

UC Davis

UC Davis Previously Published Works

Title

#let's-discuss: Analyzing Students' Use of Emoji when Interacting with Course Readings

Permalink

<https://escholarship.org/uc/item/3sw6q0h4>

Authors

Blobstein, Ariel

Facciotti, Marc T

Igo, Michele

et al.

Publication Date

2024

DOI

10.1007/s40593-024-00401-5

Peer reviewed



#let's-discuss: Analyzing Students' Use of Emoji when Interacting with Course Readings

Ariel Blobstein¹ · Marc T. Facciotti^{3,4} · Michele Igo⁵ · David Karger⁶ · Prerna Ravi⁶ · Kamali Sripathi⁴ · Kobi Gal^{1,2}

Accepted: 11 March 2024
© The Author(s) 2024

Abstract

Emoji are commonly used in social media to convey affects, emotions, and attitudes. While popular in social media, their use in educational contexts has been sparsely studied even though emoji can be a natural way for students to express what they are feeling about the learning material. This paper studies how students use instructor-selected emoji when relating to and engaging with educational content. We use an online platform for collaborative annotations where discussions are embedded into the readings for the course. We also make it possible for students to use 11 unique emoji-hashtag pairings to express their thoughts and feelings about the readings and the ongoing discussion. We provide an empirical analysis of the usage of these emoji-hashtag pairs by over 1,800 students enrolled in different offerings of an introductory biology course from multiple academic terms. We also introduce a heat map, which allows the instructional team to visualize the distribution and types of emoji used by students in different parts of the reading material. To evaluate the heat map, we conducted a user study with five instructors/TAs. We found that instructors/TAs use the heat map as a tool for identifying textbook sections that students find difficult and/or interesting and plan to use it to help them design the online content for future classes. Finally, we introduce a computational analysis for predicting emoji/hashtag pairs based on the content of a given student post. We use pre-trained deep learning language models (BERT) to predict the emoji attached to a student's post and then study the extent to which this model generated in an introductory biology course can be generalized to predict student emoji usage in other courses.

Keywords Course forums · Affect recognition in course readings · Natural language processing · Supervised learning

✉ Ariel Blobstein
arielblo@post.bgu.ac.il

Extended author information available on the last page of the article

Introduction

Students' interactions with learning materials can elicit a broad range of reactions that may include emotional, cognitive, affective, and behavioral elements. Collectively, these reactions can contribute positively or negatively to students' learning (Boekaerts & Pekrun, 2015). Identifying and mapping student responses onto specific course content can help instructors to better understand their students and empower teaching staff to optimize course design and materials for better learning outcomes (Sylwester, 1994; Kort et al., 2001; Baker et al., 2010; D'Mello & Graesser, 2012). Unfortunately, these tasks can be logistically challenging for instructors, particularly in large-enrollment courses.

Recent studies have shown ways to address these logistical challenges by using machine learning to predict students' emotions, affects, and attitudes (Kastrati et al., 2021). Prior studies created labeled data sets from discussion forums in a MOOC platform (Agrawal et al., 2015; Yang et al., 2015). The labels were collected by manually assigning reactions to student posts using crowd workers hired from platforms like Mechanical Turk. Confusion, emotion, and urgency are some of the annotated categories studied in prior work. However, paid crowd workers are far removed from the course material and can fail to identify student reactions that may be useful to instructors. A better solution to the labeling problem would be to let students annotate their reactions themselves and to provide ways for them to distinguish between issues related to specific course content versus more general feelings about the course.

The premise of this paper is that students can use emoji to communicate their reactions while commenting on the course material, thus creating a useful tool for students to report their reactions to specific parts of course content. Since emoji are conventionally used in social media to enhance the meaning of the text or as a substitute for words (Kerslake & Wegerif, 2017), emoji can be a natural way for students to express how they are reacting to the learning materials, which would otherwise be difficult for instructors to infer. A body of work in the social and computational sciences has studied people's use of emoji as a tool for conveying reactions and used them to facilitate computational tasks such as sentiment analysis and emotion recognition (Kralj Novak et al., 2015; Felbo et al., 2017). However, there is sparse research on the use of emoji for self-reporting reactions by students in educational forums. Our work directly addresses this gap by providing an empirical and computational analysis of students' use of emoji when they discuss specific sections of course material from an online textbook.

Our empirical investigations involve NotaBene (NB) (Zyto et al., 2012), a collaborative annotation-based forum, which allows students to anchor messages directly to reading material in online course textbooks. We worked closely with the course instructors of a large enrollment introductory biology course offered to ~3500 students per year at a large public research university in the western United States. Course instructors designed a set of 11 emoji-hashtag pairs (Shown in Fig. 1) that allow students to communicate reactions to the course textbook during their discussions on NB. Example of emoji-hashtag are to express curiosity, confusion, or interest about topics in the course, or inviting discussion about a topic, or flagging a passage as pertinent to a learning goal or real-world problem. Limiting students to using only these 11

emoji-hashtag pairs rather than allowing them to insert their hashtags was important for helping us standardize what students meant to convey when they selected to use a specific emoji-hashtag pair. Considering the diversity of response types represented by our emoji-hashtag pairs, we use the term **reactions** throughout this paper to inclusively refer to emotions, affects, attitudes, epistemic states, and/or requests for interaction.

Here, we provide an empirical analysis of students' use of emoji in two offerings of the course. This includes measuring the frequency of use of each individual emoji and the number of different emoji used in the same post. We found that the most common use of emoji was to request discussion and assistance from instructors or peers. We further investigated the use of emoji among different students and found that most students tended to have a favorite sub-set of the emoji set that they used in their posts rather than the full emoji set. We also analyzed the correlation between emoji frequency as used by the students. We find that specific pairs of emoji frequently appear together in posts relating to the same paragraph or portion of the text.

Our second step was to provide instructors with a "birds eye view" of students' emoji use in the course. To this end, we integrated a heat map (shown in Fig. 6), the heat map can provide a visual representation of either a specific emoji or an entire set of emoji, depicting the frequency of their usage by students. The heat map also shows how emoji are utilized by students about course content, emphasizing their engagement with posts related to that content.

We conducted a user study with five experienced instructors to investigate what types of information instructors can glean from the heat map, and how they might use this information in their instruction.

The results of the user study revealed that the heat map helped instructors identify areas of interest or difficulty in the course material that would have been much harder for instructors to discover by other means. Participants stated that without the heat map, it would have been more time-consuming and difficult to identify sections of the reading material that students found difficult or interesting, without the help of the tool and thus letting them refocus their lectures or revise content to enhance students learning. Our work can help instructors, particularly in high-enrollment courses, decide where and how to intervene in discussions.

The work we describe shows that emoji offer useful insights to instructors. However, students don't always include emoji in their comments. Therefore, our third step was to explore how information generated from the use of emoji by other students could help develop a supervised learning model for *inferring* specific emoji for a post without emoji. A predictive model that can infer key components of a student's reactions to course content could be used to better represent course-wide reactions to content and potentially be used for courses that incorporate student posts but don't provide an easy way for students to indicate their reactions to the material.

An important contribution of this work is the way we leveraged the *annotated content* to improve inference accuracy. We use a pre-trained deep-learning language model (BERT (Devlin et al., 2018)) to predict emoji from students' posts *and the course material that the post addressed*. Since students using NB link their annotations directly to a specific part of the textbook, we can use the context derived from the reading material to further enhance the model. Thus the resulting model takes into consideration both students' posts and the context of the reading. The model was

trained on students' posts from a biology course where the predicted target is an emoji. This model significantly outperforms the state of the art from the literature for predicting emoji in social media (Felbo et al., 2017; Baziotis et al., 2018; Blobstein et al., 2022).

We extend a previous conference publication (Blobstein et al., 2022) by providing new insights and improvements on the results presented therein. Specifically, we extended the initial model and improve its performance compared to the previously reported version. Moreover, we show that the new model is able to generalize to courses outside of the biology domain. Furthermore, we introduce a novel emoji heat map and evaluate its benefit to instructors in a detailed user study.

Related Work

Our work relates to past studies on collaborative platforms and learning analytics, inferring student reactions from their online text conversations, as well as computational work that analyzes the use of emoji in social media. We mention relevant prior works below and refer the reader to the review of the field by Kastrati et al. (2021) for additional details.

Collaborative platforms and Learning analytics

Over the last decade, online learning has become increasingly popular. The transition into this new educational space has been driven by various socioeconomic forces and enabled by numerous evolving platforms for online content delivery and collaborative learning. The increased use of this new medium for learning has highlighted the need for tools that can monitor students' participation and assess the skills and knowledge they have developed while online (Kuosa et al., 2016). It has also pushed the development of so-called learning analytics for online platforms (Kuosa et al., 2016; Sahin & Ifenthaler, 2021). To date, most research that utilizes data from learning analytics tends to focus on data that can be gleaned from user interaction logs, assignment, and final grades, and sometimes performance on various embedded formative assessments. While this type of research is critical for understanding how different platform features and user behaviors contribute to learning, these studies often ignore other key factors that can contribute to student engagement and learning, such as how students feel about and how they react to the online content. Although the relationship between emotions, student well-being, and self-regulated learning is well documented (Bennett & Folley, 2021; Boekaerts et al., 1999; Boekaerts, 1988), most prior work in this area relies on self-reported data from the students from student interviews or expert annotation of student work. The latter two mechanisms are difficult to scale and subject to misinterpretation. Thus, there is a need for mechanisms that enable students to accurately self-report as an additional data source on self-regulated learning.

Our work relates to other works using collaborative annotations and highlights to infer students' attitudes to learning materials (Sun et al., 2023; Sun & Gao, 2017; Farzan & Brusilovsky, 2008; Winchell et al., 2020). **Our use of emoji provides a broader range of information regards students' reactions that are useful to**

instructors, as we show in the user study (User Study). The focus of this study is to investigate how this type of platform could benefit instructors by providing them with efficient tools for examining their students' reactions to course material, both intellectually and emotionally. According to Munday (2021) a fundamental shift "in feedback methods which utilize the technological affordances beyond simply electronic text and upload" is a crucial requirement these days.

