Document Title	**Evaluation of developer seminar**
Project Title and acronym	Cyprus Center for Algorithmic Transparency (CyCAT)
H2020-WIDESPREAD-05-2017-Twinning	Grant Agreement number: 810105 — CyCAT
Deliverable No.	D5.4
Work package No.	WP5
Work package title	Designing and evaluating interventions
Authors (Name and Partner Institution)	Styliani Kleanthous (OUC) Maria Kasinidou (OUC)
Contributors **(Name and Partner Institution)**	Kalia Orphanou (OUC) Jahna Otterbacher (OUC)
Reviewers	Monica Lestari Paramita (USFD) Lena Podoletz (UEDIN)
Status **(D: draft; RD: revised draft; F: final)**	F
File Name	D5.4_Evaluation_of_Dev_Seminar_M36
Date	01 November 2021

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 810105.

Draft Versions - History of Document				
Version	**Date**	**Authors / contributors**	**e-mail address**	**Notes / changes**
v1.0	23/09/2021	Styliani Kleanthous (OUC)	styliani.kleanthous@ouc.ac.cy	Initial version
v2.0	23/09/2021	Maria Kasinidou (OUC)	maria.kasinidou@ouc.ac.cy	Circulated to OUC team
v3.0	28/10/2021	Lena Podoletz (UEDIN)	lena.podoletz@ed.ac.uk	Review
v4.0	29/10/2021	Monica L. Paramita (USFD)	m.paramita@sheffield.ac.uk	Review
v5.0	01/11/2021	Styliani Kleanthous (OUC)	styliani.kleanthous@ouc.ac.cy	Revised version
v6.0	01/11/2021	Maria Kasinidou (OUC)	maria.kasinidou@ouc.ac.cy	Revisions done
v7.0	01/11/2021	Styliani Kleanthous (OUC)	styliani.kleanthous@ouc.ac.cy	Final Version

Abstract

This document describes the evaluation of a technically-oriented seminar targeted at developers and interface designers designed and discussed in detail in D5.3 - Materials for developers seminar. It served as input to the developer seminar event described in WP2 and evaluated in two phases: immediately after the seminar completion and a year after the seminars.

Keyword(s):	Developers training, Evaluation, Intervention with computer science students, evaluation

Contents

1. Executive Summary

As described in the CyCAT DoA, the developers seminar has been designed and conducted in collaboration with the Department of Computer Science of the University of Cyprus (UCY), which is recognized as the most comprehensive and prestigious computer science program in Cyprus. Deliverable D5.4 describes the agreed-upon plan for evaluating in multiple ways the technically-oriented seminar on algorithmic transparency described in D5.3 - Materials for developers seminar. In particular, we detail the methodology by which the evaluation was developed, providing an extensive analysis of the results and some useful insights and discussion.

2. Seminar Description

The seminar was conducted in the "Software Engineering" module, which is a mandatory course offered to third-year undergraduate students of the Computer Science (CS) Degree at the CS Department, University of Cyprus (N=50). The seminar was also run in the "Advanced Software Engineering" module, which is an elective course offered to postgraduate students of master degree in CS or Advanced Information Technologies at the same institution (N = 7). Due to the COVID-19 pandemic, both courses were being offered in an online format for Fall 2020 in a three hours slot.

The seminar followed the format described in D5.3 - Materials for developers seminar. In this seminar participants: i) became aware of FATE issues in the development of (algorithmic) process/systems; ii) learned the core FATE concepts related to software development; iii) developed appreciation for the role that developers play in mitigating algorithmic bias and in promoting ethical practices; iv) became aware of techniques for auditing services/modules used in development. The seminar began with asking the students to fill in the pre-seminar questionnaire. Then a lecture-style introduction and basic definitions of the concepts that were going to be discussed during the course were provided. In order to motivate discussions between the students and the moderators, we used examples from real life systems that the students were familiar with (e.g. Google Search Engine, Facebook etc.) and have exhibited behaviour that was not fair or just to some parts of the population. Research results were used to explain to the students the methods and approaches followed for uncovering and mitigating bias in such systems and the main stakeholders who are involved i.e. developers, users. Examples of such approaches include Auditing, Fairness Management and Explainability. Moving on, the students became aware of relevant policies - national and international - that attempt to regulate issues related to algorithms FATE e.g. GDPR, ACM Principles of Algorithmic Transparency and Accountability and National Strategies on those topics.

3. Evaluation Methodology

3.1. Initial Evaluation

In order to understand how the perception of CS students on algorithmic FATE changes, we ran seminars and measured their perceptions before and after the seminar. Responses were anonymised and participants had to provide their consent for their responses to be used for this study. Participants were also informed of the purpose of the study and ethical clearance under the CyCAT project.

