29

Computer-assisted activations during physical classroom teaching

Daniel Hershcovich

Department of Comptuer Science
University of Copenhagen

Introduction

The field of Natural Language Processing (NLP) in Computer Science is growing rapidly. Educators in NLP struggle to keep up, needing to make decisions about what to teach and how with every iteration of a course, sometimes even as a course is being offered. The fast-paced nature of NLP brings unique challenges for curriculum design. We also have an increasing number of students interested in NLP, bringing with them a wide range of backgrounds and experiences.

I teach a course on NLP at the University of Copenhagen. It is a Master-level course given in English, amounting to 7.5 ECTS and stretching over one block. The usual attendance is 30-40 students. In 2021, about a third of the students came from a Social Science background and had limited Computer Science experience. This proved challenging for catering for the needs of all students, while maintaining the goal of targeting the indented learning outcomes without boring the Computer Science students and without leaving the Social Science students behind. To facilitate this, I chose to activate the students in class with exercises that are focused in scope but aim for a high level of engagement with the material, with the goal that each student should gain what they can from the exercise according to their background.

Instead of traditional student activations, I chose to use online tools that have become popular during the massive turn to online teaching during the COVID-19 pandemic, relying on the students bringing their laptops or smart devices to class, which turned out to be a valid assumption.

Related Work on Student Activation

The teaching paradigm has been moving from a teacher-centered view to a more student-centered perspective (Kaymakamoglu, 2018), meaning that instead of focusing on the role of the teacher, the focus is increasingly on what the student should do, that is, process the material through deliberate practice, collaboration, and active reflection. To effectively support this process, teaching is planned and conducted with the student's disposition in mind, considering their prior knowledge, expectations, study skills, and other conditions. With the proper planning, design, and implementation of the course, active learning can then be achieved (Cook & Babon, 2017) when students are going into lectures and tutorials prepared to engage in the learning process, and they are not just passively trying to absorb information. Active learning encourages active cognitive processing of information, and the concept is not a new one (Saini et al., 2021).

Many pedagogical theories and frameworks have been developed to facilitate effective teaching covering different aspects of teaching (Chi & Wylie, 2014; Cook & Babon, 2017; Kandlbinder, 2014). However, with the advancement of technology and globalization, the traditional pedagogical models evolved to make distance learning possible. Students can sit anywhere and learn online through the internet and connect with other students in the physical classroom or online.

Flipped Classroom (Brame, 2013; Lage et al., 2000) is a teaching approach in which students get first exposure to the material of a lecture outside of class, and the in-class activities involve applying the learned content. Provision of lecture material can be done in several ways, such as reading material, video lectures as slideshows, podcasts, and so on. Students are expected to do the homework, and most of the in-class activities fully depend on that. Because of that, the teacher should make sure that students do their homework, because even though we name it in various ways, watching lectures or reading articles is still homework (Nielsen, 2012), and the challenge of making students accomplish with that is still there. A possible idea for making sure that students do the reading homework could be to ask them to do a quiz or reward them somehow. Apart from challenges regarding students, we must not forget that all the activities, homework, readings, and so on must be retrieved, selected or produced. Furthermore, as the content covered in class is more complex, the teacher will have to be more prepared. All this results in a larger workload in the preparation of the class and its activities (Agirrezabal, 2021).

Intentions

The format of the NLP course while I have been teaching it (two years, since 2019) consists of a weekly two-hour lecture, as well as a weekly two-hour TA session (exercise). Each lecture introduces a new topic, and the TAs revisit that topic and assist students with technical issues and questions about the hand-in assignments.

Both in 2019 (when I was giving in-class physical lectures) and in 2020 (when I was giving online lectures on Zoom, which were recorded and posted on Absalon subsequently), I felt like student engagement was very low and I was talking to blank faces (in 2019) or black screens (in 2020). As a result, I often felt that I had no grasp of the students' understanding and that I was simply going through the material with no indication of whether it was absorbed. In the course evaluation, students had no strong opinion on whether I communicated the content in a clear and precise manner; were somewhat ambivalent about whether I showed interest in the learning outcomes; and indicated that I was good at expressing myself in the language of instruction, but often went too quickly over complicated topics without giving enough explanations and examples.