Inferring Student Reactions from Online Conversations

Basic approaches to infer students' sentiments in online conversations used classic NLP methods (parsing, lexical dictionaries) (Chaplot et al., 2015; Binali et al., 2009). Jena (2019) used classical machine learning methods (SVMs, Naive Bayes) to learn the sentiment polarity (positive, negative, and neutral) from students' posts, as well as predicting basic students' emotions from the text (e.g., anxiety, bored, confused, excited). Estrada et al. (2020) used deep neural networks, such as convolutional neural networks and LSTMs, as well as evolutionary generative models, to classify sentiment polarity and emotions in collaborative learning environments.

As noted by Kastrati et al., there are relatively few works on recognizing students' emotions from text, despite the pedagogical importance of this task, and the growing prevalence of online learning (Sylwester, 1994; Kort et al., 2001). One possible reason for this scarcity of work is the reliance on hand-labeled data sets, which are costly and time-consuming to obtain. We directly address this gap in our work by using students' self-reported reactions in the form of emoji as proxies for their emotional state.

Emoji in Educational Context

The use of emoji as tools in computer-mediated communication is becoming an important focus of research for interpreting posts in many different types of online forums. Santamaría-Bonfil and López (2019); Kerslake and Wegerif (2017). Emoji can be used in two different ways. They can be employed in conjunction with text to convey reactions or to enrich the significance of a sentence or multiple sentences. Emojis can also serve as replacements for words, allowing for the expression of specific concepts or ideas, such as amusement or crown emoji in replacement for the word 'king'.

There are far fewer studies exploring how emojis can be used to help interpret posts in educational settings. One such study by Dunlap et al. (2016) investigated the use of emoticons in online learning and suggested that students utilize emojis to enhance communication, foster social presence, and create a sense of community within the online educational environment. Additionally, Vandergriff (2013) studied the educational chat room setting and found that college students, who hesitated to openly disagree with their peers, often used emoticons alongside non-standard or multiple punctuation marks to clarify their non-malicious intentions. These studies demonstrate how emojis can serve as valuable tools for students to interact with one another, while also functioning as helpful tags for their posts. A different perspective is put forth by Doiron (2018), who claims that while some emojis may help to clarify messages and have widely accepted meanings or functions, the meaning of other

emojis is more complex due to cultural differences in interpretation. In addition, the usage and meaning can also be influenced by the context. To avoid some of the issues that result from this complexity, the author recommends that emojis can be utilized in educational settings in a way that is similar to how they are used in commercial contexts for branding and marketing - by using a unique set of emojis that are widely known and clearly recognized. However, the above works fall short of using emoji as a proxy to understand students' reactions to the course content via their use in forums.

Zhang et al. (2017) investigated the use of emoji among students in the Nota Bene (NB) framework. The paper leveraged students' familiarity with social media conventions and their intimate knowledge of their own reactions. Students were presented with a limited set of emoji within the annotation platform NB, which allowed them to express their reaction to the course content. Each emoji in the set was associated with a unique hashtag label, which established a connection between a symbolic representation (emoji) and a corresponding label (hashtag), thereby precisely defining the meaning of the associated symbol. The development of an interface, that supports self-coding (i.e., self-reporting of emoji), provided the authors with a mechanism for detecting students' affective states by training a classifier to distinguish between confusion and curiosity.

Geller et al. (2020) expanded on this study and defined rules for confusion detection that are based on students' use of two types of emoji, confused and question emoji. They showed that the resulting rules closely align with the ground truth judgment of educational experts. We generalize both works to the more challenging task of recognizing multiple reactions and explore ways to facilitate this task by combining clustering methods with input from the course staff.

Emoji Prediction

Felbo et al. (2017) trained an LSTM architecture to predict emoji use on Twitter as proxies for users' emotional states. They show the model was able to generalize to other datasets containing self-reported emotional states. They employed clustering methods to learn relationships between 64 different emoji. We go beyond this work in several ways. First, we study the relationship between different emoji as reflected in the course content. Second, we use the clusters to build better predictive models.

Çöltekin and Rama (2018) used SVMs to predict emoji with a bi-grams feature set, combining both character n-grams and word n-grams and weighted by the TF-IDF score. This model achieved top performance in a recent competition for predicting 20 emoji on Twitter (SemEval 2018 task 2) (Barbieri et al., 2018). Zhang et al. (2020) used a BERT model for the emoji prediction task that outperformed the (Çöltekin & Rama, 2018) approach. We directly extend Zhang et al. model in adapting BERT to a biology course setting with additional pre-training over the course's previous data.

Study Setting

In the following sections, we will describe the setting of our study, the platform we used, and the course we gathered data from.

NB

The Nota Bene (NB) web application is an open-source social annotation tool that was developed at MIT (Zyto et al., 2012). The tool creates an educational environment that enables both synchronous and asynchronous collaborative annotation of online documents. NB has been used in hundreds of university courses and includes more than 40,000 registered student users. The main feature of NB is the in-place structure, this feature gives users the ability to directly annotate course content. Course content (PDF and HTML) is uploaded to the NB website by instructors.

Students annotate content by highlighting passages in the course reading and adding a post by typing a text field that appears in the margins, as seen in Fig. 1. Classmates are encouraged to reply to other students' posts and to answer any posted questions. NB posts are organized into threads, which consist of a starting post or question followed by all the replies made by students and instructors.

Placing comments directly in the reading itself allows students to interact with each other while they are reading the course material and provides context to the discussion. This structure has been shown to be beneficial for learning (Benitez et al., 2020; Zyto et al., 2012).

NB provides a default heat map that is based on the density of posts associated with paragraphs in the text. The opacity of the color of the text is determined by the number of annotations for that section, with lighter shades indicating fewer annotations and darker shades indicating more annotations (Zyto et al., 2012). Instructors can use the heat map to track which parts of the reading material generated more or fewer conversations. This can be useful information to instructors who wish to understand specifically where in the online content their students spend their time. In small classes (less than 75 students), where comment loads are generally moderate, this type of analysis can help an instructor rapidly identify specific areas of the online content that may need further review and revision or for which additional follow-up in class clarification could be beneficial. However, in large-enrollment classes the volume of comments generated typically overwhelm the text with overlapping yellow "highlights", reducing the discriminatory power of the heat map to identify regions of the text with meaningfully different student engagement.

The NB interface is shown in Fig. 1. It offers a graphical widget (bottom right) that presents students with a limited set of emoji that are each associated with a specific textual *hashtag*. Clicking on an emoji in the panel inserts the associated hashtag into the student comment. Students also have the option to type additional hashtags manually. Students and instructors can see the emoji in other students' posts. NB also offers options to filter posts on various criteria, including whether they contain specific emoji. The particular set of emoji used was initially designed independently by the course instructors to allow students the option to express reactions and opinions about course material, as well as invite assistance or participation from their peers or instructors. The initial instructor-derived emoji set was then refined using iterative rounds of student surveys and refinement, designed to evaluate and improve sets of emoji-hashtags pairs with regard to (a) the extent to which students feel a strong and unique association between the emojis and their assigned hashtags, and (b) the extent

Metabolism in General Biology

Cellular metabolism represents roughly 1/3 of the General Biology curriculum. While this may seem like a lot, we cover very little of what a classic course in metabolism covers, and a minuscule fraction of the metabolism that occurs on the planet. What we cover, however, is the very important foundational knowledge. You will learn about common chemical reactions that are associated with the transformation of life's molecular building blocks and about different core modes of energy transfer that you will encounter often in biology. The energy story and the design challenge rubrics introduced earlier will become increasingly important in these next few modules and beyond.

What have we learned? How will it relate to metabolism?

1. We have focused on the identification and chemical properties of common biological functional groups. As we dive into metabolism, this will help you be familiar with and sometimes even predict the chemical nature/reactivity of compounds you have never seen before.
2. We have practiced recognizing and classifying molecules into four major functional groups. This will help you as we discuss how to build and break down these molecules.
3. We have learned some basic thermodynamics. This gives us a common set of concepts with which to discuss whether a biochemical reaction or process is likely to occur, and if so in which direction and how fast. This will be critical as we consider some key reactions that take place in metabolism.
4. We have learned and practiced the energy story rubric. This will allow us to study new biochemical reactions and to discuss them with a common, consistent language and approach which also reinforces the lessons we learned about thermodynamics.

An overview of this section

- We will introduce an important concept called **reduction potential** and you will be given the opportunity to use a **redox tower**. There is also a discussion on redox chemistry in your discussion manual. Use both resources.
- We will introduce two major players in metabolism, **ATP** and **NADH**. We expect you to recognize their structures if shown on an exam.
- We will cover the metabolic pathway glycolysis. Keep in mind that we want you to look at any reaction and tell us an energy story of that reaction. You should not try to memorize these pathways (though it will help to remember some big picture things - we will stress these). Often we will give you the pathway as a figure on the exams. Glycolysis produces 2 ATP via a process called substrate level phosphorylation, 2 NADH and 2 pyruvate compounds.
- We will use the reactions of the TCA cycle to create multiple examples of energy stories. The TCA cycle will also produce more ATP, NADH and oxidize glucose into CO₂.
- We will look at an alternative pathway to that of the TCA cycle, fermentation. Here, for the first time, we will see NADH used as a reactant in a metabolic reaction.
- We will follow NADH to the end of its journey, as it donates its electrons to the electron transport chain (ETC). In this module, you will need to use a redox tower. The ETC produces a proton gradient. No ATP is directly generated in this process. However, the proton gradient is then used by the cell to run an enzyme called ATP synthase, which catalyzes the reaction ADP + Pi → ATP. This method of ATP production (called oxidative phosphorylation) results in more ATP being produced than through substrate level phosphorylation.
- And finally, we will go through the process of photosynthesis.

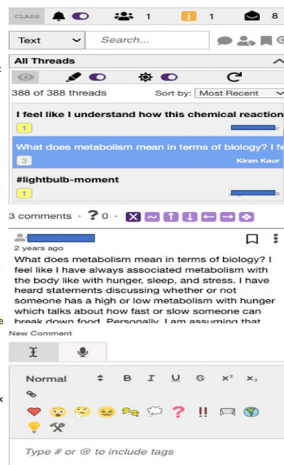


Fig. 1 Nota Bene GUI: Course reading material appears on the left, with highlighted colors representing density in terms of number of comments. Students' discussions appear on the right

to which the emoji set adequately captures the breadth of responses that students wish to be represented. The process, along with the surveys is detailed in Kim (2022), Chapter 6.