Pre-seminar Questionnaire. Participants self-assessed their knowledge on Fairness, Accountability, Transparency and Ethics in algorithmic Decision Making (DM) systems using a 5-point Likert scale (1, Not at all - 5, Very Knowledgeable) and self-reported (Yes/No/Other write-in) whether they have taken "any training/course on Fairness, Accountability and Transparency in algorithmic systems".

Then, since participants are developing their own systems, they were asked whether they would "consider dimensions of fairness in [their] system" and whether they would "need to do [their] work a certain way to make a system (more) fair.". Next, participants were asked whether they would "consider possible solutions for making a system more transparent to the user". Then, we presented the participants with three statements about who should be held accountable in case a system behaves unfairly and asked them to indicate whether they agree with each on a Likert scale: My team would be held accountable; The system would be held accountable; Neither the system nor my team would be held accountable. All Likert scale questions were ranging from 1, Strongly Disagree to 5, Strongly Agree.

Post-seminar Questionnaire. After attending the seminar, participants were asked to assess their knowledge on FATE in algorithmic decision making systems using a 5-point Likert scale (1, Not at all - 5, Very Knowledgeable). Participants were asked whether they would "consider dimensions of fairness in [their] system" and whether they would "need to do [their] work a certain way to make a system (more) fair". Then, they were asked to - in case that a system behaves unfairly - indicate "on which part of the process [they] would focus" from the following options that were explained during the seminars: Input, Output, Algorithm, Training Data, Third Party Constraints, Fairness Constraints, and User.

Participants then were asked whether they would "consider possible solutions for making a system more transparent to the user". Then participants were asked to elaborate on their answer above with free text. Then, we presented participants with the same three statements on accountability as in the pre-seminar questionnaire and asked them to indicate whether they agree with each on a Likert scale and asked them to "explain [their] answers" to the above statements with free text. All Likert scale questions were ranging from 1, Strongly Disagree to 5, Strongly Agree.

Scenarios: Participants were presented with two different scenarios in the pre-questionnaire and two different scenarios in the post-questionnaire, where algorithms made decisions that influenced humans. For these scenarios we selected contexts that our target population might be

familiar with (please see below for more information) and we alternate the scenarios with minor changes in the storyline in pre- and post- questionnaires. Specifically, we used the same context (e.g. car insurance) but different story line for Scenario A in pre-questionnaire and Scenario B in post-questionnaire; respectively Scenario B in pre-questionnaire had the same context (e.g., CV filtering) but different story line to Scenario A in post-questionnaire. The use of corresponding scenarios in the pre- and post- questionnaires aimed to examine whether the perception of the students changes after attending a short seminar on FATE. The following scenarios were used to trigger the participants' judgement on six fairness constructs extracted from the literature on FATE[1].

Scenario A: A car insurance company's premiums are dynamically-priced, based on the driver's personal details and driving behaviour. This scenario was adopted from Binns et al.[2].

Scenario B: A system is used to filter and rank CVs for the hiring manager, in order to assist in shortlisting the best candidates.

For each scenario, participants were asked to rate their agreement in five statements according to Colquitt and Rodell[1] in addition to the dimension of 'Trust' (see S6 below). A 5-point Likert scale, ranging from '1 - Strongly Disagree' to '5 - Strongly Agree', was employed for each of the six statements:

- S1 Agreement: "I agree with the decision"
- S2 Understanding: "I understand the process by which the decision was made"
- S3 Appropriateness of factors: "The factors considered in the decision were appropriate"
- S4 Fair process: "The decision-making process was fair"
- S5 Deserved outcome: "The individual deserved this outcome given their circumstances or behaviour"
- S6 Trust: "I would trust this system's decision more than a human's decision"

Participants: 50 undergraduate and 7 postgraduate students replied to the questionnaires. Twenty-six participants did not answer both questionnaires, thus 31 responses were considered in the analysis. Participation was voluntary and all participants provided us with written, informed consent for their data to be used. 77.4% of our respondents were male, with 77.4% in the age group of 18-24 and the rest were between 25-32. The majority of the participants (80.6%) identified themselves as undergraduates, and 88% of that group were in their third or fourth year of studies.

3.2 Follow-up Evaluation

The follow-up evaluation was performed a year after the seminar took place (Fall 2021) and aimed at re-evaluating the students' attitude and understanding of algorithmic FATE, thus, indirectly assessing the effectiveness of the seminar. The students were asked to reply to a short

[1] Jason A Colquitt and Jessica B Rodell. 2015. Measuring justice and fairness. (2015).

[2] Reuben Binns, Max Van Kleek, Michael Veale, Ulrik Lyngs, Jun Zhao, and Nigel Shadbolt. 2018. "It's Reducing a Human Being to a Percentage": Perceptions of Justice in Algorithmic Decisions. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–14.

