Institute of Software Engineering

University of Stuttgart
Universitätsstraße 38
D–70569 Stuttgart

Bachelorarbeit

Integrating Emotion Recognition from Real-Time Videos into Moodle e-Learning Platform

Stephan Schroth

Course of Study:	Informatik
Examiner:	Dr. Daniel Graziotin
Supervisor:	Lisa-Marie Michels, M.Sc.
Commenced:	November 16, 2021
Completed:	May 16, 2022
CR-Classification:	I.7.2

Abstract

During this thesis, a prototype plugin for the e-learning platform Moodle was developed, which allows students to have their emotions assessed via webcam video during study. Applicability and limits of this approach are then discussed. A feedback cycle with users was conducted to help iterating on the plugin's quality.

Personally interacting with students in class during in-person teaching allows teachers to react to them. Teachers may explain certain things more in-depth, change the teaching approach, as well as reach out to a specific student if they require special help. In e-Learning, such interactions are greatly reduced.

One way to receive feedback on how students are doing can be to use Facial Emotion Recognition through webcam pictures, which also works for non-live teaching sessions. This information can be used by teachers on the online platform as well as automatically by an Adaptive Learning System.

Contents

List of Figures

List of Tables

List of Listings

List of Abbreviations

ALS Adaptive Learning System. 19

AMD Asynchronous Module Definition. 28

FER Facial Emotion Recognition. 17

FR Functional Requirement. 25

GDPR General Data Protection Regulation. 21

GUI Graphical User Interface. 26

LMS Learning Management System. 19

NFR Non-Functional Requirement. 26

UI User Interface. 49

1 Introduction

1.1 Motivation

Education in online learning can often appear like a compromise. It is missing the benefits of direct communication of a human who can react to you personally and change the style of teaching if one approach at explaining proves unsuccessful.

To combine the freedom of e-learning with benefits from face-to-face meetings, Adaptive Learning Assistants can try to bridge the gap and offer variable content reacting to what each individual learner does. They can also easily be used by many more people at once than any teacher can give personal tutoring.

How feasible can it be for a machine to infer information about emotion, and to make useful choices through that to aid a student?

1.2 Problem and Goals

A plugin enabling Facial Emotion Recognition for the e-learning platform Moodle is to be developed for this thesis.

While software systems exist that adapt to what students feel and try to assist them while studying, the creation of an adaptive learning system is not part of this thesis. Instead, this work will consider how well this Moodle plugin and generally Facial Emotion Recognition work for such a system.

The project requires getting to know the platform Moodle and how it works. Software that is suitable for integrating into Moodle for emotion recognition is to be found. A goal of this thesis was also to have at least two cycles of system evaluation with real students. This way, it can be better gauged how good such a system can work today.

1.3 Structure of this Thesis

The thesis is laid out like this:

Chapter 2 – Theoretical Foundations and Related Work: Here, basic terms and concepts will be explained. Afterwards, some literature discussing e-learning and the effect of emotion in studying are looked at.

Chapter 3 – Software Design Process: This chapter describes design information and architecture of both Moodle and the software plugin that was developed for this thesis.

Chapter 4 – Technical Implementation: This chapter looks at smaller details of how the design can be implemented.

Chapter 5 – Feedback Cycle: This chapter first describes the methodology of how a feedback survey about this plugin was designed, then presents and discusses the results.

Chapter 6 – Discussion and Evaluation: This chapter regards all results (the software product and knowledge gained about it and adaptive learning systems. The implications are discussed.

Chapter 7 – Conclusion and Future Work sums up the work and presents an outlook.

2 Theoretical Foundations and Related Work

This thesis discusses the possibilities that Facial Emotion Recognition (FER) brings in the context of using it in adaptive e-learning systems. The relevant concepts and terms for this will be described in this section, as well as related work from similar projects from the past.

2.1 Facial Emotion Recognition

This section briefly explains the basic concepts necessary for a computer system to work with emotions and how a face can be used to gain information about a person's emotional state.

Many different definitions for the concept of emotions exist. [IM19]

Kleinginna and Kleinginna in 1981 said it had historically been hard to find a consensus on what characterises emotions specifically, comparing various different approaches and proposing one informed by them. (see [KK81])

A clear definition of emotions in humans is not strictly necessary for this thesis and project. An intuitive understanding would suffice as long as it is clearly defined both how human emotions are represented in a computer system, and how they can be used for an e-Learning system as defined later in this chapter, as well as how and how accurately the computer system acquires knowledge of a person's emotions.

[IM19] describes that emotion can be described in discreet or continuous models, and they can be regarded with a focus of how they manifest - through a person's body (somatic) or cognitively via thoughts and making judgements.

A common example for a discrete model for emotions consists of a certain amount of basic emotions that are all which is described. Ekman and Friesen have described 6 emotions (fear, anger, disgust, happiness, sadness, surprise) with the argument that these are felt by humanity all over the world [EF71] [Ekm92]. These emotions are often used, according to [IM19, p. 4] due to the "simplicity, interpretability, and high plausibility" of discrete emotion models, despite there also being criticism about it not being able "to capture all human emotions" with no clear consensus existing on an amount of emotions to use.

Continuous emotion models usually consist of two dimensions or more. An example is the circumplex model where various specific emotions can be found along 2 axes representing arousal and valence [IM19].

2.2 e-Learning

Many terms are commonly used in literature for similar concepts to this. A definition for *e-Learning* from [SVC12] that attempts to represent the usage of the term is: "E-learning is an approach to teaching and learning, representing all or part of the educational model applied, that is based on the use of electronic media and devices as tools for improving access to training, communication and interaction and that facilitates the adoption of new ways of understanding and developing learning."

This definition includes cases in which someone studies through media content delivered via the internet without attending any face-to-face class, but can also apply when a teacher offers such electronic learning content together with in-person approaches. [IM19, p. 7] points out that without personal meetings, positive aspects of direct personal communication are normally lost in e-learning, such as recognizing emotions and personality. Ways to adapt to an individual therefore have to be included into such a system in other ways. Creating a plug-in such as described in this paper may contribute to bridge this gap.

The term *blended learning* is sometimes used when combining both face-to-face meetings and a learner studying via e-Learning without a teacher being present. [NPB18, p. 48] The same source mentions that this can also be seen as making use of the positive aspects of both approaches. Also, blended learning may sometimes only be used as a term if both of these aspects make up a substantial part of the study process each.

Other terms are sometimes found in literature being used synonymously or to mean something related to this definition of e-Learning, such as *distance learning*. [MDG11, p. 131/132] For this thesis, any application of e-Learning is relevant and a clear distinction is only important if certain results would only apply to a specific context of applying e-learning.

2.2.1 Web based platforms

According to [AZ05], the term *Learning Management System (LMS)* describes "a software application or web-based technology used to plan, implement, and assess a specific learning process." They point out it is often enabling students to receive class content, while also giving teachers tools to "monitor student participation, and assess student performance online", and it may offer ways of communication and interactivity. Course Management System is another term that can be used for it [AZ05]. Such a system can be used for e-Learning.

In this thesis, a plugin for Moodle is developed. Moodle is an open source LMS that can run on a server as a website for students and teachers to use it[1].

Moodle has a number of features that can be used to create adaptive content, a list of which can be found in [MVTS21] on pages 6 to 10. It can also be expanded with plugins.

2.2.2 Adaptive Learning

An Adaptive Learning System (ALS) is software that changes its content to tend to personal preferences and reactions of individual students, such as adapting the presentation and content of study material to detected knowledge levels [MVTS21]. While an ALS is not meant to be developed as part of this thesis, the resulting plugin should be usable by one. Additionally, the utility of using emotion information received via FER is discussed later.

Changes can be made by such a system about the study material, the sequence in which it is presented, and to assessment (meaning that questions are altered depending on how someone is doing). To make decisions about how to adapt, various factors can be used, including "accuracy of answers, number of tries, use of additional tools, interests of a student (e.g., which subjects they prefer)" p. 6 [MVTS21].

A disadvantage can be that designing courses for adaptive environments can take more time. However, through the adaptability and recommendations from the system, more efficient learning than without it is expected. Because of these points, [MVTS21] suggests a vital part of using an ALS is to train teachers to create good, crative courses for them.

[1]Moodle Website: https://moodle.org/

2.3 Related Work

This section takes a brief look at similar projects about adaptive learning and emotion recognition from literature.

2.3.1 Similar Projects

The survey paper [IM19] has compiled many relevant publications, many of which are summed up here with relevant information for this thesis, with additional related work that was found mentioned here as well, given that it helps deciding which emotions can be used for a possible ALS and what should be expected from one when designing for it. This section also contains some studies about the effect of emotions in e-learning.

It is stated in that survey paper that the "e-learning context is not considered enough" in most emotion recognition methods' design (p. 30). Similarly, for finding appropriate emotion recognition software, the needs of e-learning are often not taken into account.

[May20] points out that finding the most relevant emotions for e-learning contexts is an important question yet to solve. Especially, finding out how learning results are affected by cognitive and emotional processing in relation to e-learning is central.

Mayer also says in [May20] that Control-Value theory is often used in literature, a theory that says perceived value from learning something and perceived control over the study outcome can positively or negatively influence the learning outcome of a student. This theory may be relevant to decide how an ALS can adapt, but may also inform the choice of emotions that are monitored.

This paper says research should focus on negative activating emotions like anxiety and positive activating emotions such as enjoyment. [May20] These terms are in reference to an emotion model considering whether emotions are more relaxing or exciting (activating/deactivating) and whether they are perceived to be pleasant (positive/negative). "Hope, enjoyment, and pride are some instances of positive activating emotions. Anxiety, shame, and anger are some instances of negative activating emotions.", as written about this model in [IM19].

[YK14] note about perceived academic control in e-learning: "The students who perceived high negative emotions with high academic control did not exhibit the benefit of having high academic control. Therefore, instructors or tutors in online learning courses need to make effort to reduce learners' negative emotions and promote their awareness of controllability".

Al-Alwani in 2016 suggested to respond to student's moods "actively in real-time when a mood leading to inattentiveness is detected". For this study, "Happy, sad, confused, disturbed, and surprised" were considered, with a detection rate of over 70 % of emotions in comparison to emotions communicated by study participants once every 10 minutes during e-learning sessions. [Al-16]

[FAM+17] implemented an adaptive learning system, also assuming that negatively perceived emotions are a bad influence for learning, trying to cheer people up through various ways if necessary. This implementation uses a facial emotion recognition library called Rekognition, which has since been acquired by Amazon. Ekman's and Friesen's basic emotions are used for emotion recognition in it, and additionally also a state of neutrality ("calmness")[2]. For this paper and study, the authors decided to "repress" these emotions: anger, sadness, confusion and disgust (see p8 of the paper) They want these emotions: Calmness/Serenity, Happiness, Surprise. Additionally, this system uses the VARK learning style model to personalize the learning experience to users. In a small sample size study, they found better performance when using their system as opposed to using the system without emotional interventions and personlization.

The VARK model used there is however not supported by sound scientific evidence, with the idea of learning styles based on VARK-preferences per person "(Visual, Auditory, Read/write, Kinesthetic)" accordingly not being something that can be expected to help improving student's learning performance. ([KKMK18] and [PMRB08]).

[RB14] presents and discusses evidence that positive emotions have "significant effects on language learners' motivation and learning", as summed up in [IM19] on page 9. According to the Broaden-and-Build theory that is discussed there, negative emotions are bad for engagement and learning, while positive ones increase it and lead to more creative thinking.

2.3.2 Ethical Consideration

The software developed for this thesis saves data on emotion recognition for the intent of applying this for personalised learning experiences. This data is stored and made available to teachers. Same privacy applies here as for other data, meaning when collecting this data, one must fulfil General Data Protection Regulation (GDPR) and other legislation in Germany, the EU and further countries. Moodle has an API for this, simplifying among other things exporting all relevant user data and deleting it [3], that outlines necessary steps to take for this plugin.

[2]Official Documentation: https://docs.aws.amazon.com/rekognition/latest/dg/faces-detect-images.html
[3]Moodle privacy API: https://docs.moodle.org/dev/Privacy_API

Consent is asked before a user's emotion data is being logged. People should know what is going on before engaging with the emotion recognition software product. This can not be entirely disabled, as asking for permission to access a webcam is enforced at least once per browser.

This data can be potentially abused like any data. For example, someone with access to it could make a guess on the personal state and emotions of a user and significantly breach privacy. Seeing someone's behaviour and face movement and inferring emotional reactions is however a normal everyday thing that people can do in face-to-face meetings themselves. The main difference lies in this happening at a large scale automatically and the data being saved, potentially allowing to make a profile of someone's emotional state beyond what is realistic to gain in person. Therefore makes sense to leave student power over whether they want it to be active at any time without forcing it. Privacy laws also already require that deletion should eventually be possible, as mentioned above.

The document [Ass+13] was assessed for any implications it could have for this project as recommended by my supervisor. That document is primarily focused on medical research that directly affects people, for example by testing an experimental medical treatment. While many of the safety and ethics concerns are good to consider, including to be aware of any damage that could be done, phisically or mentally, most of it does not apply to this thesis at all. People taking part in the feedback survey have their webcam monitor their body to receive data that is related to the person's wellbeing, which could cause discomfort. This is especially true if a person is part of a vulnerable group, e.g., depression or other mental health issues. Accordingly, it is offered to all of them to stop participation at any time, as well as to erase any data in an anonymous way. In fact, it was decided to make the entire feedback round anonymous, with a study design that allows a participant to withdraw without anyone knowing who it is, though the small group size and the author knowing the participants means that answers given could give implications on who is who, and who is absent.