A full list of the emoji, their associated hashtags, and their intended uses is shown in Table 1.¹ Some examples of emoji and their intended uses are as follows: The #lost emoji expresses confusion or the feeling of being overwhelmed with a section of the text; the #lets-discuss emoji invites students and instructors to contribute to a discussion; the #interesting-topic emoji identifies a topic or idea that the student found interesting; the #just-curious emoji expresses student curiosity about a topic; the #question emoji requests help from the course staff or peers about a section of text.

Typical posts concerning with questions about the course content and tagged with emoji are shown in the following examples:

“Does the temperature decrease because there is more pressure? But that wouldn't make sense #lost”

“Is archaea a type of bacteria? are they related in a way? #just-curious”

The FYBIO Course

FYBIO (First Year BIOlogy) is a general biology course required for all life sciences majors at a large public research university in the western United States and is typically taken by students during their first year of study. The course forms part of the core curriculum for students in the biological sciences but is also followed by students

¹ For the remainder of this paper, we will use the term 'emoji' to represent an emoji or its associated hashtag.

Table 1 Emoji in NB, their associated hashtags, intended uses, and percentage of usage

Emoji	Associated hashtag	Intended use	Percentage of use
?	#question	you would like the instructor to address a specific question in class.	18.9%
💭	#i-think	you have an idea to share but are not sure if it is accurate.	18.8%
🤔	#just-curious	you would like to learn more about a certain topic that you are curious about or intrigued by.	12.8%
!!	#important	if you found a topic or idea to be important.	11.5%
❤️	#interesting-topic	if you found a topic or idea to be interesting	11%
💡	#lightbulb-moment	something (e.g., course material, a peer's post) led you to a new understanding (an "ah-ah" moment).	7.7%
💬	#lets-discuss	you would like to invite your fellow classmates and instructors to weigh in, in an NB discussion.	4.8%
🌐	#real-world-application	A topic or idea has any implication in the real world.	4.5%
😮	#surprised	you found a topic or idea to be surprising.	4.1%
😵	#lost	you are totally confused or overwhelmed about a topic and have no idea what question to even ask.	3.4%
🗣️	#learning-goal	you identify a topic or idea related to the course objectives or reading-specific learning goals	2.3%

from more than 60 other majors, ranging from the social and engineering disciplines. Depending on the academic term it is offered, the course consists of 25 or 26 lectures.

For many years, students in the course have been required to use NB to annotate reading assignments from a custom online textbook hosted by LibreTexts² where the course staff post reading materials as individual HTML files each with content related to an individual lecture. Readings were posted before each week of instruction and the students were required to read these materials and to provide three substantial posts in NB before the end of each week. Students received additional credit (up to a maximum of 10 percent of the final course grade) for including at least one emoji in at least one of their posts. The idea behind this assignment structure was to encourage active participation in forum discussions.

For analysis, we have collected 83,380 unique student posts from two instances of the FYBIO course Spring 2021 and Winter 2022. In total, 63,702 posts contained at least one emoji. All students filled out a consent form allowing their data to be anonymized and analyzed for the sake of this study. The study was reviewed by the IRB (1456274-1) and deemed exempt.

Empirical Analysis of Students' Emoji use

In this section, we study the frequency of emoji use in students' posts, as well as the relationship between the use of emoji in different paragraphs in the reading material.

² <https://libretexts.org/>

Analysis at the Post Level

Table 1 shows the percentage of emoji use in FYBIO at the post level. The table shows that the #question and #i-think emoji were used most often. Both of these emoji reflect uncertainty about the material and invite participation from instructors and students. The emoji expressing more direct requests for participation (e.g., #lets-discuss, #lost) were used much less frequently. This may reflect resistance towards revealing information about understanding that may have an adverse effect on students (e.g., peer pressure or getting a lower grade).

By analyzing students' posts (Fig. 2) we have found that about 73% of posts included at least one emoji. Ten percent of the posts contained two emoji and less than 1% of the posts contained 3 or 4 emoji.

To delve deeper into the use of emoji by students in the discussion forum, we studied emoji use at the individual student level. Our objective was to determine the extent to which students utilize a diverse range of emoji throughout various lectures, or if they gravitate toward a particular subset of emoji. Figure 3 illustrates the distribution of unique emoji usage among students. It is evident from the data that a significant proportion of students tend to employ between 2 and 7 distinct emoji. The average number of distinct emoji used was 6.425, with a standard deviation of 2.578 and a median of 7.

Analysis at the Paragraph Level

We used the fact that students' posts are linked to a specific paragraph within the reading material to study the relationship between emoji use with specific paragraph-level units of the text. This can provide instructors with more detailed information about student reactions to discrete sections of the readings and by extension to different topics, or subtopics, in the course. A natural question is why we examined the number of unique emoji used in a paragraph. The annotated text is a subset of words extracted from a paragraph in the textbook, typically focusing on a single subject. On average,




Fig. 2 Histogram over the number of emoji used in posts

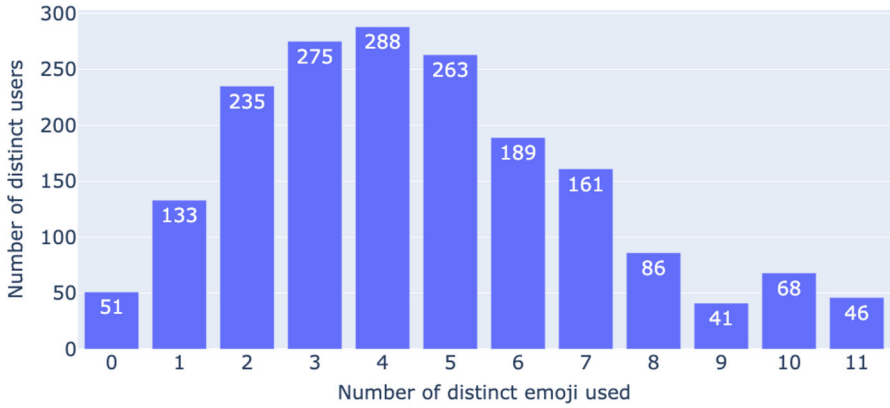


Fig. 3 Number of distinct emoji used by individual students

a paragraph consists of approximately 128 words, while the annotated text comprises about 14 words. The data to create the correlation matrix is continuous (frequency of each emoji in the paragraph), thus for each paragraph we computed the frequency of each emoji inside that paragraph, resulting in a matrix of size $N \times 11$ where N equals the number of unique paragraphs in the data. We computed a correlation matrix over this matrix.

Figure 4 shows the heat map representing the correlation between emoji used in paragraphs in the readings. We imagine at least two possible explanations for the

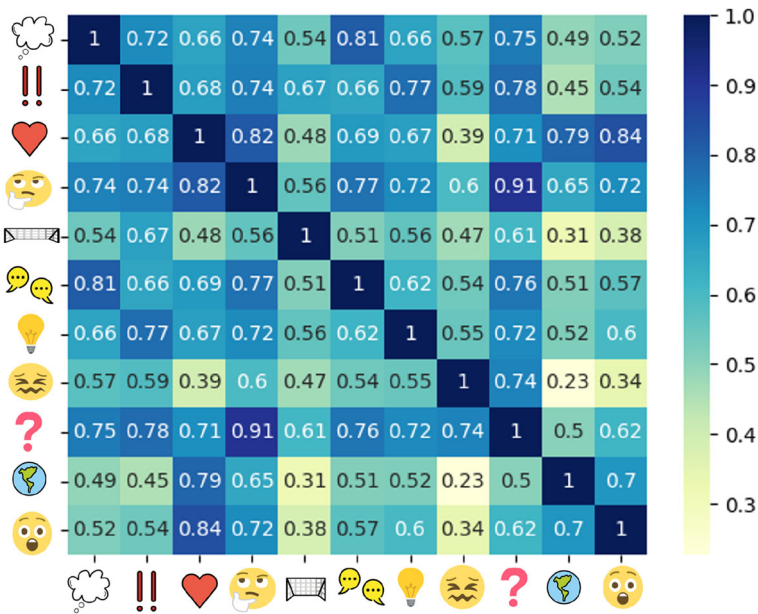


Fig. 4 Correlation Matrix of emoji within paragraphs

correlated use of emoji at the paragraph level: (a) There is a relationship between the student reactions reported by each emoji. (b) Students may use different emoji to express the same reaction.

An example for explanation (a) would be the 0.79 correlation score between the emoji #interesting-topic and #real-world-application. It seems quite reasonable that students would find reading about real-world applications interesting. Moreover, students explicitly reporting which real-world applications of course content are interesting to them is valuable to instructors and content creators by providing insight into what topics their students find relevant and meaningful.

An example for explanation (b) would be a 0.91 correlation score between #question and #just-curious. We hypothesize that these two emoji can be used to express two reactions but the result will be the same, which is a question. We can see that because when people are curious, they tend to ask questions. Although it appears that both will end up with a question, the instructors wanted to differentiate between these two emoji in order to detect where students find materials that cause them curiosity versus students' questions about course material.

Emoji and Cognitive Engagement

In order to gain more insights about the emoji and their benefit to student learning we introduce an analysis relating emoji and *cognitive engagement (CE)* (Chi & Wylie, 2014). The cognitive engagement of a student's post is a measure of how deeply a student is thinking about course material and in which way the student is interacting with the course material, which has been shown to correlate with learning gains. Yogeve et al. (2018); Chi and Wylie (2014)

We consider three types of CE categories, in increasing order of engagement, similarly to Yogeve et al. (2018).

- **A:** Active. Indicating *attention to the course in some way without referring to the content*. And, displaying *engagement with specific course materials* in the annotation by paraphrasing, repeating, or mapping resources, *without providing deep insight*. An example of a student post that maps resources (which is taken directly from the text) without providing deep insight:

"Gene expression is the appearance of a phenotype/genotype on an individual."

- **C:** Constructive. *Introducing a statement or a question* related to the reading that includes a new idea or refers to external sources, *but does not reason about them*. And, *displaying reasoning about the reading*, for example by explaining a phenomenon. An example of a student displaying reasoning about the reading:

"When blood sugar levels go down, the pancreas releases a hormone called glucagon that tells the liver to process the stored sugars and release them into the bloodstream."