(9 items) questionnaire. The questionnaires were distributed via email to all 57 students who attended the seminar and seventeen (N=17) students replied. Participation was voluntary and all participants provided us with written, informed consent for their data to be used.

Specifically students had to select from a 5-point Likert scale (1, strongly disagree - 5, strongly agree) what they believe about the following questions: “After completing the seminar I am able to better understand issues related to fairness in algorithmic systems”; “After the seminar I am able to better understand issues related to transparency in algorithmic systems”; “After the seminar I am able to better understand issues related to accountability in algorithmic systems”; “After the seminar I kept in mind the dimensions of FATE in a system that I had to develop.”; “After the seminar I tried to work in a specific way to create a (more) fair system.”; After the seminar I can perceive / identify / distinguish FATE issues when interacting with information systems (e.g. search engines, social media); “After the seminar I can perceive/detect /distinguish when a system offers transparency in its algorithmic processes.”; “I would recommend this seminar to other IT students.”.

4. Results

In this section we will discuss the results in two parts. Firstly, we will present the results of the initial evaluation using the pre- and post- seminar questionnaires. Then, we will present the results of the follow-up evaluation study.

4.1. Results from Initial Evaluation

Knowledge and Formal Training on FATE. In the pre-seminar questionnaire, we asked participants to self-report whether they have taken any kind of training/course on Fairness, Accountability, Transparency issues in Algorithmic Systems. 12.9% of our participants had taken some kind of training on the above topics, while the majority (77.4%) had not and the rest of the participants answered “Other”. We also asked participants to state their knowledge on the above topics before and after the seminar using a Likert-scale (1, Not at All – 5, Very Knowledgeable). Interestingly, Wilcoxon signed ranked test shows significant differences between the pre-seminar and post-seminar questionnaire replies, with replies prior to the seminar being significantly lower compared to their replies after the seminar, in the above questions Fairness - ($z = -3.947$, $p<0.001$); Transparency - ($z = -4.008$, $p<0.001$); Accountability - ($z = -3.857$, $p<0.001$) respectively. These results show that students felt more knowledgeable on FATE topics after they have attended the seminar.

Perception of Algorithmic Fairness. When asked whether they would consider fairness in their system most of the participants (80.7% in pre-questionnaire, 93.2% in post-seminar questionnaire) responded affirmatively (4 - 5), 16.1% in pre-seminar questionnaire and 6.5% in post-seminar questionnaire seemed undecided (3), and 3.2% in the pre-seminar questionnaire and none in the post-seminar questionnaire indicated that they would not consider fairness (1 - 2). It is important that the percentage of students who appear undecided in the pre-seminar questionnaire moved into options 4 and 5 in the after the seminar indicating that they would consider dimensions of fairness in their systems. When asked whether they would work in a certain way to

make a system (more) fair the majority of the participants (80.6% in pre-seminar questionnaire, 87.1% in post-seminar questionnaire) responded affirmatively (4 - 5), 3.2% in both pre-seminar and post-seminar questionnaire indicated that they would not consider fairness (1 - 2).

Participants were asked to choose the parts of the process they think could possibly cause unfairness in the system and explain their choices. The majority of the participants indicated that the Algorithm and the Training Data (25 out of 31) are most possible to cause unfairness in a system. Most of the participants shared the opinion that "unfairness can be caused due to the fact that we do not have enough data for all cases of the system, the way of operation and classification of the elements by the algorithm may favor some specific cases" (p13). Some participants also discussed that "developers with their own bias can affect the system, the data may not have been chosen to be representative for all sectors, just as developers and users influence the system with their biases" (p16). Other participants specifically mentioned the "Biased dataset of training data" (p4) and that "it is not the way of it's implemented that is responsible but the way of data entry and the way of its training" (p11). On the other hand, some participants discussed that "[t]he system and its developers are responsible for the proper functioning" (p8).

Often participants referred to the Input (18 out of 31) of the system as a possible cause of unfairness. Participants referred to unfairness "[d]ue to incorrect entry, for example with the Microsoft bot, where users were responsible for logging in and learning the model" (p21). 16 responses discussed the User as a possible cause of unfairness. Participant 14 pointed out that "[u]ser's biases often get in the system" (p14), and "[i]f the system learns from the users, then the system may learn based on wrong data that are given from the user causing wrong results" (p4). 14 participants chose Third Party Constraint and Fairness Constraints. Only a few participants explained why they chose these parts. Participant 4 mentioned for the Fairness Constraints that "[t]he operator of a system may have biased perceptions in a specific topic and set the system based on his beliefs". Participant 28 explained for choosing Third Party Constraints that "third parties can with their own views indirectly influence even the writing of the algorithm". Finally, only 11 out of the 31 said that the Output could cause unfairness in a system but none specifically explained their choice.