Further implications on privacy will be discussed later as part of the feedback survey design of this project, as well as in the evaluation chapter.

3 Software Design Process

This section describes major design decisions and relevant Moodle systems for creating an emotion recognition plugin. The most recent version of Moodle at this time was Moodle 3.11. To define exactly what the plugin should be able to do, a number of requirements for the final software to meet where defined and discussed with the thesis supervisor to make sure the right software is being built. Specific requirements that were discarded are discussed here in-depth as well.

3.1 Requirements Analysis

To design the plugin, at first, a number of user stories where thought up to make sure potential use cases are considered in developing the prototype.

Additionally, functional and non-functional requirements were defined to describe specifically what the final system should be able to do, how it should perform and what restrictions should apply.

A list of all user stories and requirements made can be found in appendix A.1.

3.1.1 User Stories

User Stories are simple sentences describing what a user wants do do with the software system in a certain situation, focusing on the intended outcome and not the way in which it is implemented. In a common template, they state what role a user has in the system, what they want, and why they want it. [1]

These user stories were discussed with the thesis supervisor to receive a better understanding of what the final system should be. This development process was not strictly adhering to any formal Scrum or agile methods that are often associated with user

[1] User Stories explained in the Agile Alliance glossary: https://www.agilealliance.org/glossary/user-story-template/

stories (although weekly goals were defined and worked on throughout the process and communicated with my supervisor, keeping regular progress up). The user stories were used primarily as a tool to phrase any potential goals for the software to discuss and refine them with project stakeholders (i.e., people who are affected by what the result ends up being)[2]. In this case, people who were the stakeholders in the software project to consider were my supervisor and students in the feedback cycle later on. The user stories were also used to create more formal requirements out of this information.

3.1.2 Functional Requirements

Functional Requirements define what specific features and properties the software must have regarding its functionality. This was used for early planning of the software design. These requirements were also used for having a more formal and detailed description of the software system to develop.

The full list of requirements can be found in the appendix. As development progressed, some requirements had to be changed, which is explained in this section.

The following functional requirements were not implemented as initially phrased.

Functional Requirement 5b: Disabling the option for students to choose whether the camera emotion recognition is activated or not was discarded as a requirement. This was initially conceived to give teachers a way to enforce emotion recognition for certain individual tasks where it was particularly important. In real world scenarios, a teacher would be able to react and know the most important times to gauge how their students are doing as well, which may be simulated to some extent with this. This case appears unlikely to be vital for any study session though. Users must give agreement for access to their hardware in home application anyway and can retract it at any time via controls given in any modern browser. Because of this, the feature would be extremely hard to implement effectively.

An approach could have been that the software instantly limits access to the currently loaded page content upon camera deactivation. This is unrealistic for the scope of this project and brings many limitations of its own. Learning content would need to be constantly loaded again in small chunks, as otherwise the entire page content can usually be saved locally easily. Blocking the camera view would also render any attempt to force presence futile, unless constant face recognition was made a requirement, ideally resilient to usage videos or photos disguised as webcam feed. While doable in theory

[2]Brief overview over the term 'stakeholder': https://www.agilealliance.org/wp-content/uploads/2016/01/Stakeholder-Management-Agile2012-Draft-1.pdf

(for example, some plugins exist that check for certain known faces regularly), this use case could also be perceived as disrespectful to the user who may not want to use the camera at every time. Simply asking people to please activate the feature when working on a specific thing can likely work in many cases.

Functional Requirement 5c: Defining situations in which the plugin is not allowed to be used or where it is mandatory. The latter is not necessary with Functional Requirement (FR) 5b being cut. The idea was to automate and enforce when and when not the plugin would be active, in case it was useful to teachers. Forcing it to be inactive at certain spots like when students work on very personal things is a niche use case. No privacy or security risk arises from having it active, other than perhaps when working on personal issues leading to emotions that someone may not want to share. As the plugin can be freely disabled already, this is not required anymore. This was declared optional from the beginning. The plugin has access to the necessary information to know what content is being viewed, so this could be implemented with a blacklist/whitelist added in the plugin settings.

Functional Requirement 6b: At an early point in time, the institute where this thesis was being supervised considered licensing expensive software for emotion recognition, but decided against it. For a prototype, having an open source facial emotion recognition library was already enough to test this. It is possible to replace the current emotion recognition model with a more sophisticated or specialised one later if the results are not sufficient.

Functional Requirement 7: Optionally recognizing emotions from other information in webcam video (such as movement of a person) instead of only using still pictures, is one way in which more sophisticated FER software can use more information to try and achieve better results. [IM19] The chosen library does not support this optional requirement, however. Reasoning for the library selection is found in chapter 4.

Functional Requirement 8a and 8b - It is possible for admins to access the logs containing all emotion data and - via other plugins that allow setting up standard database queries for others to use - people without administrator privileges like teachers can also see these statistics. This could be more elegantly implemented by creating a dedicated GUI that does this all automatically. This is discussed in more detail in the later section 3.3.3.

3.1.3 Non-Functional Requirements

To make sure that any piece of software is usable in the real world, it must also meet certain non-functional requirements such as running on the necessary hardware and software systems and having adequate performance for all defined tasks to actually work.

Accordingly, such requirements were defined for this plugin. The non-functional requirements also include security and privacy requirements, such as Non-Functional Requirement (NFR) 5.

Non-Functional Requirement 6 - Code for implementation of the Privacy API exists for the plugin, but is not fully complete and ready yet to automatically export all relevant collected data as it should due to the task never terminating. At the time of writing, the cause of the issue has not yet been identified. Therefore, taking care of any GDPR / privacy based requests has to be done manually until this is finished.

3.2 Software Architecture

Moodle is built in a modular way, allowing easy expansion with new modules being added[3]. Different types of modules exist, each with own properties and each located in a separate folder. Details about Moodle are defined in its documentation and the existing code.

Of the available plugin types, the "Block" plugin type was chosen for this project. It inherently allows insertion on almost all pages across Moodle, including course pages and specific learning resources that are being viewed. Therefore, a block can be used as a starting point for FER software to be implemented onto each page. The block's position can be configured via Moodle as exists and is usable as Graphical User Interface (GUI) for users to control it and see feedback on what the software is doing.

3.2.1 Moodle APIs

Moodle is written in PHP which runs on a server, creating HTTP pages and managing the database for users connecting via the web. Part of a basic Moodle installation are also a number of PHP and JavaScript libraries that can be called by any newly added plugins as well.

Among the core APIs of Moodle[4] as well as other APIs descibed in the Moodle documentation, the following were identified as particularly useful for this thesis and potential use in an adaptive system:

[3]Documentation on Moodle Architecture: https://docs.moodle.org/dev/Moodle_architecture
[4]Moodle Documentation - Core APIs: https://docs.moodle.org/dev/Core_APIs

- *Events API*: Used for logging various events on Moodle. Can be observed and reacted to by other plugins. The Core API documentation page[4] refers to this as the "recommended form of inter-plugin communication". This could be used to let any third party plugins react to changes in emotions of a specific user.
- *Data Manipulation API*: An API that allows accessing the Server database.
- *Analytics API*: Allowing Moodle to analyze data for various analyses and predictions, which could be connected to a plugin for adaptive content which is to react to certain conditions.
- *Availability API*: This manages user access to certain activities or sections on the website.
- *Privacy API*: Implementing this API allows Moodle to handle data export and deletion requests related to a plugin.

The following APIs were considered for communicating with users. Showing messages to a user could be useful when their emotions have changed and a reaction from the system is required as part of an ALS.

- *Message API*: This allows sending a message to a user, showing up at a bell icon on Moodle[5]. These messages can not be easily read while doing a different task on Moodle (e.g., studying), as they have to be opened in an own page. The notification on the bell icon also does not appear right after the message was sent, though a page load can bring it forward..
- *AMD Modal* module[6]: This can be called from JavaScript code, creating a pop-up on the relevant Moodle webpage allowing to show a message that interrupts normal usage of the website. This may be too intrusive though, students could consider such messages annoying if the emotion recognition system unexpectedly used these and interrupted a study session.
- *Notification API*[7]: Show a message box at the top of the current page, can be called by JavaScript directly.

[5]Message API: https://docs.moodle.org/dev/Message_API
[6]AMD Modal: https://docs.moodle.org/dev/AMD_Modal
[7]Notification API: https://docs.moodle.org/dev/Notifications

3.2.2 Structure of the FER Plugin

The FER plugin developed for this thesis is subject to all the normal requirements and restrictions of a Moodle plugin. The Moodle Documentation describes what the structure of a plugin has to look like.

To fulfil the non-functional requirement number 5 for privacy of the video feed, as well as to minimize load on the server itself, a good approach is to never send any video data to the server. This way, it is easier to keep users' trust in the fact that they are not being watched by any person. Sending video data would not be easy to hide from users bothering to check network traffic, and promising not to do it can give some confidence that having this run in the background for a prolonged amount of time in your own home is not a privacy issue.

To keep all video data on the client device itself, JavaScript based emotion recognition software can be used. The intended way of implementing JavaScript into Moodle plugins is via Asynchronous Module Definition (AMD).[8] Calls of these modules can be defined in a PHP script of the plugin which creates the webpage (or parts of it in the case of a block plugin) for the client device.

The documentation page for blocks[9] details what most of the basic files in the plugin source folder must be for the plugin to work properly wth Moodle. A .php-file called "block_" followed by the plugin name defines the content of the actual block. The folder "amd" contains source code for own JavaScript code that can be called from PHP scripts. Another noteworthy file in the plugin's source code is "fireEvent.php", which is called from JavaScript code to trigger an event on Moodle via the Event API on the server. The folder "lib" includes the integrated emotion recognition library. Further technical details are explained in chapter 4.

3.3 Graphical User Interface

The GUI is what allows anyone using the plugin to interact with it.

[8]https://docs.moodle.org/dev/Javascript_Modules#How_do_I_write_a_Javascript_module_in_Moodle.3F
[9]Moodle Blocks: https://docs.moodle.org/dev/Blocks

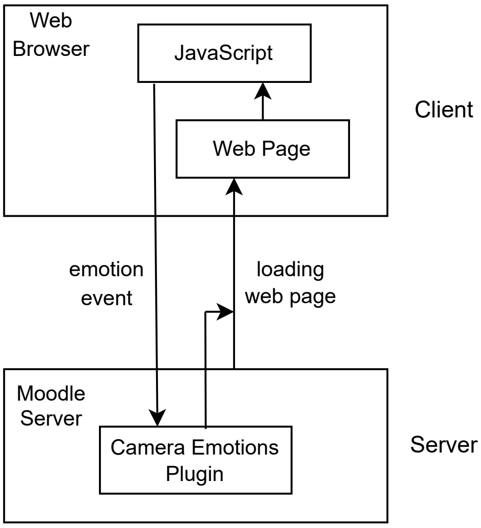

Figure 3.1: Deployment and Basic Structure of the plugin

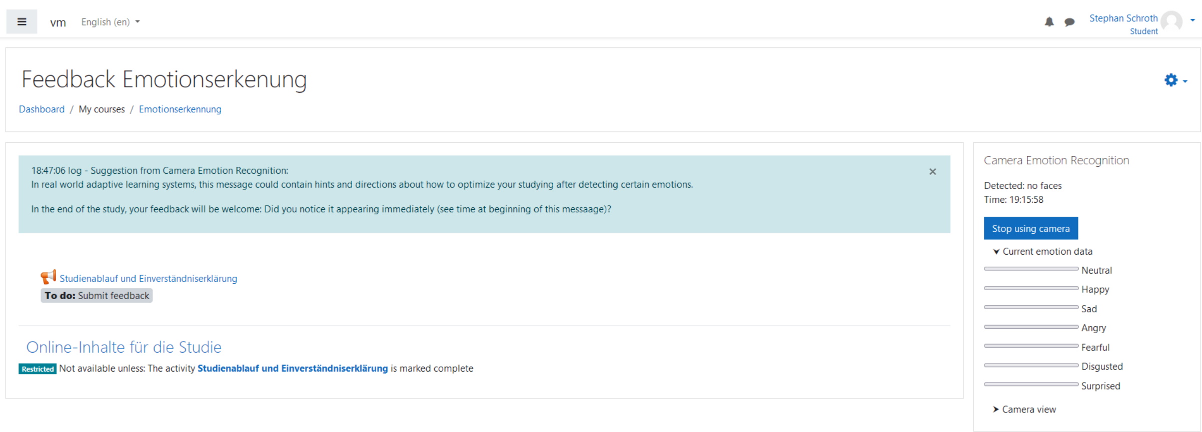

Figure 3.2: View of a Moodle course page with the Camera Emotion Recognition UI visible on the right and a notification message box shown above the main content.

3.3.1 Main GUI

The interface for the user mainly has to allow enabling and disabling the plugin's camera usage for emotion recognition.