- **I:** Interactive. Displaying constructive reasoning (C) and interacting with previous comments. An example of a student replying to another student, adding an explanation while showing reasoning:

“Additionally, unrestricted gene expression can lead to potentially malignant conditions like cancer.”

The fourth Passive category cannot apply to comments that students are actively authoring.

We used a predictive model to classify the CE label corresponding to a student post in a similar way to Yogev et al. (2018). We trained and ran the model on a subset of the data (winter 2021).

Table 2 compares emoji use with students’ CE over posts. As shown in the table, #i-think and #lets-discuss emoji mostly align with the ‘C’ category CE, which is used for posts that describe a new idea but do not reason about them. This category fits the emoji description, which states that “you have an idea to share but are not sure if it is accurate” and “invite fellow classmates in a discussion.” respectively (Table 1).

The #important and #just-curious emoji mostly align with ‘A’-category CE, which is used for posts that indicate attention to the course in some way or referring to content without providing insight This category fits the emoji description, which states that “you found a topic to be important (or curious)” (Table 1).

Similarly, the #lost and #question emoji align more with the ‘A’-category CE. These posts commonly refer to a question or express confusion in regard to the learning material, thus they are not introducing new ideas nor displaying reasoning about the reading.












Lastly, we note that the scarcity of ‘I’ labels in our findings can be attributed to the fact that most posts are new threads and do not respond to others.

Discussion

Our analysis revealed that 73% of the posts contained at least a single emoji. This shows that students’ use of this tool went beyond the assigned requirement of a single emoji per reading (see [The FYBIO Course](#)). The wide range of distinct emoji used (see Fig. 5) implies that students chose to express diverse reactions to the course content and to others’ posts.

We included an analysis at the paragraph level ([Analysis At The Paragraph Level](#)) due to the characteristic focus on a singular subject within each course paragraph. Our objective was to assess the variation in emoji usage within a specific subject, an aspect best examined when analyzing at the paragraph level. The annotated texts, averaging 14 words in length, exhibit a limited repertoire of emojis compared to the paragraph level. This observation indicates that analyzing emoji at the paragraph level

Table 2 The frequency of posts labeled as ‘A’, ‘C’ or ‘I’ for each emoji

Metric											
A	3,644	1,101	2,318	1,690	1,155	948	508	346	466	449	474
C	864	2,615	1,304	852	1,268	825	1,386	546	405	107	205
I	4	135	31	18	66	20	75	36	6	0	0




Fig. 5 Frequency of distinct emoji usage across paragraphs

reveals more insights about their attitudes toward course consent and may provide better insights to course instructors.

The analysis in [Emoji and Cognitive Engagement](#) shows that there is a level of agreement between the pedagogical information that is entailed in the emoji, and the cognitive engagement of the students' posts. This provides support to our claim that emoji reflect students' actual attitudes about the course material.

User Study

In this section, we studied whether providing instructors with information about students' use of emoji aids them in their teaching and course planning. A key tool in the study was a new heat map that displays aggregate information about students' use of emoji in the course. Course instructors identified the following requirements from the heat map: 1) identifying sections of the reading material that students are interested in, and 2) identifying sections of the reading material that students find difficult

To evaluate whether the heat map satisfied these requirements, we conducted a user study with five members of the FYBIO teaching staff: two instructors, two teaching assistants, and one class assistant.

All of the participants were experts on the use of NB, and four of them were on the instructional team of FYBIO for the past 5 years. Participation in the study was voluntary.

Emoji Heat map

Our emoji heat map is an extension of the original NB implicit heat map (shown in Fig. 1) the highlights marking the annotated text form a heat map since they are darker where there are more comments. An example of the emoji heat map is presented in Fig. 6.

Carbohydrates

Carbohydrates are one of the four main classes of macromolecules that make up all cells and are an essential part of our diet; grains, fruits, and vegetables are all natural sources. While we may be most familiar with the role carbohydrates play in nutrition, they also have a variety of other essential functions in humans, animals, plants, and bacteria. In this section, we will discuss and review basic concepts of carbohydrate structure and nomenclature, and a variety of functions they play in cells.

Molecular structures

In their simplest form, carbohydrates can be represented by the stoichiometric formula $(CH_2O)_n$, where n is the number of carbons in the molecule. For simple carbohydrates, the ratio of carbon-to-hydrogen-to-oxygen in the molecule is 1:2:1. This formula also explains the origin of the term "carbohydrate": the components are carbon ("carbo") and the components of water ("hydrate"). Simple carbohydrates are classified into three subtypes: monosaccharides, disaccharides, and polysaccharides, which will be discussed below. While simple carbohydrates fall nicely into this 1:2:1 ratio, carbohydrates can also be structurally more complex. For example, many carbohydrates contain functional groups (remember them from our basic discussion about chemistry) besides the obvious hydroxyl. For example, carbohydrates can have phosphates or amino groups substituted at a variety of sites within the molecule. These functional groups can provide additional properties to the molecule and will alter its overall function. However, even with these types of substitutions, the basic overall structure of the carbohydrate is retained and easily identified.

Nomenclature

One issue with carbohydrate chemistry is the nomenclature. Here are a few quick and simple rules:

1. Simple carbohydrates, such as glucose, lactose, or dextrose, end with an "-ose."
2. Simple carbohydrates can be classified based on the number of carbon atoms in the molecule, as with triose (three carbons), pentose (five carbons), or hexose (six carbons).
3. Simple carbohydrates can be classified based on the functional group found in the molecule, i.e. ketose (contains a ketone) or aldose (contains an aldehyde).
4. Polysaccharides are often organized by the number of sugar molecules in the chain, such as in a monosaccharide, disaccharide, or trisaccharide.

For a short video on carbohydrate classification, see the 10-minute Khan Academy video by clicking [here](#).

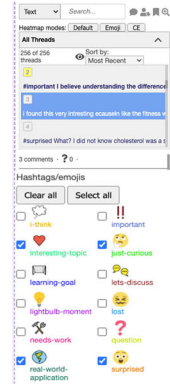


Fig. 6 Emoji heat map on Nota Bene GUI: On the left, Course reading material highlighted text represents the density of emoji usage. Right: Students' discussion along with the emoji pictures and their corresponding colors. Users can choose the comments that contain specific emoji by checking the emoji in the checkbox

The opacity of the highlighted color in the heat map is determined by the number of emoji used in the selected sections in the reading material (darker implies higher usage). The color of an area in the text is defined by a specific emoji such that each emoji is represented by a unique color. The use of a color-coding system allows instructors to identify patterns and trends directly in the reading material. All of the participants were presented with a demonstration of the emoji heat map and were provided with a general description of the study goals.³

Study procedure

The user study consisted of two parts, outlined in Fig. 7.

In the first part of the study, participants were asked to complete two tasks that required interaction with the NB platform, in line with our hypothesis.

- T1. Identify reading material that students may be interested in.
- T2. Identify reading material that students may find difficult.

The reading material in the study was sampled from lecture 5 in the course, which focused on pH and pKa, and is considered to be challenging by the course instructors. This reading contains 32 different sections. Prior work shows that students in introductory biology courses have often faced difficulties in conceptualizing pH and pKa and in the use of logarithms to solve practical problems in these topics (Watters & Watters, 2006).

The study lasted about an hour. For the first part of the study, participants were not able to use the emoji heat map. Their responses were based on an existing NB tool for filtering comments according to keywords and hashtags, as well as using the original NB heat map showing distribution over comments (see Fig. 1).

³ IRB approval for this instructor-focused portion of our study was obtained from the Human subjects research committee of Ben-Gurion University.

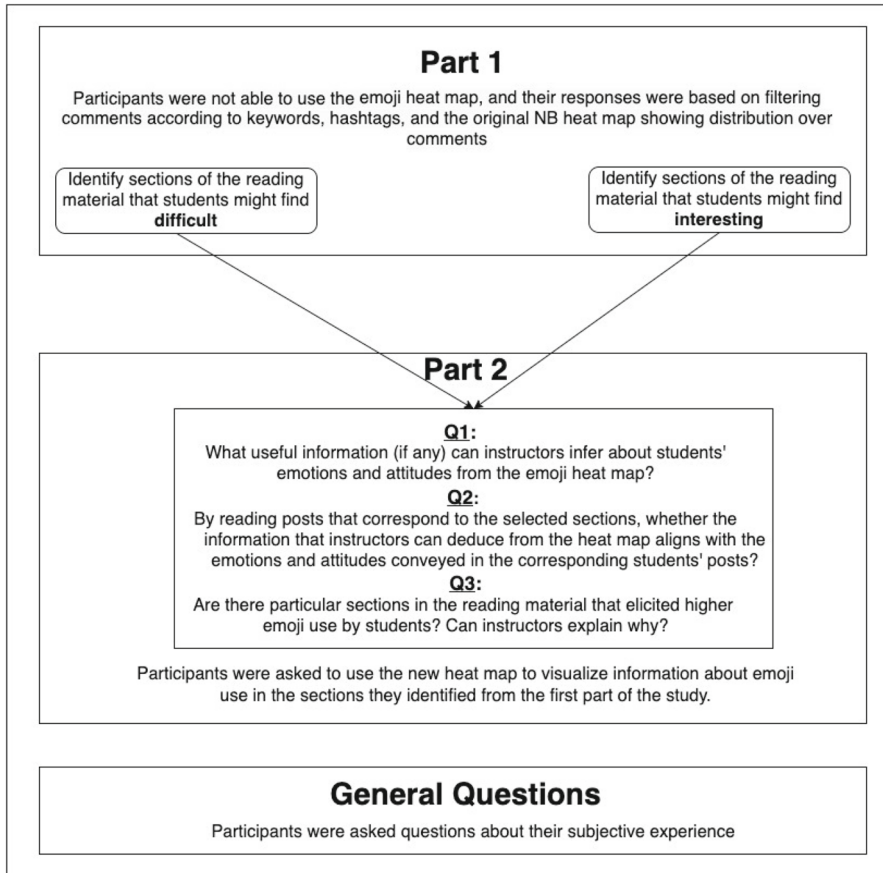


Fig. 7 Procedure of the two-part User Study

The second part of the study evaluated the additional benefit provided by the emoji heat map to help instructors understand students' reactions. For each of the first two tasks (identifying interesting and difficult readings), participants were asked the following questions.