I wanted to change this situation this year, introducing more interaction into the lectures themselves to both increase student engagement and to gauge their understanding of the content to better organize the lectures. This should already be easier than, for example, a course where all lectures are pre-recorded, since the overall format can stay the same this way. Of course, the main goal of the lectures is still to teach the content of the course to the students, while the TA sessions complement the lectures by providing an opportunity for more interaction and letting the students express themselves. That said, I believe the lectures should be less authoritative and more dialogue-based to a high degree.

While I acknowledge the importance of lectures, as opposed to reading material, for communicating content to students in a synchronous and colloquial manner, I intended to shift some of the burden of teaching from the lectures to the reading material provided to students before each lecture. Additionally, the students could watch the recorded lectures from the 2020 iteration of the course. This way I could implement a "flipped classroom" style of teaching. This would free up the time during class to focus on the points that were difficult for the students (perhaps by giving short lectures on them) and to activate them.

To activate the students, I intended to use in-class exercises with Gather.Town, Padlet, Socrative, and Google Docs as the means of interaction with the students. I intended to open each lecture with immediate activation relating the content to students' personal experience with some application domain, e.g., sharing their own business reviews for sentiment analysis, to create a dialogue among the students about possible solutions before presenting them to them, or asking students what languages they speak and simply listing them before talking about cross-lingual NLP. I intended to implement the TDS framework (Theory of Didactic Situations; Brousseau, 1997) with the activation phase being assisted by digital tools, as these had been vastly improved and made widely available when all teaching became online in 2020 due to the COVID-19 pandemic and had been shown to be beneficial for online teaching (Saini et al., 2021).

My research question in this project is whether digital tools for student activation can also benefit in-class teaching.

Format

I decided to experiment with digitally assisted activations during my three lectures in the course, respectively on (1) Information Extraction and Question Answering, (2) Machine Translation and Cross-Lingual Transfer Learning, and (3) Dependency Parsing.

I started each lecture by asking the students if they had read the reading material and watched the pre-recorded lectures, which they answered by filling in a Google Form.

In the Information Extraction and Question Answering lecture, I gave the students a 20-minute Padlet exercise at the end about Reading Comprehension vs. Knowledge-based Question answering (Figure 1). Padlet is a flexible digital tool, which in my case served as an online canvas where students could type in answers to the different "question cards" by commenting on them. The students performed the exercise individually, answering as many questions as they could and reflecting on the method they used for answering the questions, which used online resources as well. Subsequently, I spent the end of the lecture (20 minutes more) going through the answers and reflecting on the advantages and disadvantages of each of the methods the students used.

In the Machine Translation and Cross-lingual Transfer Learning lecture, I started with a 2-minute Mentimeter questionnaire about which languages

Figure 1: Padlet exercise on Question Answering.

the students speak. This served to "break the ice", relating to the students' personal experience, but also to learn more about the students, since the subsequent exercise on Machine Translation uses examples in multiple languages, which the students should be able to understand. I considered using Socrative for this questionnaire, but decided to use Mentimenter since it presents the responses as a world cloud, which was appropriate for my needs.

In the responses to the questionnaire (Figure 2), I confirmed that the students are multilingual and that the distribution of languages I chose for the Machine Translation exercise is appropriate.

Subsequently, after an introductory lecture to Machine Translation, I allocated 20 minutes for a Google Docs exercise on Machine Translation (Figure 1), which I prepared in advance. In it, students needed to determine failures of hypothetical Machine Translation systems as reflected in their outputs in multiple languages. The exercises was completed in groups of 2-3 students based on the seating arrangement. I dedicated 20 more minutes to going over the answers to this exercise, which served to teach much of the material I would usually teach in a lecture.

Figure 2: Student responses to the Mentimeter questionnaire about which languages the students speak.

After concluding with institutionalization about Machine Translation and presenting an overview of Cross-Lingual Transfer Learning, I finished this part of the lecture with a follow-up Google Form about Cross-Lingual Transfer Learning for Machine Translation (Figure 2), which immediately related to the Google Docs exercise from earlier in the lesson.

Below there are examples of machine translation system outputs.