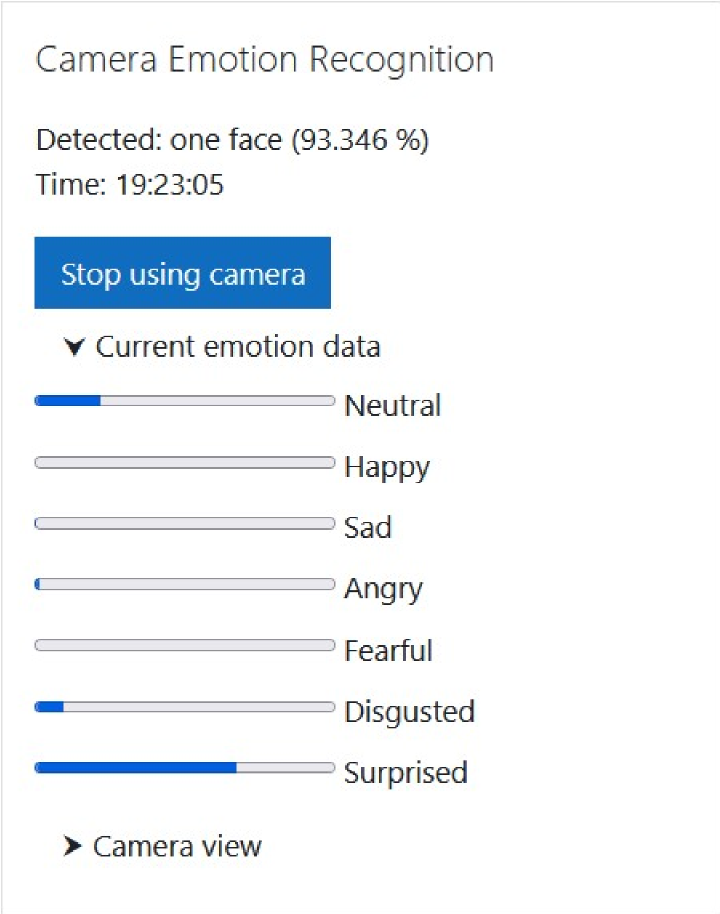

Figure 3.3: The Camera Emotion Recognition plugin's UI showing what emotions are detected by the software

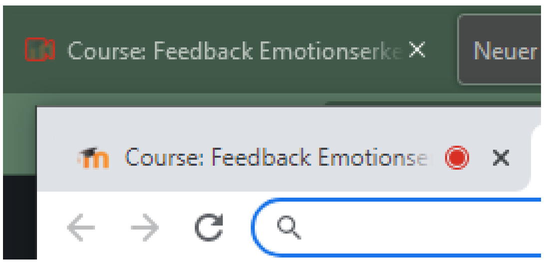

Figure 3.4: Webcam usage indicated by icons on a browser tab (Mozilla Firefox and Google Chrome)

Basic web accessibility considerations are applied to the UI elements[10], making them keyboard controllable as well.

The defined requirements make a button necessary to toggle the camera, which was then given the same HTML category as other buttons in Moodle. This way, the defined visual style of Moodle that is already used for the main web page that the user is running will also be applied to it, making the plugin fit in better. This can potentially help making the plugin less distracting, as it does not stand out from other content by using fundamentally different shapes and colours.

Whenever emotion recognition is turned on and (within the same browser tab) a new Moodle page featuring the plugin is being loaded, it will automatically re-enable emotion recognition 7 seconds after the page has loaded. This is to prevent annoying situations for users where the website can become unresponsive for a few seconds while emotion recognition is starting up. This way, a user who only briefly visits a page to immediately click towards the next will not be experiencing the unresponsiveness. If someone is going to study on the page, the few missing seconds are not likely to be extremely important to be missing.

In order to not make the plugin distracting, the GUI should also not be cluttered. However, to give users more control and a feeling of agency over the plugin, multiple ways of the software giving feedback to the user were added to the GUI. Two collapsible areas are implemented, therefore allowing the GUI to stay small in size by default, but to show as much information as needed if requested by the user.

One of the areas shows the user any current emotions that are detected at that time (Figure 3.3). This gives users a glimpse into how they are perceived by the software, helping them understand what the system and people using it will see. The plugin becomes slightly less of a mysterious black box that users do not understand at all.

Another collapsible area shows the current webcam video, enabling them to see in yet another what the system can perceive. Importantly, users can also check whether their face is positioned in frame and whether the lighting may make it hard to see anything. This should result in a simple and easy to use graphical interface.

A text line always gives users information about the current state of the plugin. When inactive, people can see that on this browser tab, the software is not active. When multiple browser tabs or windows are opened, an icon in most modern web browsers indicates where the active instance can be found (Figure 3.4). When active, the text feedback line will inform users about whether a face is currently detected. This way, it

[10]Accessibility in HTML development on MDN: https://developer.mozilla.org/en-US/docs/Learn/Accessibility/HTML

Listing 3.1 SQL code to receive a list of all stored emotion values usable with the Ad-hoc Database Queries plugin

```
SELECT log.other, log.timecreated, log.id
FROM {logstore_standard_log} log
LEFT JOIN {user} u ON log.userid = u.id
WHERE log.component = 'block_camera_emotions'
AND log.contextid = :contextid
AND log.eventname LIKE '%emotion_switched'
AND u.firstname = :firstname
AND u.lastname = :lastname
```

is easy to check for a user whether they have moved out of the camera frame, or the plugin can't detected them for any other reason. The current time is also always shown as a pulse indicator, showing that the plugin is currently working.

3.3.2 Teacher Interactions

The so called "block" in Moodle containing the main GUI described above can be placed in Moodle via the standard tools given by the platform for adding blocks to pages, allowing them to be shown next to the main content on almost every page. Moodle allows setting up an order from top to bottom in which such blocks are shown, which will take effect if multiple different kind of blocks are present on a page as well. If Moodle is used on a device with a narrow display (like a smartphone) or if the browser window is small, Moodle will show the block below the main content instead.

The plugin is meant to be used by an ALS, and saves the relevant information in the Moodle database on the server where it can be used by other systems or read later. To see the results, an additional plugin that creates a new report page can be implemented to this plugin as well, however. At the time of writing, this can be done via the existing plugin "Ad-hoc database queries" [11] which allows an administrator to set up certain SQL queries that teachers can easily use via a Moodle GUI to get this data.

The SQL code in code listing 3.1 can be used on the current implementation to receive all emotion Data for a certain user.

Wildcarts can be used in SQL in the Ad-Hoc Database Queries plugin to allow entering different names and also filtering by time periods.

[11] Custom SQL Queries Documentation: https://docs.moodle.org/311/en/Custom_SQL_queries_report

Listing 3.2 SQL code accessing emotion data via the Data manipulation API of Moodle in PHP

```
$sql = "SELECT log.other, log.timecreated, log.id
       FROM {logstore_standard_log} log
       LEFT JOIN {user} u ON log.userid = u.id
       WHERE log.component = :modname
       AND log.contextid = :contextid
       AND log.eventname LIKE '%emotion_switched'
       AND u.firstname = :firstname
       AND u.lastname = :lastname
      ";

$params = [
   'modname'   => 'block_camera_emotions',
   'firstname' => 'Adam',
   'lastname'  => 'Example',
   'contextid' => '2',
];


$rs = $DB->get_recordset_sql($sql,$params);

foreach ($rs as $record) {
   $alldata[$record->id] = (object)[
      'other' => $record->other,
      'time' => transform::datetime($record->timecreated),
   ];
}
$rs->close();
```

In the future, a PHP page could be created showing this data in a more feasibly readable way, including graphs. The same code adapted for Moodle's PHP Data manipulation API can be seen in 3.2.

3.3.3 ALS Feedback

Creating an ALS is not part of this thesis. However, to plan for how an Adaptive Learning System can work alongside this plugin and its emotion information, some ways in which an ALS in Moodle could use them were considered.

To prevent patronising students, using suggestions shown as messages to he user was chosen as a very basic ALS interface instead of an approach that forces the student to change what they are doing based on emotion data.

A test message like that is easily addable to the existing system (with no real adaptive functionality). In real adaptive applications communicate to a student any recommendation for other learning content and strategies to use instead, linking to other resources and pages as well.

Having a simple test message implemented into this plugin enables receiving feedback from students about how such a system acts during usability surveying, making the experience easier to grasp for the participants who do not get to use a full ALS but so far only a system that monitors their emotions. Figure 3.2 shows such a message appearing at the top of a Moodle page.

4 Technical Implementation

This section discusses more low-level decisions for the implementation of the Camera Emotions Plugin.

This includes consideration of Approaches for specific algorithms to use. Different possibilities, only said to make as standard as possible, so common method used, but no completely independent standard developed.

While developing the Moodle plugin, Moodle 3.11 was installed locally on a Windows device for reference, testing and debugging. This was the latest version of Moodle at the time this work began[1].

4.1 Facial Emotion Recognition Software

This section discusses why the JavaScript library face-api.js[2] was chosen to perform FER in this plugin, and why.

Face-API recognizes the following emotions, each with a percentage value that describes confidence in having found it: neutral emotional state, happy, sad, angry, fearful, disgusted, surprised. These are the Emotions described by Ekman as universal in humanity, which are commonly used also by other FER software like Rekognition, as discussed in chapter 2.

This library was chosen as it could both be implemented very pragmatically, while also fulfilling all formal requirements that were defined for this project (see Section 3.1 on page 23). It is a library that comes with various face recognition features, including face detection, face recognition and, importantly, emotion recognition on faces (FER). It is noted on the libraries' Git page that glasses in from of a face may reduce the detection accuracy[2].

[1]Moodle Release Documentation: https://docs.moodle.org/dev/Releases
[2] Face-API Git repository: https://github.com/justadudewhohacks/face-api.js

Since it is implemented in JavaScript, it can readily be run in a web browser on the client's device, which is the ideal scenario for privacy reasons. At the same time, the library is designed with performance on low end devices and mobile devices[2], which may be used for e-learning purposes. For example, there are modes of operation for it that require very little bandwidth to be loaded on the client device and run with a more basic, but effective face detection model.

The offered FER functionality working in real time was considered acceptable for this project. Creating an own library in JavaScript, using publicly available datasets, would go beyond the scope of this project and would be unlikely to outperform a popular open-source library like this which is not old yet.

Another hint that face.api.js can be expected to perform on par with other open source projects is a student research project by Figueiredo et al. that was conducted within the institute this thesis was written at. They compared many popular open-source solutions, where this library performed similarly well as most others at recognizing emotions which where compared to what subjects said they had felt, compared over time intervals. This library scored slightly lower results than the best ones, but the difference between these values was small. The mentioned research project is unpublished and not peer reviewed, and cannot be taken as evidence due to this, but as an indication of similar performance.

This library was a good fit to test emotion recognition in Moodle, but has proven not to be ideal for the task of e-learning support, which will be discussed in the evaluation section. When the library was selected, it was considered by the author to be a solid but not necessarily final choice due to it meeting all formal requirements and being an established software, relatively up-to-date, it being released a few years before this project and updated since (last in 2020 for a final release).

4.2 Storing Emotion Data

As defined in the JavaScript file webcam_access.js, emotions are checked regularly as long as the camera emotion recognition is turned on. Only individual snapshots of emotions are detected each time. However, emotions are saved in an array created at a length to fit in multiple seconds of emotion data. The amount of seconds can be changed in the code by changing the numerator value when "*const historyIndexLimit*" is defined.

This approach allows the script to access previous emotion data and regularly compute averages of emotions. This makes minor outlier values less significant if the plugin does not detect a certain emotion for multiple seconds.

The script checks whether an emotion confidence average for any of the seven emotions included in the model changes noticeably from the last time a log was created about an emotional state. The first average values of all emotions are saved to the Moodle database as the first log. Any average after that only creates a new log if at least one emotion varies from the previously saved value for it by more than 20 percentage points. This is when an emotion change is considered to have happened. This low difference requirement helps with emotions that are not going to receive extremely high confidence levels often because they may only be seen with high confidence if the emotion is particularly strong, meaning there would be a danger that subtle emotions could regularly get missed otherwise Only for the value neutral, a difference of 60 percentage points is used, as it is neither an emotion as defined by Ekman, nor is it very rare. Neutral being the absence of any clear emotional tendency appears more often than others and tends to have very high confidence values in this model.

When no face could be detected for a while, a log with all emotion values set to 0 is created.

Experimentation with the event API of Moodle has shown that too many logs being created may not come through. Systematically creating new logs once every second still only results in 10 to 12 logs per minute. This is why the script always waits for at least 5 seconds since the last change in emotions that was logged until it first considers logging an emotion change again.

An alternative approach to saving these emotions could be not to calculate an average of confidence values, but to accumulate occurrences of each emotion that crosses a set confidence threshold. This could be implemented with a higher specific time-frame between logs as well.

But having hard threshold means that anything slightly below goes unnoticed. However, with that approach the script doesn't need a minimum time defined that an emotin needs to effectively last until it is considered to have occurred. so some cases with small changes that could go unnoticed in the averages model would be visible here.

4.2.1 Data Structure

Emotion data consists of 7 percentages that are averages of the past 3 seconds of emotion recognition. Additionally, there is face recognition confidence. This information can be saved in a database using the Data manipulation API.

It can also be saved directly as part of event logs created with the Event API in the dedicated "other" field. This has the advantage that these are created with context information that allows one to know what page a person was on when this emotoin event was logged. Otherwise, this could be saved in the normal database manually as well. For cases in which the saved context may not be exact enough, it is also possible to save the current URL. In this implementation, the data was saved in the "other" data field of event logs.

Because of technological restrictions, no floating point numbers can be saved in the "other" field[3]. To circumvent this, a percentage value can be interpreted as integer with a set amount of decimals that are always enough for the cause used in this case. Percentage values are therefore currently always rounded to 3 decimals.