- Q1: What useful information (if any) can instructors infer about students' reactions from the emoji heat map?
- Q2: Does this information align with the reactions? that participants infer when reading the actual posts?
- Q3: Are there particular sections in the reading material that elicited higher emoji use by students? Can instructors explain why?

Finally, participants were asked questions about their subjective experience using the emoji heat map. These questions were taken in part from the USE questionnaire (Lund, 2001) which is commonly used in the literature to query systems' usefulness, satisfaction and ease of use. The questionnaire is constructed as a five-

point Likert rating scale, where users are asked to rate their agreement with a series of statements, ranging from strongly disagree to strongly agree.

User Study Results

We analyze participants' behavior in the first part of the study followed by the second part of the study. We identify individual participants using the notation $P\#$ where $\#$ is a participant number.

Part 1: Selecting Readings

For the T1 task (identifying interesting sections of the reading material), four out of the five participants (P1, P3, P4, and P5) filtered comments in NB for posts containing #interesting-topic, #just-curious, #real-world-application, and #surprised. The remaining participant (P2) filtered comments in NB using the keywords "interesting" and "fascinating".

For the T2 task (identifying difficult sections of the reading material), three out of the five participants (P1, P3, and P5) filtered comments in NB for posts containing #lost and # questions. The remaining two participants (P2 and P4) relied on their prior teaching experience of FYBIO and used the original NB heat map to select difficult sections in the reading material. As an example, P2 used the original heat map to visualize high frequency of students' comments relating to the topic of proteins, which they stated was "confusing for students", and proceeded to analyze the density of comments on the relevant sections in the reading material.

Part 2: Identifying Interesting Sections in Readings

We describe participants' responses to questions Q1-Q3 after applying the emoji heat map to the sections of the reading material they considered to be interesting. With respect to question Q1 (inferring useful information from the heat map), four out of five of the participants (P1- P4) agreed that the new heat map provides insight into interesting topics for students, the participants did manage to find interesting parts in the first part, however they had to link comments to the content instead of just looking at the content. They all mentioned the green highlighted reading material in the heat map that indicated posts with #interesting-topic, #just-curious. An example of a participant's response (P4) was "I was very interested that the Sterols paragraph had many more #interesting-topic comments and #important comments than I was expecting" Participant P5 stated that the heat map exhibits a high amount of green highlights, and they added that they could also observe other colors in the heat map which imply difficulty.

With respect to question Q2 (whether the information in the heat map aligns with students' posts), after reading comments associated with the paragraph they chose, all five participants agreed that the information conveyed in the heat map aligned with the reactions that students' posts exhibited in the corresponding sections and provided justification. An example of a justification by participant P4 on a post exhibiting a

#real-world-application emoji was that “real-world applications are always interesting to students.”

With respect to question Q3 (which sections elicit higher emoji use), all participants responded that the sections of the reading material with the highest emoji density were those with real-world applications or subjects that students were already familiar with. Participant P5 stated: “Yes, the students are likely familiar with the term lipid/fats, so they are interested in learning more.”

Part 2: Identifying Difficult Sections in Readings

This section describes participants’ responses to questions Q1-Q3 after applying the emoji heat map to the sections of the reading material they considered to be difficult.

With respect to Question Q1 (inferring useful information from the heat map), two of the participants (P1 and P5) stated that their selected readings exhibited a high occurrence of the #question emoji. An example of one of these statements from participant P5 stated: “The new heat map allows me to more easily see where in a paragraph a student might be confused or otherwise engaged visually, instead of having to click through the student comments” which highlights that the heat map helped the participant identify problematic sections of the reading that he did not notice.

The other participants (P2 - P4) did not make a concrete statement about difficult topics in their selected readings. Specifically, Participant P3 stated that they saw readings highlighted with green color (representing interesting topics) “Yes, the second paragraph was all shades of green, but very little confusion colors” All of the participants noted the low frequency of the #lost emoji in their selected readings.

With respect to question Q2 (whether the information in the heat map aligns with students’ posts), participants P1 and P5 agreed that the information conveyed in the heat map aligned with students’ reactions. As an example, participant P5 Found the following comment in a paragraph that exhibited question and lost colors. “Okay I think I got this but the part I’m confused about is visualizing what secondary, tertiary, and/or quaternary structures may change too. #question”

Participants P2-P4 noted that the most frequent emoji in their selected readings expressed interest rather than difficulty and that the information in the heat map aligned with interest. P3 responded that the section they chose exhibited interest and intrigue, although they said this section was difficult

With respect to question Q3 (which sections elicit higher emoji use), all of the participants successfully located readings with a high density of #lost and #question emoji in students’ posts. An example of a justification from Participant P1 to such a post was “That topic (Titration Graph) is known to cause issues for students and they often need to get help outside of lectures and readings on it.” Moreover, P5 added that the number of #lost and #question emoji emphasizes how crucial it is for instructors to examine that particular section.

Results: Satisfaction and Ease-of-use

We describe some of participants’ responses to the satisfaction and ease-of-use questionnaire. The full questionnaire and summary of participants’ responses are shown in

Table 3. Participants were requested to rate their agreement (on a scale of 1 to 5) with pre-defined statements and questions from USE questionnaires and general questions with regard to their subjective experience using the emoji heat map.

With regard to ease-of-use (S13-S15), participants' high scores imply that the heat map can be learned quickly and be used without instructions. This is also apparent in participants' responses with regard to their satisfaction with the heat map (S9, S10, S16) and the usefulness they derive from the heat map (S7, S11, S12). In particular, participants stated that the heat map will save them time and helps them be more effective.

There was general agreement among participants that the heat map will be used in the future (S10). There was a higher willingness to use it for revising course content (S9) than to use to plan classes while a course is in session (S8). One participant stated that while revising course content the heat map can show aggregated information

Table 3 General and USE questionnaires with their corresponding average score

Index	Statement/Question	Avg. Score (SD)
S1	Does the heat map align with the places I thought the students found difficult	3.6 (0.49)
S2	Does the heat map helps find difficult sections of the reading that I did not notice	4.2 (1.166)
S3	Does the heat map align with the places I thought the students found it interesting	4.4 (0.49)
S4	Does the heat map helps find interesting sections of the reading that I did not notice	3.75 (1.089)
S5	Does the heat map helps you find struggling students that I did not notice	2.4 (1.497)
S6	Does the heat map helps you find students' posts in a specific section where I would like to join the conversation or intervene	3.2 (1.327)
S7	I will act when I see a high density of emoji	4.6 (0.49)
S8	Will you use the heat map to prepare a future class?	3.6 (1.02)
S9	Will you use the heat map to analyze which sections of the reading material need to be revised for future class sessions?	4.4 (1.2)
S10	I will use the heat map next time I am teaching	4.2 (1.166)
S11	The emoji heat map helps me to be more effective, by effective we mean that the heat maps help you to find what you looked for in the reading material.	4.2 (0.4)
S12	The emoji heat map will save me time	5 (0)
S13	The emoji heat map is too complicated to use	1.8 (0.4)
S14	I can use the emoji heat map without instructions	4.8 (0.4)
S15	I have managed to learn to use the emoji heat map quickly	4.8 (0.4)
S16	I would recommend other teachers to use the emoji heat map	4.8 (0.4)

The number in parentheses denotes the standard deviation

regarding the students' reactions to a specific course content which otherwise can be achieved only by reading a large number of comments.

Discussion

All of the participants acknowledged that it was easier to identify students' reactions relating to the emoji that are more prevalent in the heat map. This statement is backed by a justification from P4 which stated "The new heat map more easily shows if students have the same or different reactions regarding specific places in the reading."

The participants pointed out the benefits of using the heat map. They stated that without the heat map, it would have been more time-consuming and difficult to identify sections of the reading material that exhibit high-engagement student reactions without the help of the tool. This is particularly true when courses have a high number of students or when courses are frequently being revised. P4 and P1 gave examples of such statements. P4 said "The new heat map more easily shows if students have the same or different reactions regarding specific places in the reading" and P1 said "The density can show me where the students may be interacting more or where they decided to comment, but the emoji map shows the possible types of comments."

We saw that the participants were willing to use the heat map for revising course content and for planning classes while a course is in session. The participants stated that it would assist them in identifying sections of the reading material that may require further examination in the context of class preparation and course revision, as it provides an overview of student understanding of those topics. These are reflected in two statements from participants. P1 stated "I think this new heat map is useful in finding topics that are confusing or need rewriting." P3 stated that "It can help point out which topics need to be expanded or connected with the real world"

Moreover, we studied whether providing instructors with information about students' use of emoji aid them in their teaching and course planning. By evaluating the participants' behavior when the new heat map was not available to them we noticed that most participants used students' emoji filtering method as tools to solve their tasks, emphasizing their benefit to instructors. Moreover, we allowed the participants to choose a paragraph based on their own experience in order to assess the alignment between the emoji heat map and the knowledge of experienced instructors.

With respect to our first hypothesis, we noticed that instructors had an easier time finding sections in the reading material that students might feel interested and it can be concluded from the participants' responses that exhibit a high score when asked whether the heat map aligned with the place they thought the students found interesting 4.4 in compared to 3.6 when asked the same question on difficult. When asking the participants for an explanation about the alignment, P5 stated "Yes, the students are likely familiar with the lipid/fats, so they are interested in learning more."

With respect to the second hypothesis, we noticed that the participants had a harder time identifying sections from the reading material that contained the #lost emoji due to students not frequently using this particular emoji. The reason #lost emoji is not frequently used may be because students feel that information about comprehension might potentially harm students, such as peer pressure or receiving a lower grade.

However, the participants reported that the heat map helped them identify problematic sections of the reading that they did not notice, and locate smaller parts of the reading material that contained students' posts that exhibits difficulty. The issue regarding finding difficulty can be addressed by the use of the filtering bottom alongside the emoji heat map to conceal other emoji. We also asked whether the heat map will help instructors find struggling students which resulted in a low score (2.4). This implies that the heat map is limited in its capabilities to aid instructors in such tasks.