1. Make a group of 2-3 people
2. Choose one unclaimed example
3. Write your names under it (before someone else does!)
4. Answer the questions under your names (5-10 lines in total):
   a. What is the correct translation? (Try your best!)
   b. Where is the system on the Vauquois triangle [1-10]? Why?
   c. Which "component" did not work properly? Explain.

The Vauquois triangle:

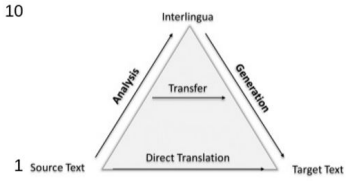

Spanish to English
Source: Maria no daba una bofetada a la bruja verde
Output: Mary not gave a slap to the witch green

Figure 3: Google Docs exercise on Machine Translation.

In the Dependency Parsing lecture, I started with an overview lecture of Dependency Syntax, and continued with a 20-minute Draw.Chat exercise on it (Figure 5). Draw.Chat is a free online tool for sharing a whiteboard with other users, which proved useful for this exercise, where the students worked in groups again, annotating dependency trees of given sentences in multiple languages. I spent another 20 minutes going through the trees that the students drew, and using the opportunity to teach important distinctions in Dependency Syntax. Subsequently, after a lecture about Transition-based Parsing (Figure 6), I concluded with a 5-minute Quizalize exercise on it, where students were presented with a series of parser states and had to apply their knowledge from the lecture to determine what an arc-hybrid transition-based parser would need to select as the next transition each time. Quizalize is a free online tool for "gamification" of multiple-choice quizzes, giving points and optionally even ranking participants based on their scores.

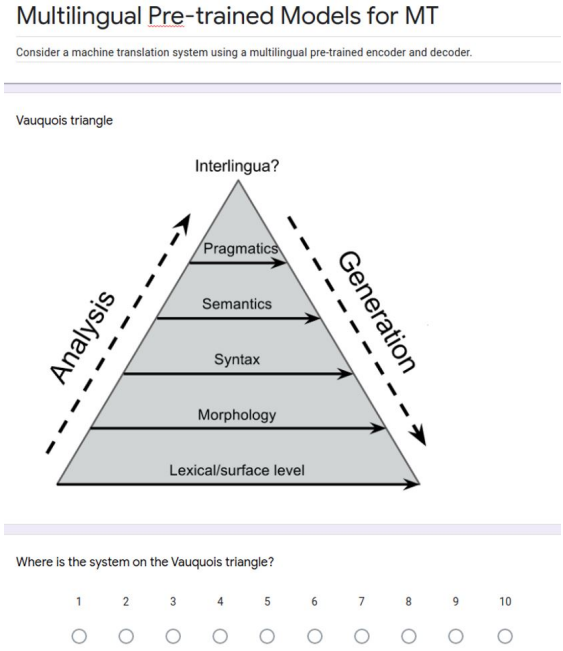

Figure 4: Google Forms quiz on Cross-Lingual Transfer Learning for Machine Translation.

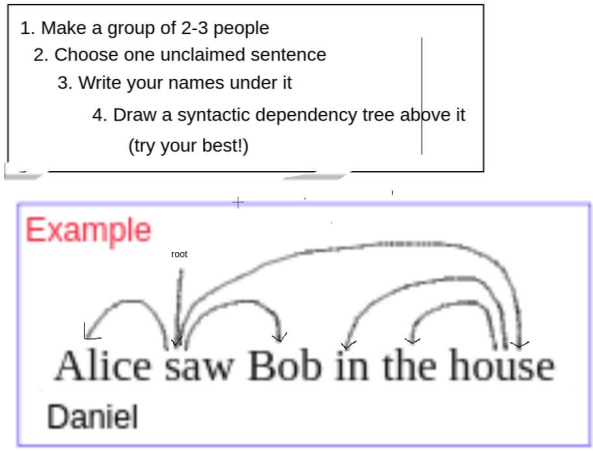

Figure 5: Draw.Chat exercise on Dependency Syntax.