4.3 Quality Assurance

To enable further development and keep code quality high, comments are found in the code describing each function. The readme file and this thesis also allow getting a quick idea of the important concepts and most of the tools Moodle has that are used.

On local installs, the features implemented were tested. Then, on a virtual machine server made available by the institute ISTE, in collaboration with other students who had already installed Moodle on it, this VM moodle install was changed to use HTTPS instead of an insecure connection.

To have real users connect to the server and have them test the plugin (next chapter), an HTTPS connection is necessary and was enabled for this Moodle install and the server. This is because modern browsers demand a secure context such as using HTTPS to make hardware available to a website[4] [5].

Using this VM server, a feedback survey was conducted with a small amount of real students.

[3]Documentation for the Events API: https://docs.moodle.org/dev/Events_API
[4]Firefox: https://blog.mozilla.org/webrtc/camera-microphone-require-https-in-firefox-68/
[5]Chromium: https://www.chromium.org/Home/chromium-security/deprecating-powerful-features-on-insecure-origins/

5 Feedback Cycle

A feedback survey was conducted between four participants. The goal of this feedback cycle was to receive feedback on the plugin's usability for further iteration of the software development process. Additionally, (non-representative) surveying of how the concept is perceived by students was used to assist in prioritisation and designing further features and assessing the applicability of the concept.

This feedback was only one cycle. After the first one, not much time was left and not many changes were yet to be made to get use out of a second iteration, even though one was intended initially.

5.1 Survey Design

All four participants were students from the University of Stuttgart from various courses of study, two male and two female, aged somewhere between 18 and 27. They were given anonymous accounts for an install of Moodle 3.11.5 (Build: 20220117) running on a virtual machine on a server of the University of Stuttgart. They were asked to connect to the site from their home, with a VPN client giving them access to the university network in which the site was hosted.

Initially, participants were instructed on how to activate the plugin and to check whether the software could detect them, with instructions to make sure enough light was covering their face and they don't move out of frame too often. They were also instructed on how to use the plugin.

They were given two simple maths problems to solve while making sure they keep the plugin recognizing them. This was to teach the participants to remember not moving away from the frame of where the webcam could see their faces even when concentrated on working on something. Any potential problems of the FER system not recognizing them should therefore be detected in this stage. This is especially important since the study was conducted without supervision by anyone. Participants were told that working on these maths problems was more important than getting correct results, to ensure this step of working on it is not skipped (e.g., by using an online tool to solve it).

To simulate working on online learning content, a short lecture video from a real online-only university lecture produced during the COVID-19 pandemic was shown to the participants. A brief quiz asked questions about what was taught to further simulate university work.

A second, longer video about the same topic is shown as part of the survey next. This time, users weren't quizzed on the content anymore, but they may have been more attentive to prepare for any possible future questions they might have expected due to having just completed such a quiz before. This was hoped to help get users more engaged with the content with the goal to be able to know enough to answer questions about it afterwards, similar to how a student may approach lecture videos for their own studies.

Instead, the survey after the video focused on usability feedback and how comfortable they were using this plugin prototype as well as rating some use cases as defined in the design process. This is the core part of the study, the simulation of work previously having prepared the users for it. For usability feedback, questions about the plugin itself, and questions from the feedback grid[1] were used. They were available to participants at the end of the survey. Other questions in the survey asked for answers given on a numerical scale usually consisting of an even or odd number of options, depending on whether neutral answers were intended to be possible to that question.

Both the survey in the end as well as the quiz in the middle of the study's time frame also ask participants which emotions they felt while having worked on the study content so far. It was not a focus of this survey to verify that the measured emotions exactly match those reported. However, the emotions that are found overall should be similar in both, and it may be worthwhile discussing why any discrepancies exist and use this for further development of the software and its concept. Asking for emotion feedback in the middle of the feedback survey's time frame as well as in the end means fewer emotions will have been forgotten by that time as opposed to only asking once in the end. Also, clearer comparisons can be made to what was logged by the plugin automatically in that time frame. This was done similarly in [Al-16, p. 450] where 10-minute time frames were used to assess emotion and the participants were asked to note at which part of the lecture this occurred.

Whenever a user had spent more than 3.5 minutes on a single page during the study, a message box as a feedback interaction from an ALS as described in section 3.3.3 was shown, prompting them to remember their impressions of the message appearing on the page while they may have been working on it for later feedback.

[1] Feedback Grid: https://nativdigital.com/feedback-grid/

Before starting the study, informed consent was received via a text describing the study content. Each participant was automatically granted access to the main study content only after accepting an agreement found in the appending along with the survey questions. This agreement is based on [HSE13, pp. 46-47]

The same book also recommends on p. 71 and later that, if personal opinion or sensitive data is asked from people, anonymity should be made an option to whatever extend is possible. Most things asked are not sensitive, but emotion data once again could be considered that. It was decided to make the evaluation and survey as anonymous as possible. The identification number used for participants in this chapter is also not identical to the number given to their accounts.

Since every participant's user account is fully anonymous, other students (or the administrator) with root access to the server at the institute can not gain any information about them.

5.2 Results

The feedback from all participants on each topic will be summed up and emotion data from the plugin is presented.

5.2.1 Feedback Results

For usability evaluation, participants were asked what they liked about their experience of using it, what they were wishing for, what things they found themselves wondering and what ideas may have come up while using it (see Feedback Grid[Footnote 1 on page 40]).

This is free text input for the participants. They were prompted share their thoughts about the emotion recognition interface display and about using it. Examples were given for what could be discussed: clarity in usage, information shown by the software, place of feedback messages, which was presented as a non-comprehensive list of examples.

It was reported by three of the four participants that they liked that detected emotions could be seen on the interface. It was also described as easy to use by two of the participants. The fact that face recognition success confidence was shown was also positively mentioned by one person.

Participant 4 wished for more accurate recognition of sadness, reporting noticing false positives of it when they were really feeling confusion. Participant 2 wished for more

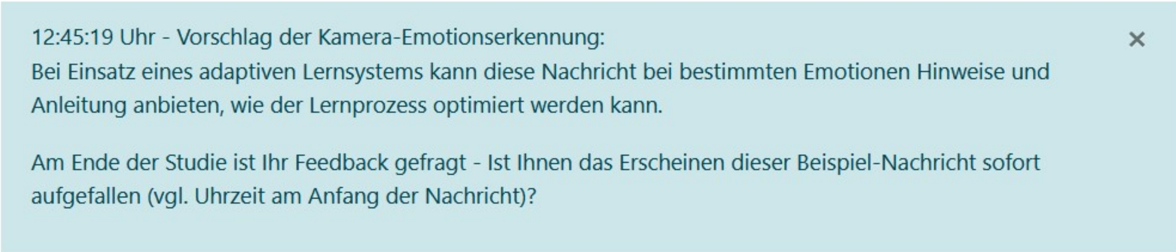

Figure 5.1: A message box in German appearing in Moodle.
It is part of the evaluation cycle that was conducted, for which it is telling users to remember it. The English localised version can be seen in context in figure 3.2.

emotions to be supported by the systems, noting some of the emotions shown not being the same ones they felt.

Participant 4 mentioned they were wondering how accurate the software is.

One participants presented the idea that an option to see themselves through the camera's perspective while the plugin could help with setting up a good position in the camera's view, but might also be distracting while trying to study.

In text feedback about the messages usable for a simple ALS integration (Figure 5.1), the participants reported they had seen the messages on the video pages immediately with it not disturbing them, only briefly breaking their concentration while they paused to read it. One participant saw it only on the second video page and missed it during the first one.

Participants were usually quick enough with filling in the feedback and getting to the next page according to Moodle logs so that the messages did not appear on those. No one except participant 3 saw a message appearing on the feedback page, where it did not disturb that participant according to their given answer.

Asked about whether the plugin being active in the background or the user interface in general was distracting people, participants answered with 0, 1, 1 and 2 indicating little distraction on a scale ranging to 5 and, if there was distraction, further saying that it occurred only very rarely.

All participants responded to questions about their positioning that they changed their posture or position initially to be in the frame, with participant 2 not thinking about it anymore later.

Half of the participants answered that being in sight of a camera for this plugin to work led to them behaving differently. One said this was related to being more concentrated from being watched. The other person wrote that it made them thoughtful about how they reacted to things, always looking at their reactions in the GUI.

Asked how much they felt under watch by the software on a scale from 0 to 5, one person responded with 5 ("very"), one answered with 3, and two with 2.

The answers came from laptop or desktop computers equally often running Windows and macOS using the browsers Safari, Brave or Firefox.

Aked about what other things were noticed that were different from ordinary e-learning experiences (the question prompt giving the example of device reaction times, performance, or own feelings), nothing was put forward other than one person reiterating they were a little more nervous than usual due to feeling they are under observation.

Participants were asked about specific features or use cases for application of this plugin in real-world scenario.

When asked how important it was to the participants to have full control over when the camera is active for emotion recognition, three of them replied with the highest number of the used scale labeled as "very" important. One participant responded with the number 1, which is one off the smallest value of the scale which would be "not important at all).

Participants were very generally asked about these scenarios: A teacher uses emotion data to see that a specific person has issues while learning and the teacher can react to this. This was rated on a scale from 0 (very bad) to 6 (very good) with 3 being neutral. Participants rated this with 3, 6, 0 and 2.

A focus on positive emotions was in the question of the scenario: A teacher sees emotion data that shows which parts of an online lecture are a success with certain students and can thus improve on the course with this feedback. This was rated with neutral (3) as well as 4, 4 and 6 (very good).

A system detecting automatically that a student needs help with frustrating content and acting on it was rated with 1, 2, 4, 4.

5.2.2 Emotion Log Results

The table 5.1 shows that the amount of instances when the camera was unable to detect a face (for the majority of a 3 second time frame) is making up just under 4 % overall of logged time frames.

In the table 5.2, it is shown that every emotion was at one point detected with a confidence value between 95 and 100 %, except for disgust. Most of the emotions were

Emotion Logs:	All	Participant 1	Participant 2	Participant 3	Participant 4
Total	1830	505	372	703	250
No face found	72	20	5	40	7
	3.93 %	3.96 %	1.34 %	5.69 %	2.80 %

Table 5.1: Recorded emotion logs compared to the number of cases in which the applied face-api.js software did not manage to detect a face.

	Mean	Median	Maximum	Lower Quartile	Upper Quartile
Neutral	0.87047	0.996667	1.0	0.95529	0.99960
Happy	0.0128	0,00003	0.999 9	0.00001	0.00024
Surprised	0.01919	0.00010	0.95423	0.00001	0.00108
Sad	0.05070	0.00043	0.99995	0.00003	0.00629
Angry	0.00436	0.00006	0.99942	0	0.00059
Fearful	0.05070	0	0.99926	0	0
Disgusted	0.00206	0	0.44998	0	0,00001
Face Detection	0.83679	0.90898	0.9987	0.79747	0.96269

Table 5.2: Face Detection and Emotion Recognition confidence described through auxiliary values

however rarly receiving values far from 0, except for neutral, which dominated the statistics.

Figure 5.2 shows that filtering the logged emotions further for which emotions go beyond a certain threshold will eliminate different emotions when raising the threshold from 0.2 to 0.5 with all emotions becoming less frequently represented, but neutral losing little in relation to its overall quantity of logs. Meanwhile, other emotions like happiness or sadness get halved while disgust gets eliminated.

Focusing on emotion changes gives smaller ratios between emotions, mainly by eliminating most of the neutral logs.

Figure 5.3 is split in session A and session B, with A being until the first emotion questionnaire, and B from then to the end of the survey and compares all emotion logs (left) to the reported emotions (middle) to emotion changes deteted (right).

