [7] Alexis Deveria. *Can I use Web Notifications?* URL: <https://caniuse.com/notifications> (visited on 2020-10-14).
- [8] Pierre Dillenbourg and Frank Fischer. "Basics of Computer-Supported Collaborative Learning". In: *Zeitschrift für Berufs- und Wirtschaftspädagogik* 21 (2007), pp. 111–130.
- [9] Pierre Dillenbourg, Sanna Järvelä, and Frank Fischer. "The Evolution of Research on Computer-Supported Collaborative Learning". In: *Technology-Enhanced Learning: Principles and Products*. Ed. by Nicolas Balacheff et al. Springer Netherlands, 2009, pp. 3–19.
- [10] J. Euzenat. "Eight questions about Semantic Web annotations". In: *IEEE Intelligent Systems* 17.2 (2002), pp. 55–62.
- [11] M. W. Gallagher. "Self-Efficacy". In: *Encyclopedia of Human Behavior*. Ed. by V. S. Ramachandran. Second Edition. Academic Press, 2012, pp. 314–320.
- [12] Nabeel Gillani and Rebecca Eynon. "Communication patterns in massively open online courses". In: *The Internet and Higher Education* 23 (2014), pp. 18–26.
- [13] Ian Glover, Glenn Hardaker, and Zhijie Xu. "Collaborative annotation system environment (CASE) for online learning". In: *Campus-Wide Information Systems* 21.2 (2004), pp. 72–80.
- [14] Anuradha A. Gokhale. "Collaborative Learning Enhances Critical Thinking". In: *Journal of Technology Education* 7.1 (1995), pp. 22–30.

- [15] Gene Golovchinsky and Catherine C. Marshall. "Hypertext Interaction Revisited". In: *HYPERTEXT '00: Proceedings of the Eleventh ACM Conference on Hypertext and Hypermedia*. Association for Computing Machinery (ACM), 2000, pp. 171–179.
- [16] Jan Grbac et al. "Collaborative Annotation Sharing in Physical and Digital Worlds". In: *Technology and Intimacy: Choice or Coercion. 12th IFIP TC 9 International Conference on Human Choice and Computers, HCC12*. Ed. by David Kreps, Gordon Fletcher, and Marie Griffiths. Springer International Publishing, 2016, pp. 303–313.
- [17] Charlotte N. Gunawardena and Frank J. Zittle. "Social presence as a predictor of satisfaction within a computer-mediated conferencing environment". In: *American Journal of Distance Education* 11.3 (1997), pp. 8–26.
- [18] Raija Hämäläinen and Katja Vähäsantanen. "Theoretical and pedagogical perspectives on orchestrating creativity and collaborative learning". In: *Educational Research Review* 6.3 (2011), pp. 169–184.
- [19] David J. Harrison. "Assessing experiences with online educational videos: Converting multiple constructed responses to quantifiable data". In: *The International Review of Research in Open and Distributed Learning* 16.1 (2015), pp. 168–192.
- [20] C. Hoff, U. Wehling, and S. Rothkugel. "From paper-and-pen annotations to artefact-based mobile learning". In: *Journal of Computer Assisted Learning* 25.3 (2009), pp. 219–237.
- [21] Wu-Yuin Hwang, Chin-Yu Wang, and Mike Sharples. "A study of multimedia annotation of Web-based materials". In: *Computers & Education* 48.4 (2007), pp. 680–699.
- [22] David W. Johnson and Roger T. Johnson. "Creative and Critical Thinking Through Academic Controversy". In: *American Behavioral Scientist* 37.1 (1993), pp. 40–53.
- [23] Roger T. Johnson, David W. Johnson, and Mary Beth Stanne. "Effects of cooperative, competitive, and individualistic goal structures on computer-assisted instruction." In: *Journal of Educational Psychology* 77.6 (1985), pp. 668–677.
- [24] J. Kahan et al. "Annotea: an open RDF infrastructure for shared Web annotations". In: *Computer Networks* 39.5 (2002), pp. 589–608.
- [25] Karel Kreijns, Paul A. Kirschner, and Wim Jochems. "Identifying the pitfalls for social interaction in computer-supported collaborative learning environments: a review of the research". In: *Computers in Human Behavior* 19.3 (2003), pp. 335–353.
- [26] Karel Kreijns, Paul A. Kirschner, and Wim Jochems. "The Sociability of Computer-Supported Collaborative Learning Environments". In: *Journal of Educational Technology & Society* 5.1 (2002), pp. 8–22.
- [27] David G. Lebow and Dale W. Lick. "HyLighter: An Effective Interactive Annotation Innovation for Distance Education". In: *20th Annual Conference on Distance Teaching and Learning*. 2005, pp. 1–5.
- [28] Jian-Wei Lin and Yuan-Cheng Lai. "Harnessing Collaborative Annotations on Online Formative Assessments". In: *Journal of Educational Technology & Society* 16.1 (2013), pp. 263–274.
- [29] Janet Macdonald. "Assessing online collaborative learning: process and product". In: *Computers & Education* 40.4 (2003), pp. 377–391.
- [30] Sebastian Mader. "An Annotation Framework for a Collaborative Learning Platform". Master thesis. Institute of Informatics, Ludwig-Maximilians University of Munich, 2016. URL: <https://www.pms.ifi.lmu.de/publikationen/#DA-Sebastian.Mader>.