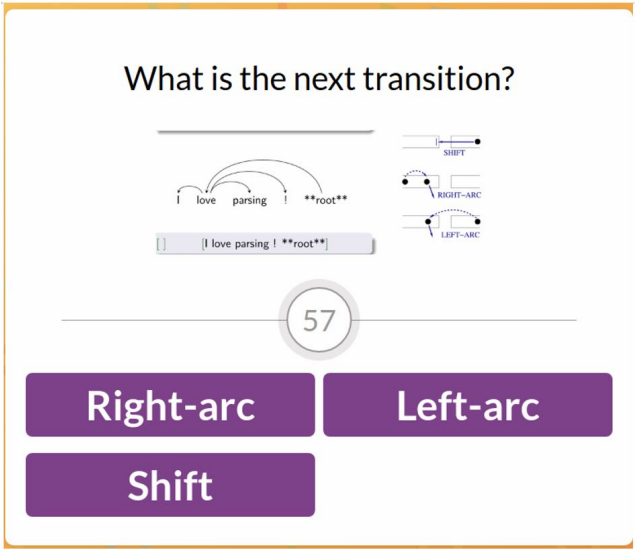

Figure 6: Quizalize exercise on Transition-based Parsing.

At the end of the last lecture, I asked the students to fill in an anonymous Google Form survey on the benefit of in-class activations, asking them to rate how much it contributed to their learning outcomes compared to a hypothetical "normal" lecture on the same subject.

Results

In the responses to the anonymous survey on the benefit of in-class activations (Figure 7), I observed that some exercises were more beneficial than others. Specifically, the Quizalize exercise turned out to be the most beneficial, although it was also the shortest in terms of time. Based on the students' free text responses to this survey, as well as to the teaching evaluation survey (see Appendix), they appreciated the material taught in the lectures more than the more indirect teaching through the exercises, although they found them engaging and the format and tools used were appropriate.

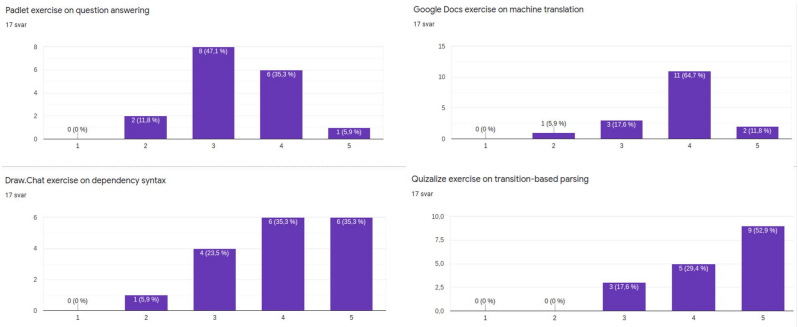

Figure 7: Student responses to the anonymous survey on the benefit of in-class activations.

Discussion

Based on the experience in my teaching, I realized it is important to set expectations with the students before class: are they expected to watch the online recorded lectures that are available from last year? Are they expected to

read the reading material thoroughly? Furthermore, it is important to check with the students at the beginning of class if they have watched the lectures and/or read the material. To my surprise, I found that although I tried to set the expectations appropriately, most students had no more than a brief look at the material. Accordingly, I had to make the exercises more superficial than intended, to avoid having to teach the material comprehensively in a lecture format (which is what I was trying to avoid).

The decision to use Padlet as a tool for in-class exercise had both positive and negative aspects, as there may not be enough time to go through all answers, even if work is done in groups. I had to either prepare efficient validation methodology (e.g., grouping the examples by "type of error" and then simply categorizing them) or announce that I will process them during the break (and thus will not be available for questions etc.). I chose to use the former approach, but students still found the validation phase tedious.

In the more free-form exercise using Google Docs for group work, it was important to clarify how elaborate students' answers should be – describing *how* the solutions would work or simply listing *what* methods they would rely on? Additionally, I had to plan carefully when to go over the answers and when to provide feedback.

Conclusion

Overall, I was satisfied with the students' engagement, as expressed by asking questions during the lecture and participating in exercises. I was also satisfied with the timing, as I did not have to skip almost anything. However, students did not read the reading material almost at all, so the exercises had to be superficial. Nevertheless, students were overall satisfied with the exercises and thought that they contributed to their learning outcomes more than a "normal" lecture would. While there was much more interaction in the Machine Translation lecture than the Dependency Parsing one, (at least some) students were overall more satisfied with the latter, as it covered more material due to more lecturing and less dialog. A good balance between the two must be found to provide the best learning experience.