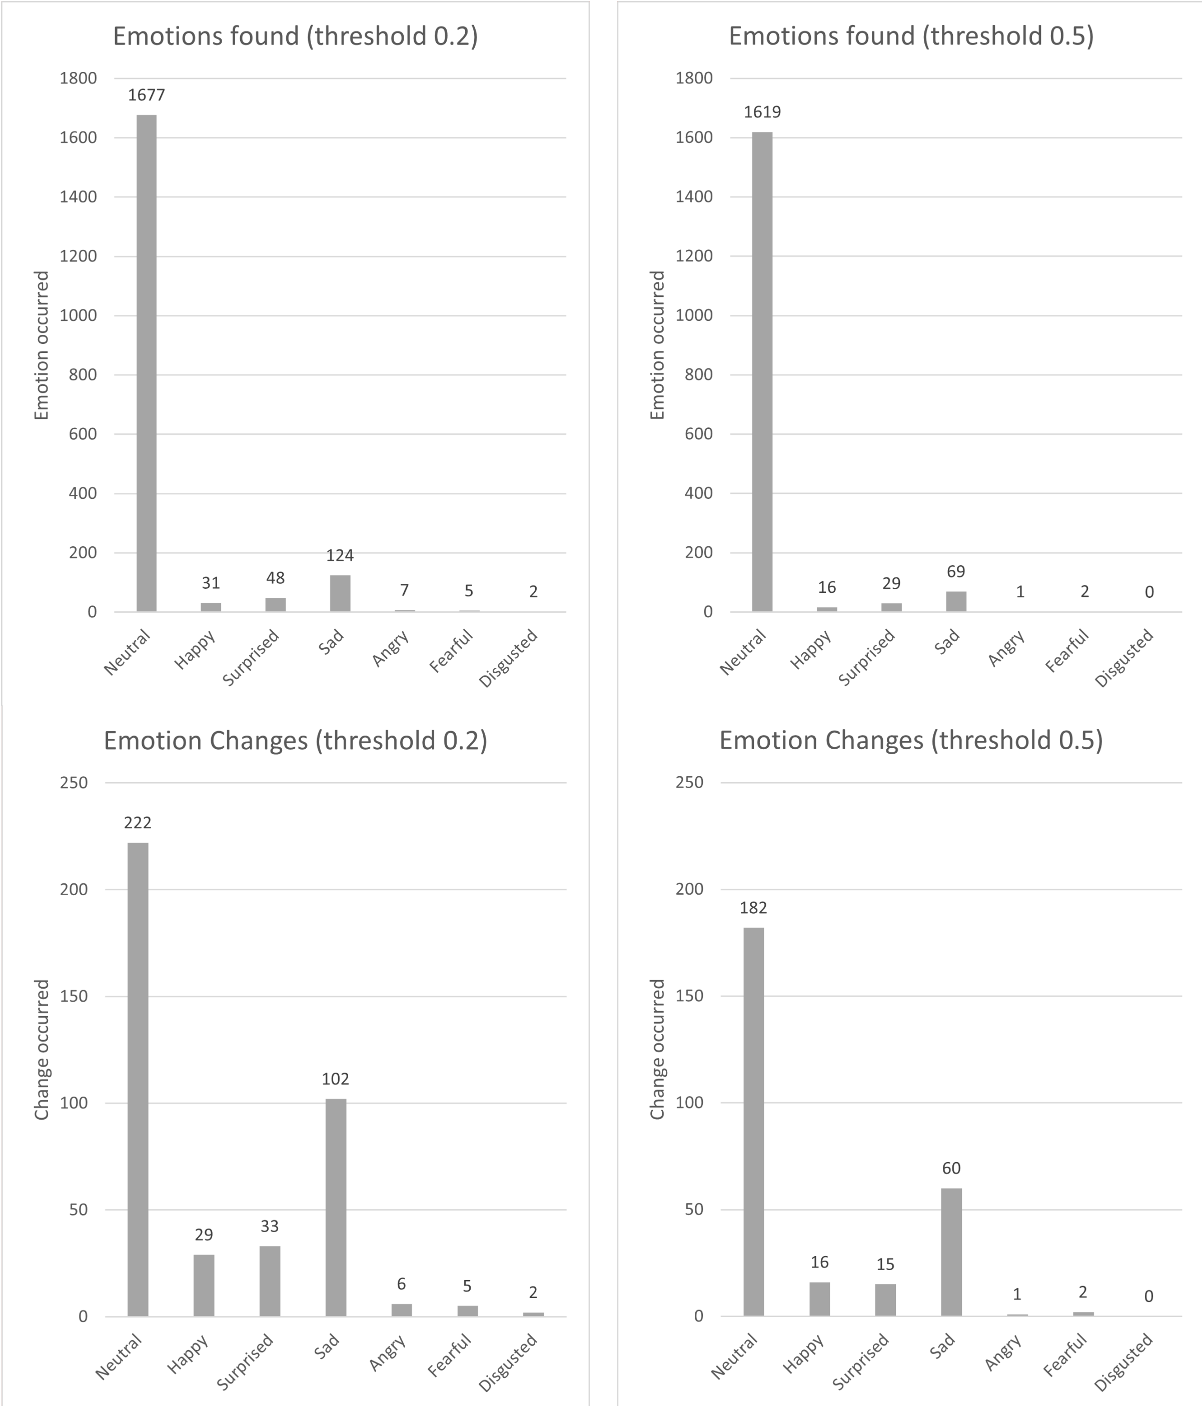

Figure 5.2: Comparing all detected emotion instances with all detected emotion changes. Detected emotion data where the average detection confidence value in a 3 second time frame was over the given threshold is included.
Emotion detection confidence can be values between 0 and 1 in this data. Changes are defined in the software code as the average confidence having a difference of 0.2 from the previous change or from 0 if none exists. Neutral required a difference of 0.6 instead, as described in section 4.2.

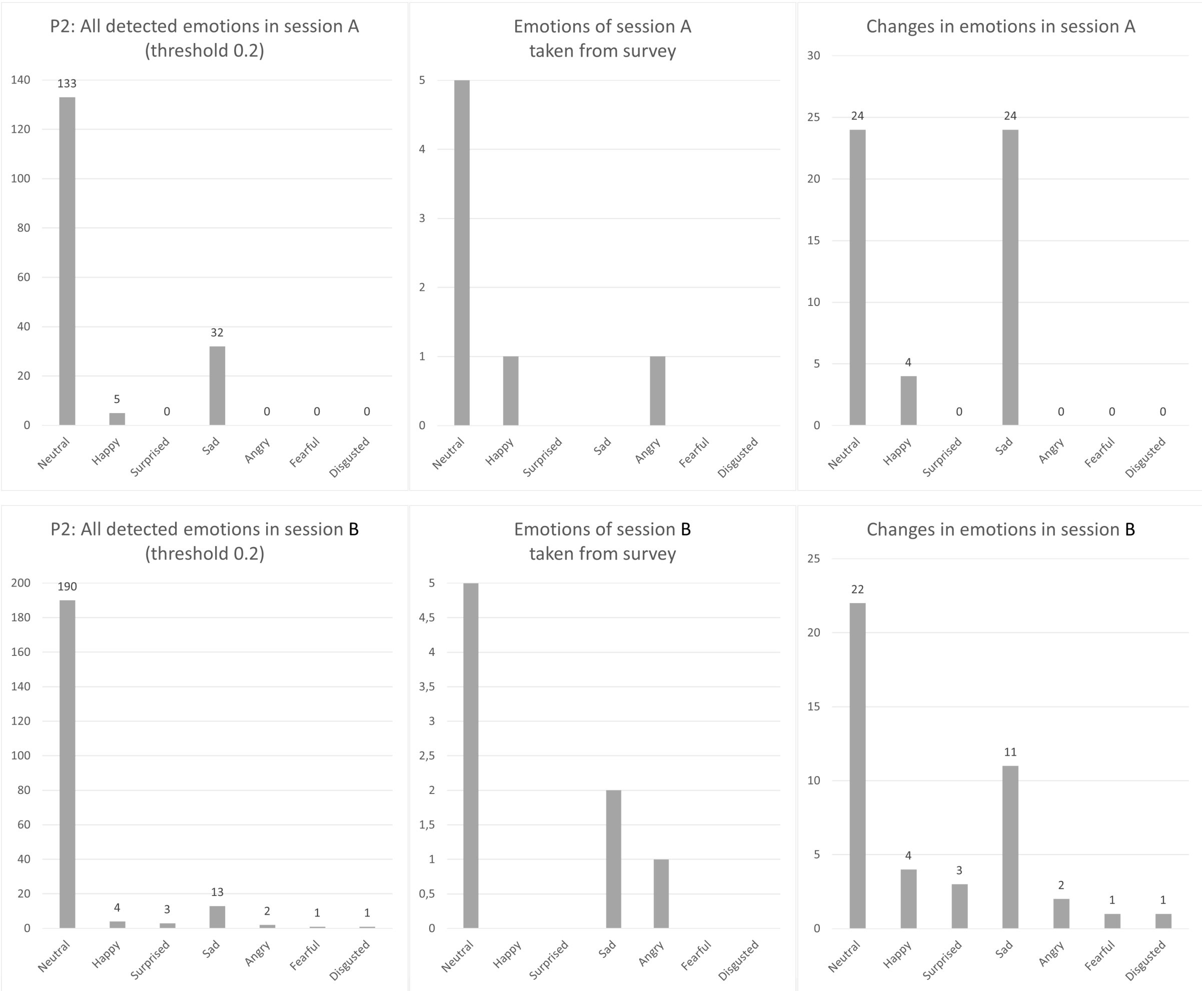

Figure 5.3: Emotions detected versus emotions given directly in survey versus emotion changes detected for Participant 2.

6 Discussion and Evaluation

6.1 Feedback Cycle Evaluation

The results and methodology in the previous chapter are evaluated.

6.1.1 Emotion Log Discussion

The results in table 5.1 show that the face detection was usually able to find faces. Even though light and moving out of frame could be a problem, there were few instances when this prevented the FER system from working.

Statistics of detected emotions are inflated by neutral states reducing other emotions to a small role in comparison. Data of emotion changes compared to all emotions shows that there is less redundancy, since "neutral" as an emotion often repeats.

Change detection works as expected, fewer occurrences have to be saved in the current implementation where only changes are saved, since the same emotion can stay currently active for a prolonged time. But this affects "neutral" more than any other emotion. The occurrence of neutral being greatly reduced means there is less overhead, and more "real" emotions are considered in relation. This approach gives other plugins as much information about the emotion as they may want to deal with the information themselves.

It would make sense to remove "neutral" and only log every instance of other emotions, instead of relying on emotion changes to mark the end of the emotion that started with the previous log. Neutral does not hold much usable information that could not also be inferred from the absence of finding any other emotions (which is, that there are no noteworthyly strong emotions at that time).

While 5.2 shows the method of logging changes reduces the amount of logs a lot and gives a barely comparable scale between the remaining emotions, it may be more interesting to quantify directly how often a certain emotion appeared instead of only noting down changes to it. It may therefore be a good choice to use low thresholds when later selecting from the emotion logs while not saving neutral values.

6.1.2 Feedback Discussion and System Applicability

While feedback on usability was overall positive, one participant missed both the explanation in the beginning of the survey process saying that the own video could be seen by clicking a specific UI element. This participant also did not discover this by themselves until at least the end of the study. They suggested adding it for help with setting up the environment for the camera to work well, but proposed that this could lead to distraction during studying. Both these points were considered while designing the interface.

Multiple participants gave feedback about the emotion recognition not necessarily being accurate, which has also been the authors perception. The suggestion of improvements to emotion recognition appear adequate. In particular, the claim that sadness is overrepresented in comparison to given values of the participants can be seen in the data as well (see also appendix for emotion data of all users).

No general major issues came up during the test. No further comments about what was noticed about using their device in comparison to normal e-learning also imply acceptable or good performance on the devices with no significant disruptions. The strategy to wait for a few seconds until emotion recognition is activated again may work to prevent a brief lag spike just when a user is starting to use a newly loaded page, or it might not be occurring for them at all.

Since it was reported that the distraction was little and rare, the general face recognition UI appears to work as intended.

While participants generally listened to the advice about making sure there is some light on one's face and to be in the camera view, one participant said they did not look after this later on in the survey. This was participant 2, who also had the lowest rate of logs with no face being detected. Participant 1 also did not consider this much, but has only got the second highest ratio there.

This may suggest that the face recognition worked very well overall and bad angles of the face or obstructions in front of the face were no problem.

From feeling supervised, some participants mentioned being more concentrated and more nervous. People's responses to being "watched" especially in the later use cases seemed mixed. This system was used only for the specified goals and their data was anonymous, which lead to no general ideas or questions or complaints about the concept in feedback. But the different scenarios were rated in a mixed way, and had big differences, with some people seeing no problem in what others didn't like.

People might get used to it later or decide they don't want this regularly at all. Results from the feedback are not representative due to the small sample size and more reliable and long term data is clearly required to draw any conclusions.

6.1.3 Limitations and Reproducibility

While the low number of participants makes completely anonymous answers hard to achieve, this was still considered desirable to show good faith and respect for their privacy when they don't initially know what kind of emotions may be recorded. This way, the most authentic responses could be achieved for calibrating what emotions are detected, especially when participants are asked about their own emotions for comparison. Achieving this could be done feasibly by giving participants randomly drawn account credentials, so only they know which account they use. Moodle logs a lot of information about any user visiting the website. Due to users connecting to the site only via VPN, the only identifiable information could be in writing style (which can be countered by people being aware they should not write in a clearly identifiable way whenever a text prompt would appear), as well as how they did in the quiz. Having multiple people with similar backgrounds would have added more ambiguity to who would be more likely to give correct answers quickly. This limits the degree of anonymization somewhat, but not entirely.

For better comparability with real life application, assuming that emotion data is not anonymous there, this could be repeated without anonymous accounts

Inviting participants to just work on any of their own (real) university tasks and lectures while having the plugin active was considered to get more realistic results. However, due to the small sample size, it could not have been a representative result either way. For the sake of maintaining reproducibility, the study would have had to rely on people having similar work to do in any replication study, too. Therefore, clearly defined lecture content was prepared for the study instead.

One participant had a background in software development, meaning they will have known about the lecture content shown during the study before. All others were from different fields of study and therefore had to work on this with very little background knowledge. The chosen lecture video allowed for this scenario by mostly discussing concepts that require little prior knowledge.

The feedback on the User Interface (UI) gave participants the option to get an idea of what data was stored on them. If this ever were to affect how students felt during the study, it would not make the results less comparable to real world application as long as users there have access to such information as well.

However, the current approach allows the script itself to notice changes in emotion and react to them. This could be useful for an ALS, where such changes can be monitored and then be used to inform the system.

Combining event logs and emotion data directly is straightforward, but has turned out to have disadvantages as well: floating point values can not be stored there, leading to small inconveniences in having to convert them to integers and lose some precision in the process. However, this precision is not needed. For the future, creating an own database and only linking exact emotion data to it in the log may be a more elegant solution though.

No standard exists for communication in e-learnig - literature shows that different systems are used for emotions and no consensus exists on which ones to use. There is no clear way o say that an ALS should expect one specific standard of emotions. But rough tendencies, like positive or negative emotions can generally be detected and used.