- [31] Sebastian Mader. "Promoting Interactivity and Engagement in Tertiary STEM Education using Technology". Dissertation. Institute of Informatics, Ludwig-Maximilians University of Munich, 2020. URL: https://www.en.pms.ifi.lmu.de/publications/index.php#DISS_Sebastian.Mader.
- [32] A. Maedche and S. Staab. "Ontology learning for the Semantic Web". In: *IEEE Intelligent Systems* 16.2 (2001), pp. 72–79.
- [33] Ioannis Magnisalis, Stavros Demetriadis, and Anastasios Karakostas. "Adaptive and Intelligent Systems for Collaborative Learning Support: A Review of the Field". In: *IEEE Transactions on Learning Technologies* 4.1 (2011), pp. 5–20.
- [34] H. B. Mann and D. R. Whitney. "On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other". In: *The Annals of Mathematical Statistics* 18.1 (1947), pp. 50–60.
- [35] Catherine C. Marshall. "Annotation: From Paper Books to the Digital Library". In: *DL '97: Proceedings of the Second ACM International Conference on Digital Libraries*. Association for Computing Machinery (ACM), 1997, pp. 131–140.
- [36] Catherine C. Marshall. "Toward an Ecology of Hypertext Annotation". In: *HYPERTEXT '98: Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia. Links, Objects, Time and Space*. Association for Computing Machinery (ACM), 1998, pp. 40–49.
- [37] Patrick E. McKnight and Julius Najab. "Mann-Whitney U Test". In: *The Corsini Encyclopedia of Psychology*. American Cancer Society, 2010.
- [38] Merriam-Webster.com Dictionary. *back channel*. URL: <https://www.merriam-webster.com/dictionary/back%20channel> (visited on 2020-08-02).
- [39] Evan Miller. *How Not To Sort By Average Rating*. 2016. URL: <https://www.evansmiller.org/how-not-to-sort-by-average-rating.html> (visited on 2020-10-05).
- [40] George A. Miller. "The magical number seven, plus or minus two: some limits on our capacity for processing information." In: *Psychological Review* 63.2 (1956), pp. 81–97.
- [41] Mozilla and individual contributors. *Canvas API*. URL: https://developer.mozilla.org/en-US/docs/Web/API/Canvas_API (visited on 2020-10-17).
- [42] Mozilla and individual contributors. *Notifications API*. URL: https://developer.mozilla.org/en-US/docs/Web/API/Notifications_API (visited on 2020-10-14).
- [43] Mozilla and individual contributors. *XPath*. URL: <https://developer.mozilla.org/en-US/docs/Web/XPath> (visited on 2020-10-17).
- [44] P. Nokelainen et al. "Evaluating the role of a shared document-based annotation tool in learner-centered collaborative learning". In: *Proceedings of the The 3rd IEEE International Conference on Advanced Learning Technologies (ICALT'03)*. Institute of Electrical and Electronics Engineers (IEEE), 2003, pp. 200–203.
- [45] Ilia A. Ovsianikov, Michael A. Arbib, and Thomas H. McNeill. "Annotation technology". In: *International Journal of Human-Computer Studies* 50.4 (1999), pp. 329–362.
- [46] Ozlem Ozan and Yasin Ozarslan. "Video lecture watching behaviors of learners in online courses". In: *Educational Media International* 53.1 (2016), pp. 27–41.
- [47] Mar Pérez-Sanagustín et al. "Discovering the campus together: A mobile and computer-based learning experience". In: *Journal of Network and Computer Applications* 35.1 (2012), pp. 176–188.