We can conclude that the heat map is useful, satisfying, and easy to use due to the fact that questions regarding these points exhibited high scores. A notable point is a unanimous answer to the question of whether the emoji heat map will save me time, where all the participants answer it will (5). The participants state that most of the saved time will be due to not having to go over comments in order to understand specific reactions.

We list several limitations of this study. First, we did not directly compare the use of two different versions of the NB GUI, with and without the emoji heat map. Because we had limited access to teachers' time, we preferred to collect teachers' subjective opinions about the benefit of emoji in their own courses. On the other hand, by conducting semi-structured interviews with the actual teaching staff, the results are ecologically valid. Also, the small number of participants who participated in the study did not allow us to perform a statistical analysis of the benefit of the emoji heat map.

Predicting Emoji Directly From Students' Posts

In this section, we study the use of computational language models based on students' discussions. We use language models to classify emoji in students' posts in the FYBIO course. This model could be potentially used to infer students' attitudes in settings where posts don't contain emoji, and to generalize the emoji classification to courses other than biology.

Classifying Emoji using ML

We've discussed the useful insights that instructors can draw from student emojis. But students don't always take the time to enter those emoji. Thus, in this section, we describe the use of language models and supervised learning to *infer* students' use of emoji at the post level. This can potentially aid instructors in making sense of students' affective states in situations where emoji are not used, by individual students or by entire classes. We define a multi-class prediction task where the target class is the set of emoji shown in Table 1. Our approach provides a mapping from a student post, combined with the annotated text to the most relevant emoji based solely on the text from the post.

We employed BERT, a pre-trained, transformer-based language model that is commonly used in state-of-the-art natural language tasks (Devlin et al., 2018). We used an open-source BERT configuration ("bert-base-uncased") that was trained on broad

domain corpora (English Wikipedia and BooksCorpus).⁴ The architecture contains 12 layers, 768 hidden units, and 12 attention heads. BERT is trained on a pre-training task where it predicts masked words given their context, in a way that it learns the relationship between words in a sentence. This pre-training task allows the model to learn the context-dependent representation of words in a language. Since we are interested in utilizing BERT for a classification task, the pre-trained BERT model can be used as an encoding layer, where the input sentences are transformed into a fixed-length vector representation. The encoded representation then serves as an input to a classifier to make the final prediction (Devlin et al., 2018).

For predicting the emoji, we defined a neural network, on top of the pre-trained BERT model, that was trained for the emoji classification task using 768 nodes with a softmax activation function. The reason for the 768 nodes is that the size of the representation is determined by the BERT model architecture, and in the case of the BERT base, it is 768-dimensional. This architecture, which from now on will be called BERT-FYBIO will be used in all of our experiments. The input to this network is the student's post, alongside the annotated text using the language representation model. The output of the network is a probability distribution over the 11 emoji in the target class.

Because we employed a multi-class method that defines a single output for each instance, a duplicate instance was created for each post that has more than one unique emoji in the training set. An identical post will be shared in each of these situations but with a different emoji label. We have taken the FYBIO dataset and further preprocessed it by cleaning HTML patterns, stripping web links, and removing posts that only contain emoji as a post.

In order to verify our result, an 8-Fold Cross-Validation was employed by random sampling fold at each iteration.

The nature of the distribution of the emoji in the data is roughly imbalanced (Table 1) To address this, we modified the model's architecture by implementing weighted cross-entropy loss, which allows us to assign a weight to each class in proportion to its frequency in the data set. This helps the model to account for the imbalance and adjust its predictions accordingly. This strategy was used in all of our models in the paper.

We trained the BERT-FYBIO for three epochs with a learning rate of $2e - 5$, using a maximal sentence length of 300 and batch size of 32. Table 4 breaks down the performance of the model according to each of the emoji. The table shows a positive relationship between the amount of training data for a given emoji (according to Table 1) and the prediction performance of the model for the emoji.

We compared BERT-FYBIO to a bi-directional LSTM architecture, similar to the one used by Felbo et al. (2017); Baziotis et al. (2018) to classify emoji in social media, which achieved best performance in their experiments. The model included five layers. One layer was used to embed words in students' posts as high-dimensional vectors; three layers for classification consisting of 64 LSTM units (32 units in each direction); the final layer was an attention layer that connects words in posts with preceding and succeeding words while computing the importance of each word with the corresponding label. The model was implemented using Python's keras package (Chollet et al.,

⁴ <https://huggingface.co/transformers/modeldoc/bert.html>

Table 4 BERT-FYBIO prediction score, divided by prediction output score for each emoji

Metric	?	☁️	🤔	!!	❤️	💡	💬	🌐	😮	😓	🗨️
<i>Precision</i>	0.54	0.54	0.41	0.44	0.39	0.41	0.06	0.37	0.42	0.41	0.3
<i>Recall</i>	0.82	0.7	0.24	0.57	0.39	0.25	0	0.23	0.27	0.19	0.01
<i>F-1 Score</i>	0.65	0.6	0.3	0.5	0.39	0.31	0	0.28	0.33	0.26	0.03

2015). We used a separate pre-training process using word2vec (Mikolov et al., 2013) to construct a high dimensional 200-sized vector representation of words in students' posts. The pre-training used the FYBIO textbook, as well as students' posts from 2020, and was implemented using Python's Gensim package (Rehurek & Sojka, 2011).

Table 5 compares the BERT and LSTM models when classifying emoji according to precision, recall, and weighted F-1 scores. We can see that the BERT model outperforms the Bi-LSTM model in all three metrics by a significant margin (McNemar's test, $p < 1.16 \cdot e^{-35}$). We attribute this difference to the pre-trained language model in BERT which allows it to generalize to domains with low amounts of training data (Devlin et al., 2018).

Classifying Groups of Emoji

In this section, we study the nature of emoji that are often used together in paragraphs to see if we can cluster the emoji into pedagogically meaningful categories and predict the categories, rather than individual emoji for better-capturing reactions from students' posts while still capturing the reactions that are important to the instructors. Felbo et al. (2017) clustered emoji based on the relationships between predicted probabilities of a computational model. We follow their approach and apply a hierarchical clustering algorithm with average linkage on the correlation matrix on the predicted probabilities of the BERT-FYBIO model depicted in [Classifying Emoji using ML](#). Each emoji is represented as a vector of predicted probabilities for each post in the test set. To illustrate this approach, consider the following post from the test set:

Is there an operon for every single metabolite? [trp] operon applies to tryptophan while lac operon applies to lactose. how many different types of operons are there and how do they function differently?

The BERT-FYBIO model, applied to this post, outputted the highest probability for #question and #just-curious (0.467 and 0.45 probability, respectively) and low proba-

Table 5 Performance comparison

Model	Precision	Recall	F-1 Score
BERT	40.2%	43.3%	40.7%
LSTM	26.6%	32.8%	29.3%

bility for the rest of the emoji (e.g., 0.008 probability for #interesting-topic, and 0.003 probability for #lightbulb-moment emoji). Note that because we are using the softmax function, the output vector is transformed to a vector of probabilities that add up to 1, although the two emoji have a combined probability of approximately 0.92. The true label for this post was #question. The fact that the model is likely to assign a high or low probability to both #question and #just-curious emoji for a given post suggests that they should form a single category since the model is having a hard time distinguishing between these labels from a given student post.

Figure 8 shows the dendrogram that is outputted by applying hierarchical clustering. The height of each node in the y-axis is proportional to the value of the intergroup dissimilarity between its child nodes. The distance threshold of the clustering is used to determine which nodes to put in the same cluster. The challenge is to find a threshold that improves predictive power but does not lose the ability to discriminate between reactions by over-clustering.

The distance between the merged clusters generally increases with the level of the merging, and the clusters become less robust. For example, the cluster containing the (#question, #just-curious) emoji has a lower height than the cluster containing the (#question, #just-curious, #lost) emoji, indicating that the dissimilarity between the emoji in the former cluster is lower than in the latter cluster. Higher thresholds in the dendrogram creates larger emoji categories and can potentially facilitate prediction at the category label (the threshold parameter is a number in the range [0, 2.5] in our experiment). On the other hand, setting a high threshold reduces the granularity of the information we can provide to teachers. Essentially, we are trading precision (exactly which emoji to predict) for recall (predicting the proper general category of emoji).

To address this trade-off, in concurrence with the FYBIO instructions, we determined a threshold of 1. This threshold grouped the 11 emoji into three emoji pair clusters and 5 singleton emoji clusters: (#question, #just-curious), (#lets-discuss,

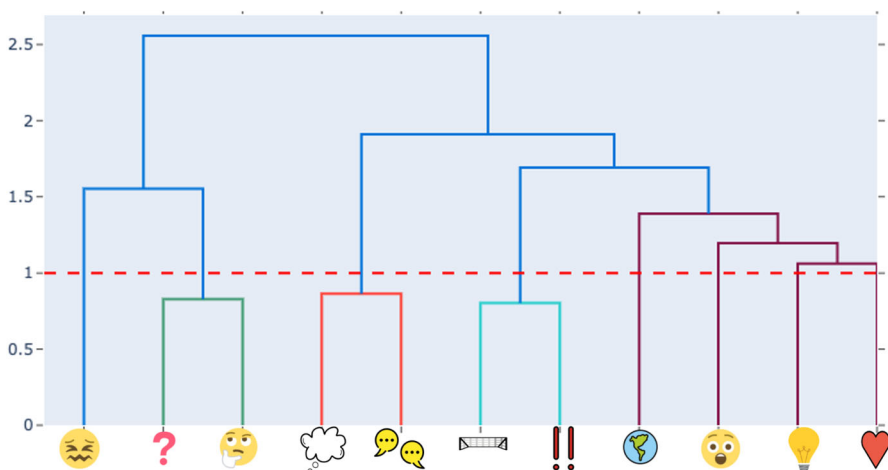


Fig. 8 Hierarchical clustering of the emoji based on the model's prediction. The dashed red line represents the confidence threshold

Table 6 Cluster of emoji-pairs and justification from course instructors

Emoji cluster	Instructor Justification
(#just-curious, #question)	Both emoji invite a response
(#lets-discuss, #i-think)	Both emoji invite follow-up discussion around a topic
(#learning-goal, #important)	Both emoji identify specific reading material to be important

#i-think), (#important, #learning-goal), (#surprised), (#lightbulb-moment), (#lost), (#interesting-topic), and (#real-world-application).