6.2 Performance and Scalability

System evaluation participants had no negative remarks regarding performance. The test was short and conducted with only four participants, meaning real users running it regularly for a long time might have more observations.

Moodle can be used via the web on most devices, including mobile devices as well. The plugin was able to function on a more than five years old smartphone running with Android 7 as part of testing the system with no noticeable negative effect during the brief use of emotion recognition. Exact performance benchmarks to assess, e.g., whether the emotion recognition application lowers battery charge more quickly than ordinary e-learning was not performed though. Face-api.js promising to be designed for mobile devices as well should give this plugin enough flexibility to run on any device. How big the effect of this is is not quantified, however.

During testing, regarding scalability, over 700 new event logs increased the space taken in the database by around 0.1 MiB. This should mean no significant load is coming from the plugin even with frequent logs appearing. Naturally, Moodle already logs every user (logged in or not) when accessing many pages or doing specific actions on them, meaning frequent new entries from a single person in a short time are to be expected already.

However, regarding database size, the Moodle documentation[1] mentions: "Accessing logs on large Moodle instances will be slow and will cause load on the entire system". Should this turn out to be causing issues on a large install, a version of the plugin that uses rarer reports would be desirable. This would be comparable to the load of a user changing from one activity to another regularly if the frequency would be reduced to logs only occurring once every few minutes.

A small downside is potential loss in accuracy as more data is summed up in one log. Also, more data could get lost if any one log happened to fail being received, perhaps due to connectivity issues. In the case of performance issues, altering the frequency may likely be worth it as long as individual emotions can still be discerned from the logs, for example by changing the format to highlight which emotions were detected as individual instance over a certain threshold in this time frame and how often this occurred, instead of using averages over a few seconds that are tied to reports each.

Moving some of the information to a separate database table from the logs may help, though logging emotion events when important changes occur is still desirable for inter-plugin-communication. It is not needed if an ALS is implemented as part of this same plugin, however.

6.3 Improvements to Software Concept

To implement an e-learning system on Moodle with Facial Emotion Recognition (FER), the current plugin can be expanded and/or separate plugins can read the data from this one. To make the process more smooth and fix remaining issues, this is a list of considerations:

- The import of face-api.js in a Moodle AMD JavaScript file requires either a hard coded path adapted to the server file structure, of to adapt face-api's distributed library to work with Moodle's naming scheme.[2]
- A fix for the plugin's Privacy API implementation to avoid manual interventions being necessary.
- Emotion data should be stored in a database table that can be linked in logs, instead of integrating it in the event logs themselves. Logs are meant to allow other plugins access to the emotion data. This means there will be less need to

[1] https://docs.moodle.org/dev/Migrating_log_access_in_reports

[2] Related Moodle Documentation: https://docs.moodle.org/dev/Javascript_Modules#Including_an_external_javascript.2Fjquery_library

adapt to restrictions from using the logging system such as lack of floating point number support in the relevant data field.

- The current system uses particularly common and established emotions as basis, which seem to work well overall to guess positive or negative emotion tendencies. But neither this nor any individual emotion of face-api.js is fully reliable. Especially sadness is over-estimated by face-api in many cases so that it appears too often.
- Other emotions that can be relevant to e-learning can be implemented for detection. This could be done additionally to what already exists, or instead of the face-api implementation. [DSGG20] could be used to attempt more directly useful attentiveness analysis, which is also suggested in [Al-16] for a strategy to intervene whenever a student is at danger of becoming inattentive, which may be detected with an attentiveness model or through other factors (including emotions). To detect whether a student is frustrated, confused, bored or engaged the DAiSEE data set also contains ways to create a fitting model [GDAB16].

7 Conclusion and Future Work

7.1 Summary and Contributions

Moodle is an open-source system that allows much freedom and expansion. A plugin was created in it that allows placing blocks on course pages and activities that use Facial Emotion Recognition FER to store emotion data of Moodle users and allow reacting to it. While the size and complexity of Moodle makes some concepts hard to immediately fully understand, Moodle worked well for developing this plugin. All central points for thesis were implemented, though some not in a functional way: the Privacy API for automating issues like data information and deletion requests is not usable yet, and manually deleting data can be annoying for systems that run on a big scale with many users.

To understand how an ALS may work, some own use cases were designed and UI for teachers to get feedback directly so far not designed, but outlined and functionality exists with established Moodle tools, though currently data are not compiled or interpreted in a somewhat usability friendly way. This could be used as a way of feedback for teachers, but it not required for an ALS.

This thesis also compiled some relevant literature for FER in e-learning specifically. For this, no full literature review was conducted. But existing recent survey papers were used to filter the most relevant information while other non-included literature was searched for and used, too. Requirements, implications and suggestions for privacy were discussed. Steps were outlined for refining this plugin for actual use. The overall concept was verified via a survey.

7.2 Future Work

This plugin can be expanded to either offer emotion data to a Moodle site for other modules to use it, or to become the driver of an ALS itself. Better emotion detection implementations with more and more relevant emotions can be implemented. Creating a better understanding of how adaptive systems can react to emotions in meaningful and evidence based ways that increase user friendliness and learning ability in e-learning can be a good next step.

A Appendix

A.1 Requirements and User Stories

A list of all user stories and software requirements made to describe the scope and functionality of the software product to develop.

A.1.1 User Stories

1. As a teacher, I want to know if a student needs personal attention to not fall far behind so that I can help them.
2. As a teacher, I want to get feedback on what works well in my teaching methods and on which audience, so I can adapt my methods if required.
3. As a teacher, I want the emotion tools to be easy to use so that I don't waste time whenever I interact with them and don't have to concentrate on figuring the system out.
4. As a student, I want to decide whether the camera is active or not, so that I have a feeling of control over my privacy.
5. As a student, I want to get done with the task quickly without the process draining my device's battery so that I can study without depending on a power supply nearby.
6. As a student, I want the study process to be straightforward so that I don't have to worry about any plug-ins getting in the way of what I am concentrating on.
7. As a student, I want to maximize my focus while learning on this system and be efficient so that I can learn better and spend my time on other things then.
8. As a student, I want to know which parts of a course I personally may need to refresh my knowledge on so that I can avoid forgetting about a concept.
9. As a student, I want frustrating parts of lecture videos to be easier to understand, so that I can get through them quicker and understand it.

A.1.2 Functional Requirements

1. Plug-In Accesses real-time webcam input for emotion recognition
2. Plug-In should ask for necessary agreement before use, possibly once or each time it gets activated, without being annoying. This may be done via existing interfaces in browsers to ask for permission to access a webcam. See also Non-Functional Requirement on Privacy / GDPR
3. Plug-In works in Moodle given real-time video feed.
   a) Must be possible when studying on the website
   b) optional: Usable during any task when seeing as many different Moodle web pages as possible with aim to work during any interaction with the platform.
4. Don't make webcam footage itself accessible to anyone
5. GUI for starting/stopping webcam feed as user
   a) To give users direct control of when the software "sees" them (more privacy, more acceptance)
   b) This feature is to be togglable in plug-in settings.
   c) optional: Teachers may be able to set certain cases (e.g., exercises, videos, URLs) as places where the plug-in will not be allowed to operate. Also specify when it is required to be running, if at all.
6. Recognize emotions from face
   a) at least to extent that's possible and feasible through available free Open Source libraries.
   b) or possibly using more advanced licensed software if necessary
7. optional: Recognize emotions from other factors visible on camera (motion, gestures, posture)
8. Plug-In must make its info available to be usable by the Moodle system
   a) Via GUI: Allow teachers seeing emotion statistics for tasks.
   b) Via GUI: Optional: Allow teachers emotion statistics per person. This lets a teacher personalize and adapt any course by themselves, and therefore do their job well informed.
   c) Via API for Adaptive Learning Assistant (ALS)

- access information about what emotion occured during what exercise, at which time.
- read about ALAs to find out what else could make sense, get exact details later.
- This should be open to other plugins as well by offering this data in a standard way on Moodle via logs.

Note: Creating an emotion-based adaptive learning assistant to teach Artificial Intelligence is not part of the project. Only emotion recognition and an API are implemented.

A.1.3 Non-Functional Requirements

1. Works on all normal end user devices (e.g. computer, phone) and at least tested on the browsers Chrome, Edge, Firefox
2. Works on current Moodle version 3.11 and up
3. Performance: Must be able to actually run on a Moodle server
4. Performance: Handle multiple users on server at same time. No concrete number is defined here, this partially depends on server capacity too. The system must be designed with scaling it up to many users in mind.
5. Privacy / Security: keep video feed and all other data secure in connection (also see functional requirements above: video not visible for anyone)
6. Privacy / Security: System is GDPR compliant let users see / delete their data look into any further legal requirements
7. Localisation: In English and German. This way, it can be used by study participants in German, but has language consistent with the publication as well. Moodle has features for localisation support that any new plugin is recommended to use.
8. Usability: Has to be usable for teachers and students with sufficient documentation
9. Maintainability / Expandability: Code is documented and has comments.
10. Maintainability / Expandability: Code generally allows future updates and expansion.
11. Persistent: Data is kept even if system shuts down unexpectedly, given that database is intact.

A.2 German Consent Form

This is the consent information and study process contract that was used as part of the software evaluation cycle. The study was conducted in German, so this text is in German as well.

A.2.1 Ziele und Ablauf der Studie

Ziel der Studie ist es, Feedback zur Nutzungserfahrung von einer Emotionserkennungs-Software zu erhalten, die auf dieser Webseite eingerichtet wurde. Dazu gehört: Wie gut lässt sich die Emotionserkennung von Ihnen bedienen? Was kann daran noch verbessert werden?

Es werden einige Vorlesungsvideos angezeigt sowie Fragen und Aufgaben gestellt, um Situationen aus Online-Vorlesungen zu simulieren. Durch Bilder Ihrer Webcam wird dabei regelmäßig versucht, Emotionen aus Ihrem Gesicht abzulesen. An zwei Stellen werden Sie zum Vergleich gefragt, wie Sie sich bei der Bearbeitung gefühlt haben. Schließlich werden Sie um gezieltes Feedback gebeten.

A.2.2 Vertrag zwischen Versuchsleiter und Versuchsteilnehmer/in

(This text is based on [HSE13, p. 46] / Text in Anlehnung an Hussy et al., 2013, S. 46.)

Liebe/r Untersuchungsteilnehmer/in, es folgt eine Information über Ihre Rechte und Pflichten als Versuchsperson sowie über die Verpflichtungen des Versuchsleiters:

1. Freiwilligkeit der Teilnahme Sie haben sich freiwillig entschieden, an dieser Umfrage teilzunehmen. Wir sind sehr dankbar für Ihre Teilnahme und dass Sie uns helfen, diese Studie erfolgreich durchzuführen.

2. Schutz vor Überforderung und Belastungen Wir versichern Ihnen, dass wir jegliche psychische oder körperliche Belastung vermeiden wollen und dass Ihr Wohlbefinden wichtiger ist als unser Interesse an aussagekräftigen Ergebnissen.

3. Abbruchrecht Sollten während der Umfrage Belastungen auftreten, die Sie als schwerwiegend empfinden, können Sie jederzeit aussteigen. Der Ausstieg hat keine negativen Folgen. Erfasste Daten können auf Wunsch gelöscht werden.

4. Recht auf postexperimentelle Aufklärung Wir versichern, dass nach Abschluss der Erhebungsphase der einzelnen Experimente auf Nachfrage sämtliche gewünschten Informationen über Ablauf, Zweck und Ergebnis des Versuchs gegeben werden können.

Kontaktieren Sie hierzu Stephan Schroth (e-mail address redacted from this released version of the text)

5. Pflichten als Studienteilnehmer Mit diesen Rechten korrespondieren auch Pflichten für die Teilnehmer. Die Planung und Durchführung einer Studie erfordert viel Zeit und Mühe. Es ist daher wichtig, dass Sie versuchen, die gestellten Fragen so gut wie möglich zu beantworten. Dazu gehört, dass Sie die Fragen offen und ehrlich beantworten. Eine uninteressierte und oberflächliche Teilnahme gefährdet die Ziele der Studie.

6. Gewährleistung der Anonymität Wir möchten betonen, dass Ihre Antworten streng vertraulich behandelt werden.

7. Datenschutz Ihre im Rahmen dieser digital durchgeführten Umfrage erhobenen Daten werden in privaten Computern gespeichert, verarbeitet und in schriftlicher anonymisierter Form zu wissenschaftlichen Zwecken veröffentlicht. Ihre Daten werden bis zur Beendigung der Arbeit an der Studie gespeichert und anschließend wieder vollständig gelöscht.

Die Auswertung der Kamerabilder geschieht im verwendeten Browser, und damit ausschließlich lokal auf Ihrem Endgerät. Kein Mensch kann diese Bilder Ihrer Webcam jetzt oder später sehen und sie werden nicht übertragen oder gespeichert. Wann die Erkennung via Webcam aktiv ist, haben Sie selbst durch die Steuerungselemente jederzeit unter Kontrolle.

Die erkannten Emotionen und von ihnen in der Studie gegebenen Antworten werden über eine verschlüsselte Verbindung auf den Server der Website übertragen und sind dort für die Betreiber der Seite einsehbar. Moodle speichert dabei auch Ihre IP-Adresse, die jedoch aufgrund der VPN-Nutzung bereits anonymisiert ist. Niemand mit Zugriff auf die Daten weiß, welche Person hinter der Adresse oder dem entsprechenden Moodle-Account steht.