If I were to repeat this experiment, I would focus on what reading material is necessary, and cut down on non-essential material. This would make sure that the students come prepared to class and can engage with the exercises fully. I would also add a quiz before class that must be solved based on reading material, to clarify to the students what level of preparedness is

expected. This will also serve for assessing their knowledge and interest. This will also be an opportunity to source questions from students about the content, which they are interested in hearing further explanations for during class. To be able to complete the quiz, the students will have to go over the main points in the reading material, providing a kind of motivation that is lacking when simply asking them to read it.

While many of these conclusions could apply to non-digital in-class activations as well, the ultimate conclusion I reached was that indeed, digital activations are beneficial in physical teaching as well, if done right.

References

Agirrezabal, M. (2021). The flipped classroom model for teaching conditional random fields in an nlp course. *Proceedings of the Fifth Workshop on Teaching NLP*, 80–86.

Brame, C. (2013). Flipping the classroom [retrieved last 2021-03-15.].

Brousseau, G. (1997). *Theory of didactical situations in mathematics. didactique des mathématiques 1970–1990*. Kluwer Academic Publishers.

Chi, M. T., & Wylie, R. (2014). The icap framework: Linking cognitive engagement to active learning outcomes. *Educational psychologist*, *49*(4), 219–243.

Cook, B. R., & Babon, A. (2017). Active learning through online quizzes: Better learning and less (busy) work. *Journal of Geography in Higher Education*, *41*(1), 24–38.

Kandlbinder, P. (2014). Constructive alignment in university teaching. *HERDSA News*, *36*(3), 5–6.

Kaymakamoglu, S. E. (2018). Teachers' beliefs, perceived practice and actual classroom practice in relation to traditional (teacher-centered) and constructivist (learner-centered) teaching (note 1). *Journal of Education and Learning*, *7*(1), 29–37.

Lage, M. J., Platt, G. J., & Treglia, M. (2000). Inverting the classroom: A gateway to creating an inclusive learning environment. *The journal of economic education*, *31*(1), 30–43.

Nielsen, L. (2012). Five reasons i'm not flipping over the flipped classroom. *Technology and Learning*, *32*(10), 46–46.

Saini, R., Kovács, G., Faridghasemnia, M., Mokayed, H., Adewumi, O., Alonso, P., & Liwicki, M. (2021). Pedagogical principles in the online teaching of text mining: A retrospection. *Proceedings of the Fifth Workshop on Teaching NLP*, 1–12.

A

Free-text responses to the anonymous survey on the benefit of in-class activations

Padlet exercise on question answering

- A bit too many examples
- I think it may have been more usefull to look more at the query language
- A bit long
- Although I enjoyed it very much, we probably didn't need to go through all the questions as thoroughly.
- It was a good exercise, but it was too long. We don't need to go through all the answers
- After a few reviews of questions it became repetitive. I found that I learned the most from the comparison between information gathering methods, both when shown in the lecture and as a task done in groups.

Google Docs exercise on machine translation

- A bit long walkthrough afterwards - we got the point
- I found it very useful when learning the different components of a sentence in translation.

Quizalize exercise on transition-based parsing

- good - also took an appropriate time
- The interactive quiz with this was great! Really helped
- This worked very well

General comments

- Doing short exercises helps understand the topic explained during the lecture
- Because the lectures we so intense and information heavy (too much) in the beginning these last lectures has seemed almost to light, and a few of us are nervous that because we have spend a lot of time on thouroughly going through these interactive lectures that we are missing out on important stuff that we should have learned for the multiple choice test.

Free-text responses to the teaching evaluation

• The interactive examples were good and engaging.

• The tempo was a bit slow at times, it was good too do some interactive learning, but it took a lot of time at much, and we only had 2 hours lectures a week, so it felt a bit "weak".

• Some of the exercises you made are a bit trivial. The one about the varqouis triangle did not have to take a whole lecture, a few example would have been enough. The same with the wikepedia articles