The justifications provided by the course instructors for this division are shown in Table 6. With respect to the cluster containing the emoji (#question, #just-curious), instructors claimed that they invite a response from peers or the instructors, remarking that “They are both requesting responses but with different urgency.” leading instructors to agree to put them in the same category since this cluster captures the reaction they wish to seek in students.



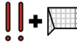





With respect to the cluster containing the emoji (#lets-discuss, #i-think), instructors claimed that “both of these emoji indicate enthusiasm to continue around a topic, either for curiosity or sometimes to clarify [...] they are expressing how they understand it and want clarification or alternate views from peers.”

With respect to the cluster with the emoji (#important, #learning-goal), instructors claimed that both emoji depict students’ perspectives on exam-related content. They claimed that “these emojis are used when students identify parts of the reading material that they believe is important for them to perform well in the course - that may be linked to assessment or the development of knowledge that build towards good performance on assessments.”

Interestingly, the emoji-pair (#surprised, #interesting-topic) exhibited a high correlation on the paragraph level but was deemed sufficiently distinct by instructors to warrant a separate reaction category for each emoji. Instructors claimed that “Both emoji are expressions of “enjoyment” of the text - though as noted slightly different.”

Table 7 shows the performance for the emoji-pair clusters as well as the performance over each one of the single-cluster. There was an improvement of 15% in the average F-1 score for the prediction performance (from 40% to 55.8%) when compared to

Table 7 BERT-FYBIO performance for emoji clusters

Metric								
Precision	0.83	0.67	0.49	0.39	0.23	0.32	0.36	0.27
Recall	0.67	0.57	0.58	0.31	0.65	0.41	0.31	0.49
F-1 Score	0.74	0.61	0.53	0.35	0.34	0.36	0.33	0.34

using a target set containing the original 11 emoji. We note that an improvement in prediction performance was to be expected, given the reduction in the size of the target set. What is interesting is that the two emoji #lets-discuss and #learning-goal, which achieved very low performance when predicted as individual emoji (see Table 4), exhibited significant improvement when clustered together with another emoji (F-Score of 64% and 53% respectively). The cluster with the emoji pair (#question, #just-curious) received the greatest score across all metrics. The two next clusters have gained an improvement in both of the emoji, but the major improvement happened in the emoji #learning-goal and #lets-discuss, which according to Table 4 while these emoji initially performed poorly, the clustering process enabled them to be predicted more precisely.

In terms of predicting the clusters with a single emoji, we manage to improve the prediction score for #lost, #interesting-topic, and #surprised, but we saw a slight reduction in performance in the emoji #lightbulb-moment and #real-world-application.

Generalizing the Model to other Courses

In this section, we evaluated the model's ability to generalize across different subjects. We examined whether the model relies on FYBIO-specific aspects for predictions or if it learns broader, generalizable patterns from the data.

To test this approach, we took 1,722 students' posts from NB which consisted at least of one emoji from 25 courses such as Computer science, Physics, Math, Bio-Informatics, and more. Emoji are only rarely used in these courses, which can be attributed to students not being motivated to use them because they are not required to, as well as a lack of instruction on the use of emoji as a tool in these classes. For our preliminary experimentation with a limited dataset, we opted for the Support Vector Machine (SVM) as our baseline model instead of a language model. The rationale behind this choice was that language models struggled to effectively capture meaningful patterns and signals given the small data size.

Initially, we evaluated our model on the new data set without any additional training. We found that the model performed worse when compared to its performance on FYBIO data, as shown in Table 8. This result indicates that the model had difficulty generalizing from the FYBIO data to new non-biology courses. However, based on the results in Table 8 we can deduce that the model was able to predict some of the examples correctly, showing it is able to exploit the previous attitude signals that it learned through training. This indicates that the model can still perform better than training a new model from scratch.

We decided to follow previous research that suggests additional fine-tuning on new data to account for domain drift (Ma et al., 2019). We used 1,012 posts for training, with 10% set aside for validation, and 259 posts for the test set. After additional fine-tuning, we observed an improvement, as seen in Table 8. This suggests that our model is able to adapt to different domains with a limited amount of training data.

The significance of this result lies in the implication that our model can effectively learn from a small sample size and generalize well to unseen data. This is especially

Table 8 Model performance (F-1 Score) on non-Biology courses. BERT-FYBIO denotes the original model trained on FYBIO Biology course, and W-Training denotes the BERT model with additional fine tuning

Strategy								
SVM	0.6	0.44	0.11	0.32	0	0	0.12	0.12
BERT-FYBIO	0.76	0.58	0.4	0.44	0.33	0.22	0.15	0.2
BERT-FYBIO+ Additional fine tuning	0.77	0.6	0.47	0.42	0.33	0.29	0.11	0.25

crucial when working with educational datasets that often have a limited amount of labeled data available.

Discussion

The premise for predicting students' emoji from posts (which was confirmed in the data analysis) was that the same text would trigger different responses from different students (see Fig. 5). Therefore it makes sense to predict students' reactions from their comments rather than the textbook itself.

Although the BERT classifier was able to outperform the LSTM baseline, its overall performance (See Table 5) was not sufficiently high to allow the deployment of the tool in real classrooms. This may be due the low number of occurrences of individual emoji and the non-uniform distribution over the emoji classes. Another possible reason is that some of the emoji may not have reliably reflect students' attitudes. They may have been generated to meet the quota required for class. This corresponds to documented cases in which students have been shown to game-the-system in educational technologies (Baker et al., 2008). The clustering analysis (see [Classifying Groups of Emoji](#)) allowed us to reduce the number of emoji to a smaller set of meaningful groups, as determined by experts, and to improve the performance of the classifier. Future versions of NB can be improved by reducing the number of emoji to a more meaningful set and to increase performance.

Because posts can include more than a single emoji, we considered using a multi-label prediction approach, also common in social media settings (Felbo et al., 2017; Çöltekin & Rama, 2018; Zhang et al., 2020). However, in the NB domain the number of instances containing more than a single emoji is rather few (19% of all instances, see Fig. 2), whereas 60% of all instances contain a single emoji. Due of this imbalance, we predicted a single emoji using a multi-class approach.

The ability to generalize emoji classification across courses, even to settings where emoji is not available, demonstrates the potential of our approach to providing meaningful pedagogical insights in courses other than FYBIO, and with minimal overhead for instructors. This also provides a potential contribution to the Educational Data Mining field, as educational datasets often have a limited amount of labeled data available and by using such a model, we can apply the emoji mechanism to new domains.

We list several limitations of the classification approach. First, we predicted a single emoji for each post, despite the fact that students were able to tag posts with several emoji. Extending our model to the multi-label approach is not trivial given the scarcity of the data. Second, the forced use of emoji may have generated patterns that did not reliably reflect students' actual attitudes about the material.

Conclusion and Future Work

Our research aimed to explore the use of emoji as a pedagogical tool for instructors to better understand students' reactions and engagement in online educational platforms. To achieve this, we analyzed the various emoji that students used in response to different course topics and grouped them using hierarchical clustering. With the help of course instructors, we identified the most pedagogically relevant emoji clusters and developed a BERT-based language model that effectively classifies student posts into the appropriate cluster.

This study further explores how instructors could utilize emoji as a pedagogical tool to improve course design and to guide students. In order to test the usability of this system in real-world scenarios, we designed a user study that visualizes the usage of emoji directly on the course material. This allows instructors to gain a more detailed understanding of student engagement and interaction with the course content. Moreover, the heat map allows instructors to quickly identify the different types of reactions that students are expressing, and understand how they are distributed throughout the text. By providing instructors with this level of insight, our system can help them to adjust their teaching methods and create a more effective learning environment for their students.

In future work, we intend to extend our work in both the computational approach and the user study approach. In the computational approach, we intend to collect more students' use of emoji from various of the different domains in order to improve the models' performance to a certain level where we will embed this model in the platform itself to further improve. Moreover, we intend to future investigate the benefit of our classifier on online courses which lacks emoji in their forums, and investigate whether the instructors of these courses could make benefits from the classifier.

In the user study approach, we wish to improve the distribution by having more aggregate information for each section, such as the number of comments, numbers for each emoji, and not just the density. We believe that combining these two future works could potentially improve the task of monitoring hybrid classrooms and help teachers to understand the reactions of students when interacting with the course content even without the need of face to face interaction. Lastly, we would like to perform live AB testing in order to better display the benefits of using the emoji heat map compared to the regular heat map.

Acknowledgements This material is based upon work supported by the National Science Foundation under Grant No. 1915724.

Funding Open access funding provided by Ben-Gurion University.