8. Unterbrechung / Fortschritt

Es ist jederzeit möglich die Studie zu pausieren und später weiterzuführen.

Mit dem Setzen des Hakens zum Akzeptieren unter dieser Erklörung und der anschließenden Teilnahme an dieser Umfrage erklären Sie, dass Sie den Vertrag gelesen, Ihre Rechte und Pflichten zur Kenntnis genommen haben und mit der Datenschutzerklärung einverstanden sind. Weiterhin erklären Sie, dass Sie uber die Ziele und den Ablauf der Studie informiert wurden und diese verstanden haben.

A.3 More Emotion Data

This section contains emotion data detected from Participants 1, 3 and 4, as well as the data they submitted themselves. Data for participant 2 is found in the main work in figure 5.3

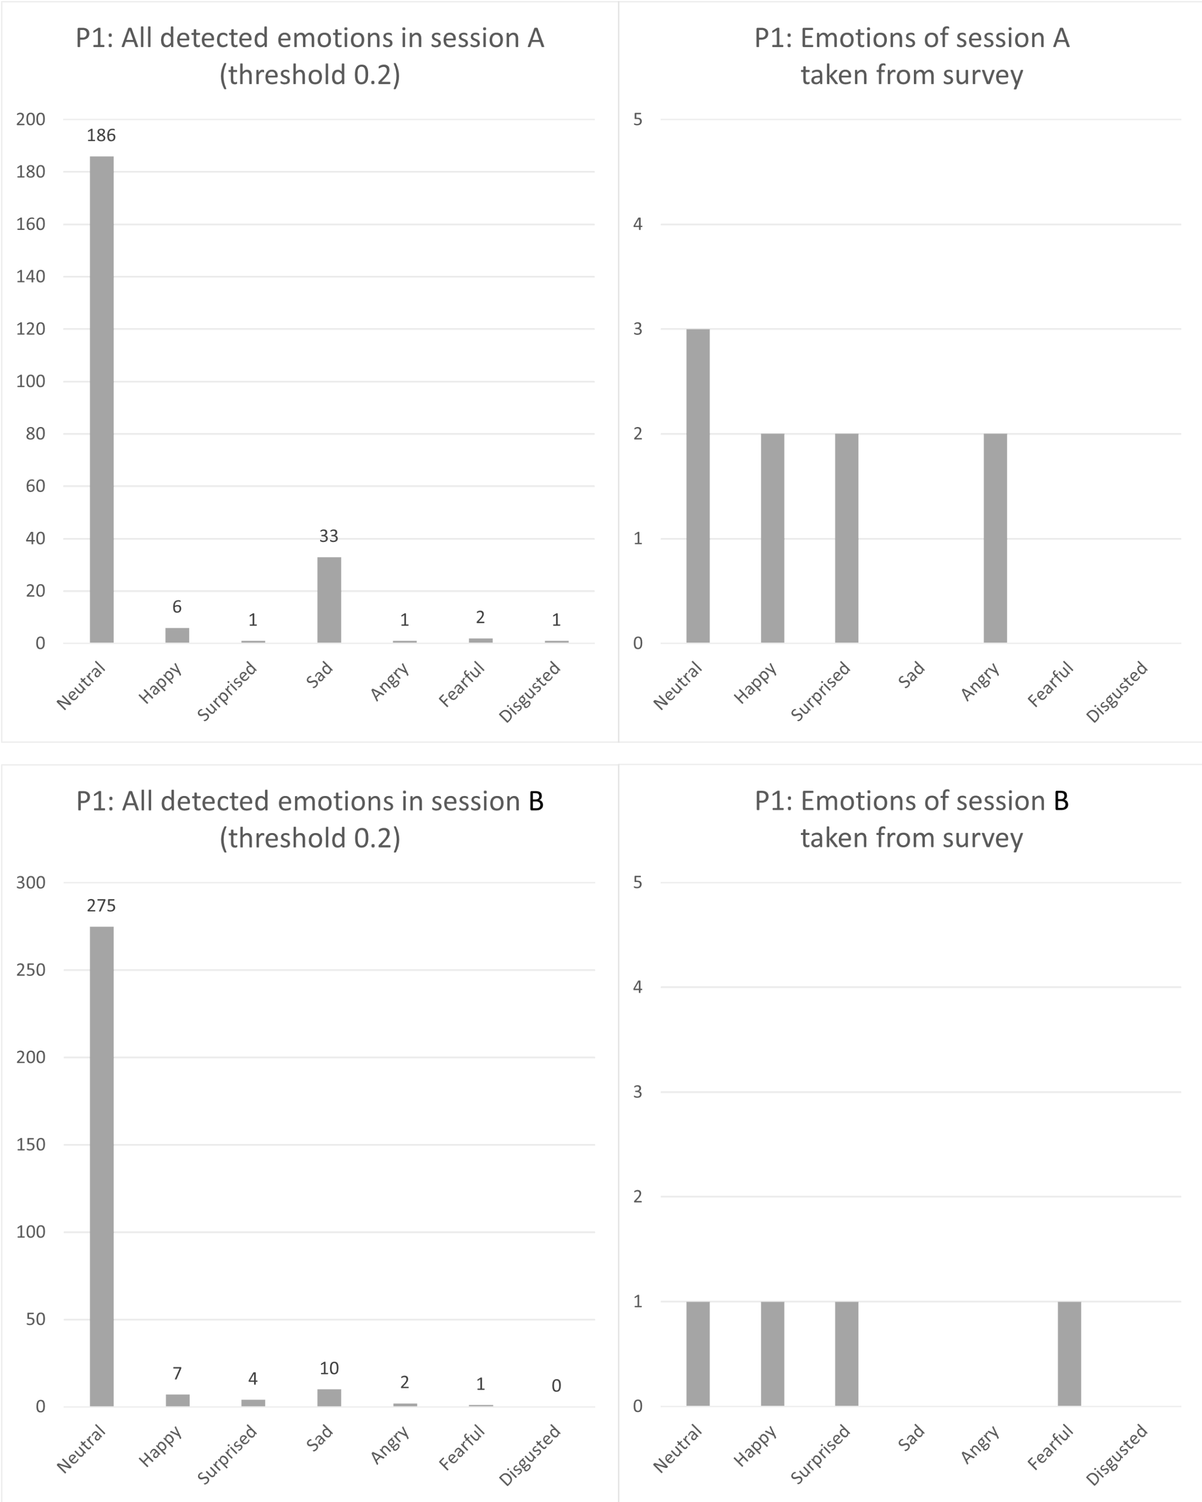

Figure A.1: Emotions detected versus emotions given directly in survey for participant 1.

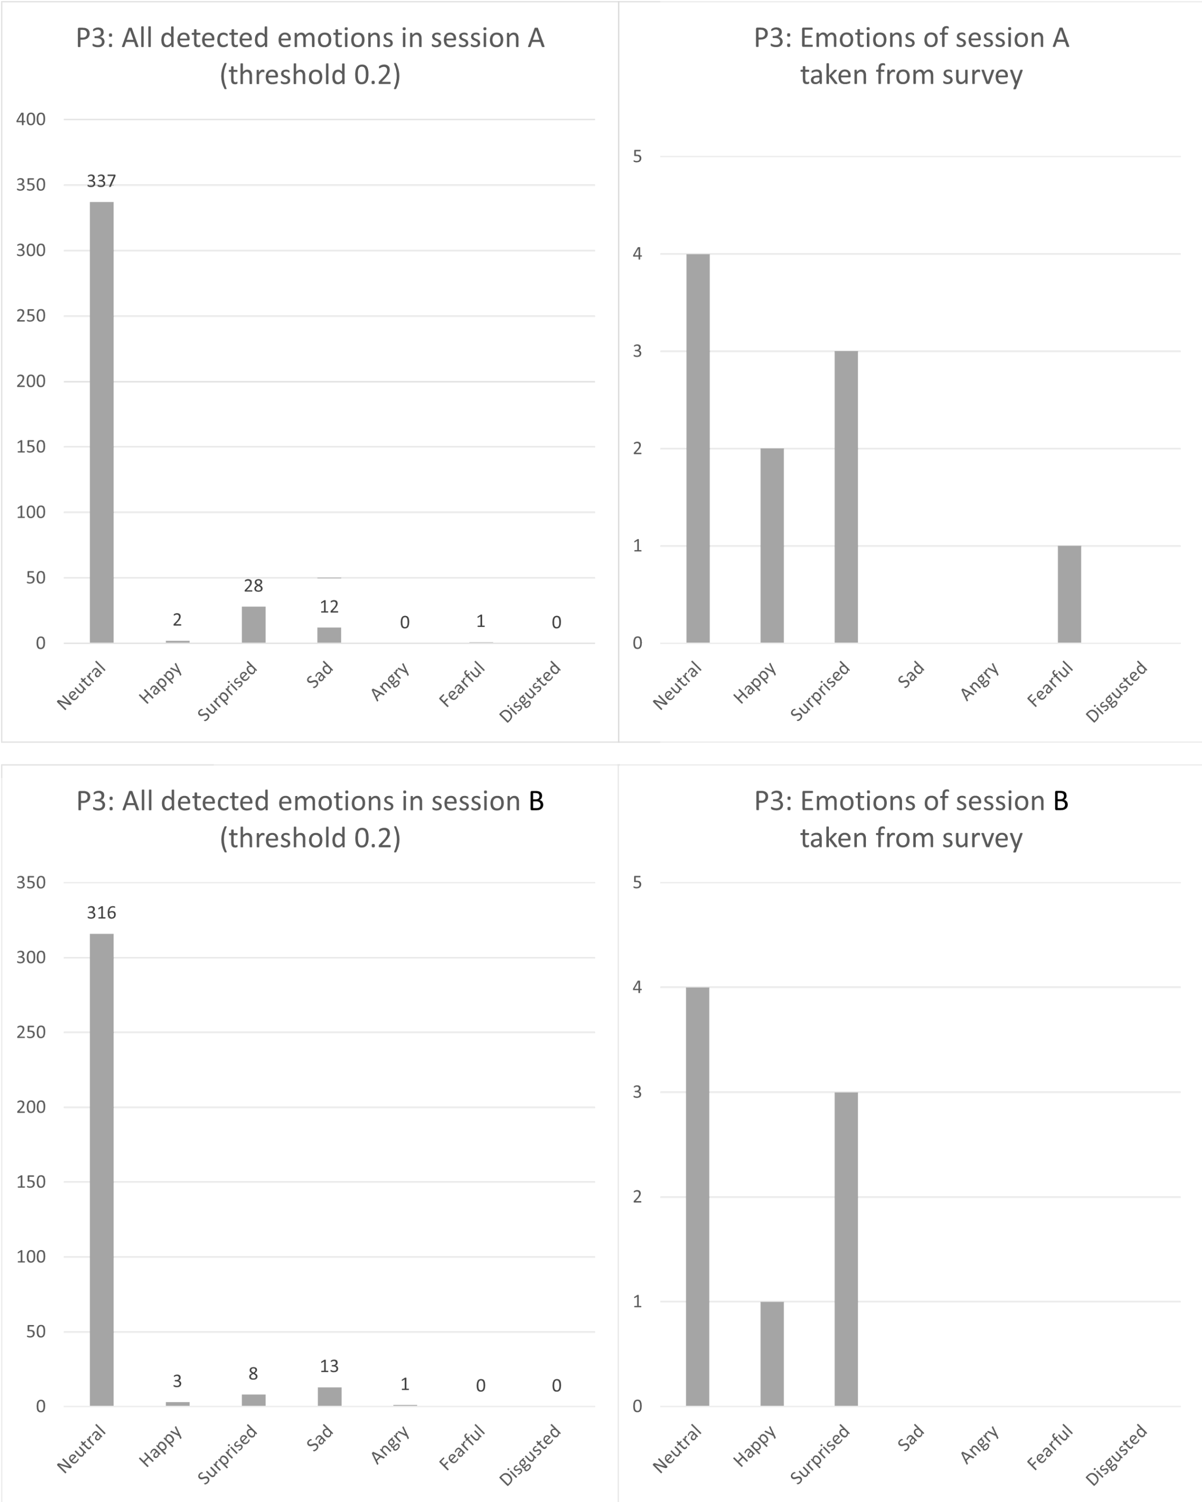

Figure A.2: Emotions detected versus emotions given directly in survey for participant 3.