- [48] Lloyd Peterson and Margaret Jean Peterson. "Short-term retention of individual verbal items." In: *Journal of Experimental Psychology* 58.3 (1959), pp. 193–198.
- [49] Alexander Pohl. "Fostering Awareness and Collaboration in Large-Class Lectures. Principles and Evaluation of the Backchannel Backstage". Dissertation. Institute of Informatics, Ludwig-Maximilians University of Munich, 2015. URL: https://www.pms.ifi.lmu.de/publikationen/#DISS_Alexander.Pohl.
- [50] P. M. Poortvliet. "Mastery Goals". In: *Encyclopedia of Personality and Individual Differences*. Ed. by Virgil Zeigler-Hill and Todd K. Shackelford. Springer International Publishing, 2016.
- [51] Vincent Quint. *Amaya Overview*. 2012. URL: <https://www.w3.org/Amaya/Amaya.html> (visited on 2020-04-05).
- [52] Chrysi Rapanta et al. "Online University Teaching During and After the Covid-19 Crisis: Refocusing Teacher Presence and Learning Activity". In: *Postdigital Science and Education* 2.3 (2020), pp. 923–945.
- [53] Selen Razon et al. "Effects of a collaborative annotation method on students' learning and learning-related motivation and affect". In: *Computers in Human Behavior* 28.2 (2012), pp. 350–359.
- [54] reddit, inc. *reddit's new comment sorting system*. 2009. URL: <https://redditblog.com/2009/10/15/reddits-new-comment-sorting-system/> (visited on 2020-10-12).
- [55] Alfred P. Rovai. "Building and sustaining community in asynchronous learning networks". In: *The Internet and Higher Education* 3.4 (2000), pp. 285–297.
- [56] Alfred P. Rovai. "Building Sense of Community at a Distance". In: *The International Review of Research in Open and Distributed Learning* 3.1 (2002).
- [57] T. Sannomiya et al. "A framework for sharing personal annotations on web resources using XML". In: *Proceedings of the Workshop on Information Technology for Virtual Enterprises. ITVE 2001*. Institute of Electrical and Electronics Engineers (IEEE), 2001, pp. 40–48.
- [58] John Short, Ederyn Williams, and Bruce Christie. *The social psychology of telecommunications*. John Wiley & Sons, 1976.
- [59] Gerry Stahl, Timothy Koschmann, and Daniel D. Suthers. "Computer-Supported Collaborative Learning". In: *The Cambridge Handbook of the Learning Sciences*. Ed. by R. Keith Sawyer. Cambridge Handbooks in Psychology. Cambridge University Press, 2005, pp. 409–426.
- [60] Jan-Willem Strijbos. "Assessment of (Computer-Supported) Collaborative Learning". In: *IEEE Transactions on Learning Technologies* 4.1 (2011), pp. 59–73.
- [61] Addison Y. S. Su et al. "A Web 2.0-based collaborative annotation system for enhancing knowledge sharing in collaborative learning environments". In: *Computers & Education* 55.2 (2010), pp. 752–766.
- [62] John Sweller. "Cognitive Load Theory". In: *Psychology of Learning and Motivation*. Ed. by Jose P. Mestre and Brian H. Ross. Vol. 55. Academic Press, 2011, pp. 37–76.
- [63] John Sweller. "Cognitive load theory, learning difficulty, and instructional design". In: *Learning and Instruction* 4.4 (1994), pp. 295–312.
- [64] John Sweller. "Element Interactivity and Intrinsic, Extraneous, and Germane Cognitive Load". In: *Educational Psychology Review* 22.2 (2010), pp. 123–138.