Declarations

Financial Interests The authors have no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Agrawal, A., Venkatraman, J., Leonard, S., & Paepcke, A. (2015). Youedu: Addressing confusion in mooc discussion forums by recommending instructional video clips
- Baker, R. S., d., Corbett, A. T., Roll, I., & Koedinger, K. R. (2008). Developing a generalizable detector of when students game the system. *User Modeling and User-Adapted Interaction*, 18, 287–314.
- Baker, R. S., D’Mello, S. K., Rodrigo, M. M. T., & Graesser, A. C. (2010). Better to be frustrated than bored: The incidence, persistence, and impact of learners’ cognitive-affective states during interactions with three different computer-based learning environments. *International Journal of Human-Computer Studies*, 68(4), 223–241.
- Barbieri, F., Camacho-Collados, J., Ronzano, F., Espinosa-Anke, L., Ballesteros, M., Basile, V., Patti, V., & Saggion, H. (2018). SemEval 2018 task 2: Multilingual emoji prediction. In: Proceedings of The 12th International Workshop on Semantic Evaluation, pp. 24–33. Association for Computational Linguistics, New Orleans, Louisiana. <https://doi.org/10.18653/v1/S18-1003>. <https://aclanthology.org/S18-1003>
- Baziotis, C., Athanasiou, N., Paraskevopoulos, G., Ellinas, N., Kolovou, A., & Potamianos, A. (2018). NTUA-SLP at semeval-2018 task 2: Predicting emojis using rnns with context-aware attention. [arXiv:1804.06657](https://arxiv.org/abs/1804.06657)
- Baziotis, C., Athanasiou, N., Paraskevopoulos, G., Ellinas, N., Kolovou, A., & Potamianos, A. (2018). Ntua-slp at semeval-2018 task 2: Predicting emojis using rnns with context-aware attention. [arXiv:1804.06657](https://arxiv.org/abs/1804.06657)
- Benitez, C., Quinones, A., Gonzalez, P., Ochoa, C., & Vargas, A. (2020). The impact of online annotation tools on students’ academic performance in a distance university program. *Turkish Online Journal of Distance Education*, 21(2), 167–177.
- Bennett, L., & Folley, S. (2021). In: Sahin M., Ifenthaler D. (eds.) Students’ emotional reactions to social comparison via a learner dashboard, pp. 233–250
- Binali, H. H., Wu, C., Potdar, V. (2009) A new significant area: Emotion detection in e-learning using opinion mining techniques. In: 2009 3rd IEEE International Conference on Digital Ecosystems and Technologies, pp. 259–264. <https://doi.org/10.1109/DEST.2009.5276726>
- Blobstein, A., Gal, K., Karger, D., Facciotti, M., Kim, H., Almahmoud, J., & Sripathi, K. (2022). #lets-discuss: Analyzing student affect in course forums using emoji. In: Proceedings of the 15th International Conference on Educational Data Mining
- Boekaerts M., & Pekrun, R. (2015). Emotions and emotion regulation in academic settings. *Handbook of educational psychology*, 90–104
- Boekaerts, M. (1988). Motivated learning: Bias in appraisals. *International Journal Of Educational Research*, 12(3), 267–280.
- Boekaerts, M., Zeidner, M., & Pintrich, P. R. (1999). *Handbook of Self-regulation*. USA: Elsevier.
- Chaplot, D. S., Rhim, E., & Kim, J. (2015). Predicting student attrition in moocs using sentiment analysis and neural networks. *CEUR Workshop Proceedings*, 1432, 7–12.
- Chi, M. T., & Wylie, R. (2014). The icap framework: Linking cognitive engagement to active learning outcomes. *Educational Psychologist*, 49(4), 219–243.

- Chollet, F., et al. (2015). Keras. <https://keras.io>
- Çöltekin, Ç., & Rama, T. (2018) Tübingen-oslo at semeval-2018 task 2: Svms perform better than rnns in emoji prediction. In: Proceedings of The 12th International Workshop on Semantic Evaluation, pp. 34–38
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. [arXiv:1810.04805](https://arxiv.org/abs/1810.04805)
- D'Mello, S., & Graesser, A. (2012). Dynamics of affective states during complex learning. *Learning and Instruction*, 22(2), 145–157.
- Doiron, J. (2018). Emojis: Visual communication in higher education. *PUPIL: International Journal of Teaching, Education and Learning*, 2(2), 1–11
- Dunlap, J. C., Bose, D., Lowenthal, P. R., York, C. S., Atkinson, M., & Murtagh, J. (2016). *What sunshine is to flowers: A literature review on the use of emoticons to support online learning* (pp. 163–182). Technology, Design, And Learning: Emotions.
- Estrada, M. L. B., Cabada, R. Z., Bustillos, R. O., & Graff, M. (2020). Opinion mining and emotion recognition applied to learning environments. *Expert Systems with Applications*, 150, 113265.
- Farzan, R., & Brusilovsky, P. (2008). Annotated: A social navigation and annotation service for web-based educational resources. *New Review of Hypermedia and Multimedia*, 14(1), 3–32.
- Felbo, B., Mislove, A., Søgaard, A., Rahwan, I., & Lehmann, S. (2017). Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. [arXiv:1708.00524](https://arxiv.org/abs/1708.00524)
- Geller, S. A., Hoernle, N., Gal, K., Segal, A., Zhang, A. X., Karger, D., Facciotti, M. T., & Igo, M. (2020). # confused and beyond: detecting confusion in course forums using students' hashtags. In: Proceedings of the Tenth International Conference on Learning Analytics & Knowledge, pp. 589–594
- Jena, R. (2019). Sentiment mining in a collaborative learning environment: capitalising on big data. *Behaviour & Information Technology*, 38(9), 986–1001.
- Kastrati, Z., Dalipi, F., Imran, A. S., Pireva Nuci, K., & Wani, M. A. (2021). Sentiment analysis of students' feedback with nlp and deep learning: A systematic mapping study. *Applied Sciences*, 11(9), 3986.
- Kerslake, L., & Wegerif, R. (2017). The semiotics of emoji: the rise of visual language in the age of the internet (book review). *Media and Communication*, 5(4), 75–78.
- Kim, H. G. (2022). Utilization and regulation of carbon sources and their metabolism in the halophilic archaeon halogeometricum borinquense. Phd thesis, US Davis. Available at <https://escholarship.org/uc/item/6c95q6bs#author>
- Kort, B., Reilly, R., & Picard, R. W. (2001). An affective model of interplay between emotions and learning: Reengineering educational pedagogy-building a learning companion. In: Proceedings IEEE International Conference on Advanced Learning Technologies, pp. 43–46. IEEE
- Kort, B., Reilly, R., & Picard, R. W. (2001). An affective model of interplay between emotions and learning: Reengineering educational pedagogy-building a learning companion. In: Proceedings IEEE International Conference on Advanced Learning Technologies, pp. 43–46. IEEE
- Kralj Novak, P., Smailović, J., Sluban, B., & Mozetič, I. (2015). Sentiment of emojis. *PLoS ONE*, 10(12), 0144296.
- Kuosa, K., Distant, D., Tervakari, A., Cerulo, L., Fernández, A., Koro, J., & Kailanto, M. (2016). Interactive visualization tools to improve learning and teaching in online learning environments. *International Journal Of Distance Education Technologies (IJDET)*, 14(1), 1–21.
- Lund, A. M. (2001). Measuring usability with the use questionnaire12. *Usability Interface*, 8(2), 3–6.
- Ma, X., Xu, P., Wang, Z., Nallapati, R., & Xiang, B. (2019). Domain adaptation with bert-based domain classification and data selection. In: Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019), pp. 76–83
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. [arXiv:1301.3781](https://arxiv.org/abs/1301.3781)
- Munday, D. (2021). Teaching and learning post pandemic. Research-publishing. net
- Rehurek, R., & Sojka, P. (2011). Gensim–python framework for vector space modelling. NLP Centre, Faculty of Informatics, Masaryk University, Brno, Czech Republic 3(2)
- Sahin, M., & Ifenthaler, D. (2021). Visualizations and dashboards for learning analytics: A systematic literature review. *Visualizations and dashboards for learning analytics*, 3–22
- Santamaría-Bonfil, G., & López O. G. T. (2019). Emoji as a proxy of emotional communication
- Sun, Y., & Gao, F. (2017). Comparing the use of a social annotation tool and a threaded discussion forum to support online discussions. *The Internet and Higher Education*, 32, 72–79.

- Sun, C., Hwang, G.-J., Yin, Z., Wang, Z., & Wang, Z. (2023). Trends and issues of social annotation in education: A systematic review from 2000 to 2020. *Journal of Computer Assisted Learning*, 39(2), 329–350.
- Sylwester, R. (1994). How emotions affect learning. *Educational Leadership*, 52(2), 60–65.
- Sylwester, R. (1994). *How emotions affect learning*. *Educational leadership*, 52(2), 60–65.
- Vandergriff, I. (2013). Emotive communication online: A contextual analysis of computer-mediated communication (cmc) cues. *Journal of Pragmatics*, 51, 1–12.
- Watters, D. J., & Watters, J. J. (2006). Student understanding of ph:“i don’t know what the log actually is, i only know where the button is on my calculator”. *Biochemistry and Molecular Biology Education*, 34(4), 278–284.
- Winchell, A., Lan, A., & Mozer, M. (2020). Highlights as an early predictor of student comprehension and interests. *Cognitive Science*, 44(11), 12901.
- Yang, D., Wen, M., Howley, I., Kraut, R., & Rose, C. (2015). Exploring the effect of confusion in discussion forums of massive open online courses. In: Proceedings of the Second (2015) ACM Conference on Learning@ Scale, pp. 121–130
- Yogev, E., Gal, K., Karger, D., Facciotti, M. T., & Igo, M. (2018). Classifying and visualizing students’ cognitive engagement in course readings. In: Proceedings of the Fifth Annual ACM Conference on Learning at Scale, pp. 1–10
- Zhang, A. X., Igo, M., Facciotti, M., & Karger, D. (2017). Using student annotated hashtags and emojis to collect nuanced affective states. In: Proceedings of the Fourth (2017) ACM Conference on Learning@ Scale, pp. 319–322
- Zhang, L., Zhou, Y., Erekhinskaya, T., & Moldovan, D. (2020). Emoji prediction: A transfer learning approach. In K. Arai, S. Kapoor, & R. Bhatia (Eds.), *Advances in Information and Communication* (pp. 864–872). Cham: Springer.
- Zyto, S., Karger, D., Ackerman, M., & Mahajan, S. (2012). Successful classroom deployment of a social document annotation system. In: Proceedings of the Sigchi Conference on Human Factors in Computing Systems, pp. 1883–1892

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Ariel Blobstein¹  · Marc T. Facciotti^{3,4} · Michele Igo⁵ · David Karger⁶ · Prerna Ravi⁶ · Kamali Sripathi⁴ · Kobi Gal^{1,2}

Marc T. Facciotti
mtfacciotti@ucdavis.edu

Michele Igo
mmigo@ucdavis.edu

David Karger
karger@mit.edu

Prerna Ravi
premar@mit.edu

Kamali Sripathi
ksripathi@ucdavis.edu

Kobi Gal
kobig@bgu.ac.il

- ¹ Software and Information Systems Engineering, Ben-Gurion University, Be'er Sheva, Israel
- ² School of Informatics, University of Edinburgh, Edinburgh, United Kingdom
- ³ Biomedical Engineering, University of California-Davis, Davis, CA, USA
- ⁴ Genome Center, University of California-Davis, Davis, CA, USA
- ⁵ Microbiology and Molecular Genetics, University of California-Davis, Davis, CA, USA
- ⁶ Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, US, USA