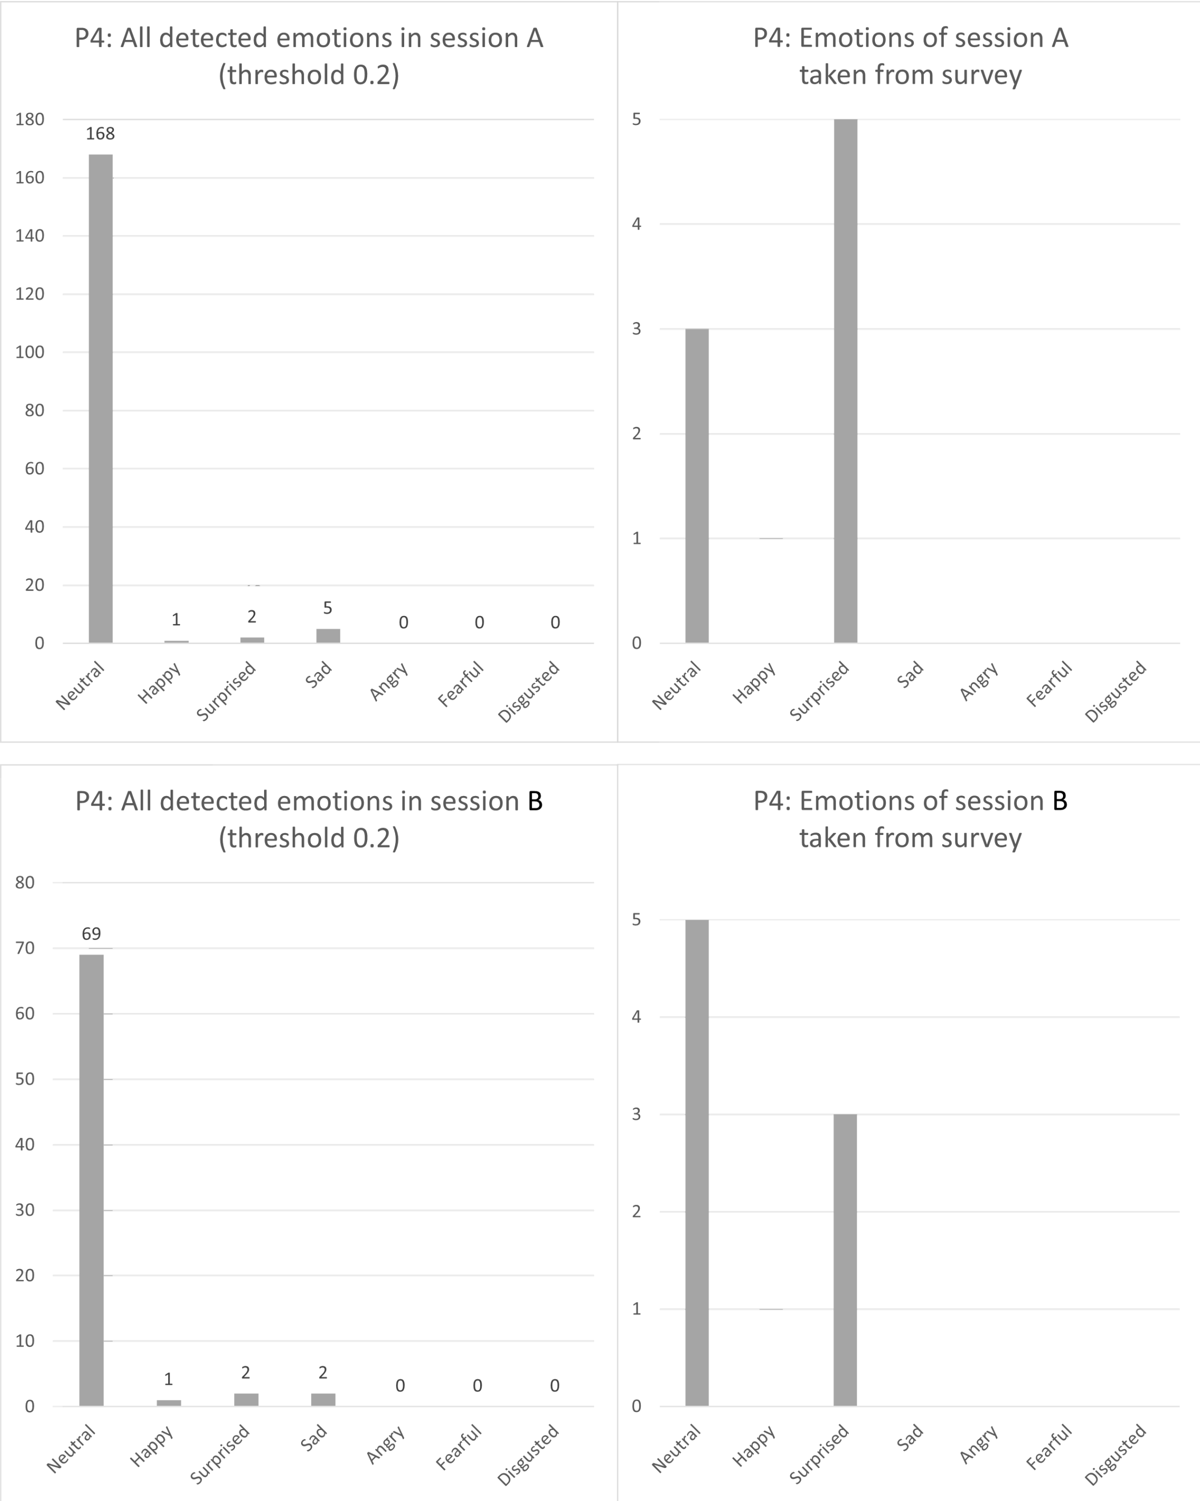

Figure A.3: Emotions detected versus emotions given directly in survey for participant 4.

Bibliography

[Al-16] A. Al-Awni. “Mood extraction using facial features to improve learning curves of students in e-learning systems.” In: *International journal of advanced computer science and applications* 7.11 (2016), pp. 444–453. DOI: https://dx.doi.org/10.14569/IJACSA.2016.071157 (cit. on pp. 21, 40, 52).

[Ass+13] W. M. Association et al. “World Medical Association Declaration of Helsinki: ethical principles for medical research involving human subjects.” In: *Jama* 310.20 (2013), pp. 2191–2194. DOI: 10.1001/jama.2013.281053 (cit. on p. 22).

[AZ05] N. A. Alias, A. M. Zainuddin. “Innovation for better teaching and learning: Adopting the learning management system.” In: *Malaysian online journal of instructional technology* 2.2 (2005), pp. 27–40 (cit. on p. 19).

[DSGG20] A. Dhall, G. Sharma, R. Goecke, T. Gedeon. “EmotiW 2020: Driver Gaze, Group Emotion, Student Engagement and Physiological Signal Based Challenges.” In: *Proceedings of the 2020 International Conference on Multimodal Interaction*. New York, NY, USA: Association for Computing Machinery, 2020, pp. 784–789. ISBN: 9781450375818. DOI: https://doi.org/10.1145/3382507.3417973 (cit. on p. 52).

[EF71] P. Ekman, W. V. Friesen. “Constants across cultures in the face and emotion.” In: *Journal of personality and social psychology* 17.2 (1971), p. 124. DOI: https://doi.org/10.1037/h0030377 (cit. on p. 17).

[Ekm92] P. Ekman. “An argument for basic emotions.” In: *Cognition and Emotion* 6.3-4 (1992), pp. 169–200. DOI: 10.1080/02699939208411068. eprint: https://doi.org/10.1080/02699939208411068 (cit. on p. 17).

[FAM+17] A. R. Faria, A. Almeida, C. Martins, R. Gonçalves, J. Martins, F. Branco. “A global perspective on an emotional learning model proposal.” In: *Telematics and Informatics* 34.6 (2017). SI: IT Education Training, pp. 824–837. ISSN: 0736-5853. DOI: https://doi.org/10.1016/j.tele.2016.08.007. URL: https://www.sciencedirect.com/science/article/pii/S0736585316302040 (cit. on p. 21).

[GDAB16] A. Gupta, A. D'Cunha, K. Awasthi, V. Balasubramanian. *DAiSEE: Towards User Engagement Recognition in the Wild*. 2016. DOI: 10.48550/ARXIV.1609.01885. URL: https://arxiv.org/abs/1609.01885 (cit. on p. 52).

[HSE13] W. Hussy, M. Schreier, G. Echterhoff. *Forschungsmethoden in Psychologie und Sozialwissenschaften für Bachelor*. German. Springer-Verlag, 2013. ISBN: 978-3-642-34362-9. DOI: 10.1007/978-3-642-34362-9 (cit. on pp. 41, 58).

[IM19] M. Imani, G. A. Montazer. "A survey of emotion recognition methods with emphasis on E-Learning environments." In: *Journal of Network and Computer Applications* 147 (2019), p. 102423. DOI: 10.1016/j.jnca.2019.102423 (cit. on pp. 17, 18, 20, 21, 25).

[KK81] P. R. Kleinginna, A. M. Kleinginna. "A categorized list of emotion definitions, with suggestions for a consensual definition." In: *Motivation and emotion* 5.4 (1981), pp. 345–379. DOI: 10.1007/BF00993889 (cit. on p. 17).

[KKMK18] A. H. A. Khan, S. Kulkarni, T. Mahmood, A. A. Khan. "Evidence-based approaches to learning." In: *Advances in Medical Education and Practice* 9 (2018), pp. 581–582. DOI: https://doi.org/10.2147/AMEP.S171499 (cit. on p. 21).

[May20] R. E. Mayer. "Searching for the role of emotions in e-learning." In: *Learning and Instruction* 70 (2020). Understanding and measuring emotions in technology-rich learning environments, p. 101213. ISSN: 0959-4752. DOI: https://doi.org/10.1016/j.learninstruc.2019.05.010. URL: https://www.sciencedirect.com/science/article/pii/S095947521930324X (cit. on p. 20).

[MDG11] J. L. Moore, C. Dickson-Deane, K. Galyen. "e-Learning, online learning, and distance learning environments: Are they the same?" In: *The Internet and Higher Education* 14.2 (2011). Web mining and higher education: Introduction to the special issue, pp. 129–135. ISSN: 1096-7516. DOI: https://doi.org/10.1016/j.iheduc.2010.10.001. URL: https://www.sciencedirect.com/science/article/pii/S1096751610000886 (cit. on p. 18).

[MVTS21] N. Morze, L. Varchenko-Trotsenko, T. Terletska, E. Smyrnova-Trybulska. "Implementation of adaptive learning at higher education institutions by means of Moodle LMS." In: *Journal of Physics: Conference Series*. Vol. 1840. 1. IOP Publishing. 2021, p. 012062. DOI: 10.1088/1742-6596/1840/1/012062 (cit. on p. 19).

[NPB18] A.-M. Nortvig, A. K. Petersen, S. H. Balle. "A Literature Review of the Factors Influencing E-Learning and Blended Learning in Relation to Learning Outcome, Student Satisfaction and Engagement." In: *Electronic Journal of E-learning* 16.1 (2018), pp46–55 (cit. on p. 18).

[PMRB08] H. Pashler, M. McDaniel, D. Rohrer, R. Bjork. "Learning Styles: Concepts and Evidence." In: *Psychological Science in the Public Interest* 9.3 (2008). PMID: 26162104, pp. 105–119. DOI: 10.1111/j.1539-6053.2009.01038.x. eprint: https://doi.org/10.1111/j.1539-6053.2009.01038.x (cit. on p. 21).

[RB14] A. Rahimi, R. A. Bigdeli. "The Broaden-and-Build Theory of Positive Emotions in Second Language Learning." In: *Procedia - Social and Behavioral Sciences* 159 (2014). 5th World Conference on Psychology, Counseling and Guidance, WCPCG-2014, 1-3 May 2014, Dubrovnik, Croatia, pp. 795–801. ISSN: 1877-0428. DOI: https://doi.org/10.1016/j.sbspro.2014.12.451. URL: https://www.sciencedirect.com/science/article/pii/S1877042814065811 (cit. on p. 21).

[SVC12] A. Sangrà, D. Vlachopoulos, N. Cabrera. "Building an inclusive definition of e-learning: An approach to the conceptual framework." In: *International Review of Research in Open and Distributed Learning* 13.2 (2012), pp. 145–159. DOI: 10.19173/irrodl.v13i2.1161 (cit. on p. 18).

[YK14] J. W. You, M. Kang. "The role of academic emotions in the relationship between perceived academic control and self-regulated learning in online learning." In: *Computers Education* 77 (2014), pp. 125–133. ISSN: 0360-1315. DOI: https://doi.org/10.1016/j.compedu.2014.04.018. URL: https://www.sciencedirect.com/science/article/pii/S0360131514001031 (cit. on p. 20).

Additionally, references to articles of the official Moodle documentation (https://docs.moodle.org) or official websites of software projects are given as footnotes directly in the text. The following other web resources are used as references:

- p. 19: Moodle Website: https://moodle.org/
- p. 21: Amazon Rekognition Documentation: https://docs.aws.amazon.com/rekognition/latest/dg/faces-detectimages.html
- p. 23: Agile Alliance glossary: User Stories https://www.agilealliance.org/glossary/user-story-template/
- p. 24: Agile Alliance brief presentation on 'stakeholder': https://www.agilealliance.org/wp-content/uploads/2016/01/Stakeholder-Management-Agile2012-Draft-1.pdf
- p. 31: MDN: Accessibility in HTML development: https://developer.mozilla.org/en-US/docs/Learn/Accessibility/HTML
- p. 35: !! Find out name: Face-API Git repository: https://github.com/justadudewhohacks/face-api.js

- p. 38: Mozilla Blog: Secure Context, https://blog.mozilla.org/webrtc/camera-microphone-require-https-in-firefox-68/
- p. 38: Chromium.org: Secure Context, https://www.chromium.org/Home/chromium-security/deprecating-powerful-features-on-insecure-origins/
- p. 40: nativdigital.com: Feedback Grid, https://nativdigital.com/feedback-grid/

All links were last followed on May 16, 2022.

Declaration

I hereby declare that the work presented in this thesis is entirely my own and that I did not use any other sources and references than the listed ones. I have marked all direct or indirect statements from other sources contained therein as quotations. Neither this work nor significant parts of it were part of another examination procedure. I have not published this work in whole or in part before. The electronic copy is consistent with all submitted copies.

---

place, date, signature