- [65] The University of Texas at Austin. *Spearman's Rank Correlation*. URL: <http://sites.utexas.edu/sos/guided/inferential/numeric/bivariate/rankcor/> (visited on 2020-08-12).
- [66] Keith Topping. "Peer Assessment Between Students in Colleges and Universities". In: *Review of Educational Research* 68.3 (1998), pp. 249–276.
- [67] Twitter, Inc. *Counting characters*. URL: <https://developer.twitter.com/en/docs/basics/counting-characters> (visited on 2020-08-05).
- [68] L. S. Vygotsky. *Mind in Society: Development of Higher Psychological Processes*. Harvard University Press, 1978.
- [69] L. S. Vygotsky. "The genesis of higher mental functions". In: *The concept of activity in Soviet psychology*. Ed. by J. V. Wertsch. Sharpe, 1981, pp. 144–188.
- [70] L. S. Vygotsky. *Thought and language*. M.I.T. Press, Massachusetts Institute of Technology, 1962.
- [71] W3C. *About W3C*. URL: <https://www.w3.org/Consortium/> (visited on 2020-04-06).
- [72] Mark Warschauer. "Computer-Mediated Collaborative Learning: Theory and Practice". In: *The Modern Language Journal* 81.4 (1997), pp. 470–481.
- [73] Frank Wilcoxon. "Individual Comparisons by Ranking Methods". In: *Biometrics Bulletin* 1.6 (1945), pp. 80–83.
- [74] Joanna Wolfe. "Annotation technologies: A software and research review". In: *Computers and Composition* 19.4 (2002), pp. 471–497.
- [75] Stephen J. H. Yang, Irene Ya-Ling Chen, and Norman W. Y. Shao. "Ontology Enabled Annotation and Knowledge Management for Collaborative Learning in Virtual Learning Community". In: *Journal of Educational Technology & Society* 7.4 (2004), pp. 70–81.
- [76] Shiou-Wen Yeh and Jia-Jiunn Lo. "Using online annotations to support error correction and corrective feedback". In: *Computers & Education* 52.4 (2009), pp. 882–892.

APPENDIX A

Appendix

A.1 Attribution for Icons Shown in Mockups

The mockups displayed in this thesis make use of the following unmodified icons from *Font Awesome* (<https://fontawesome.com/>) which are licensed under the *Creative Commons Attribution 4.0 International Public License* (<https://creativecommons.org/licenses/by/4.0/legalcode>):

- thumbs-up
- thumbs-down
- comment
- pencil-ruler
- font
- asterisk
- trash
- times
- edit

Both the icons and the license were accessed on 2020-10-11.

A.2 Attribution for Icons Shown in Screenshots

The screenshots displayed in this thesis make use of the following unmodified icons from *Font Awesome* (<https://fontawesome.com/>) which are licensed under the *Creative Commons Attribution 4.0 International Public License* (<https://creativecommons.org/licenses/by/4.0/legalcode>):

- pencil-ruler

- font
- arrow-right
- arrow-left
- arrow-up
- arrow-down
- comments
- thumbs-up
- thumbs-down
- trash
- edit
- comment
- asterisk
- times
- plus
- minus
- exclamation-triangle
- search
- caret-right
- caret-left

Both the icons and the license were accessed on 2020-11-19.