Perception of Algorithmic Transparency. Participants' view on considering possible solutions for making the system more transparent did not significantly differ between the pre-seminar and the post-seminar questionnaire. The majority of the participants, in the pre-seminar questionnaire (87.1%) and post-seminar questionnaire (83.9%) agreed that (4-5 on the Likert scale) they would consider possible solutions for making the system more transparent, compared to 6.5% for pre-seminar questionnaire and 3.2% post-seminar questionnaire who indicated that they would not (1-2 on the Likert scale). We can also observe here that some students moved to the positive part of the scale after the seminar.

In the free-text explanations, participants were asked to explain ways they would use to make the system more transparent. Participants free-text responses were coded and thematically analysed[3].

[3] David R Thomas. 2006. A general inductive approach for analyzing qualitative evaluation data. American journal of evaluation 27, 2 (2006), 237–246

Two researchers analyzed the participants' free-text responses independently to define emerging categories. We allowed multiple categories per answer.

Six themes emerged (see Table 1). The majority of the participants (13 out of 31) suggested that they would **explain** to the user how **the System/Algorithm** works and other (5 out 31) they would **explain the Output**. Some participants mentioned that the "user must know how the system works" (p4), and that "every user has the right to observe how they interact with the system" (p8). Other participants specifically discussed that they would make the system more transparent by letting the user know how the data are used (p13, p23) and the procedures followed by the system (p14, p15). Two participants though mentioned that the user should not know how the algorithm works (p18, p30). The participants who chose to explain the output to promote transparency often mentioned that the user should know how the system concluded the specific output (p12, p14, p15). Four participants suggested that they would focus on **Training Data**. They discussed that they would "re-examine the training data and the algorithm" (p1) and others briefly mentioned that they would explain to the users how they collected the data used to train the algorithm (p10, p13). One participant stated the use of "more accurate and complete training data" (p20) for training the system. Four participants discussed the development of **Unbiased** algorithms. They mentioned that "[they] would try to limit the unfairness as much as possible" p(2) and "avoiding injustice such as gender, skin color" (p22). Others briefly mentioned they would develop fair, ethical, transparent and without biases systems (p26, p27). The least common theme, **Third Party** received only one response and discussed the need to "check for 3rd party influence" (p1). Ten responses fell under the catch-all other category, which includes thoughtful responses where the participant indicated they would make the system more transparent because this is important but they do not specify how (p8, p11) or did not give any response (p21,p25).

Table 1: Themes emerged from Transparency Strategies question: name, description and frequency

Category	**Description**	#
Explaining the System/Algorithm	*explaining the process followed by the system*	13
Explaining the Output	*explaining the output to the user; why a specific decision was made*	5
Training Data	*the dataset/information used for training the algorithm*	4
Unbiased	*algorithm/outcomes without social biases or discrimination*	4
Third party	*the impact of third parties on the system*	1
Other	*[falls outside of the established themes]*	10

Perception of Algorithmic Accountability. The last part of the study examined the concept of accountability and how the participants perceive it. Participants' view on accountability did not significantly differ between the pre-seminar and the post-seminar questionnaire. Before the seminar, most of the participants (70%) agreed with (4 - 5 on the Likert scale) the statement that "their team" would be held accountable, compared to 41.9% who agreed that "the system" would be held accountable and 16.1% who agreed that "neither the system nor my team" would be held accountable. The majority of the participants after the seminar (87.1%) agreed with (4 - 5 on the Likert scale) the statement that "their team" would be held accountable, compared to 48.9% who agreed that "the system" would be held accountable and 9.7% who agreed that "neither the system nor my team" would be held accountable. Indicating that the seminar had an impact in their perception of algorithmic Accountability.

In the post-seminar questionnaire participants were asked to explain their responses using free-text. In the free-text explanations of their choices, participants remarked that "We implement the system so we are responsible for the system" (p6) and that their team should be held accountable since "[our] team may have made [the] mistake on the algorithm or choose the wrong data set" (p1). Some participants justified their responses that the team would be held accountable with the fact that the system is not autonomous, and instead a human chooses the factors that the system uses to make decisions and the data that they use to train the system. For instance, participant 12 noted that "The system works based on how it is programmed to do and the data which was given to it for training". Participants sharing this opinion felt that the humans that developed the system should be held accountable for the unfairness of the system. "The system cannot be held accountable in any case, if someone is responsible [it] is the development team, unless there was a wrong or malicious use of the system, wherein in this case the user is responsible" (p18). On the other hand, some participants felt that both the team that developed the system and the system itself should take the responsibility: "I believe that both the team and the system itself will be held responsible because the team in part allowed discrimination to occur and the system can also learn in this way from the users who use it" (p26).

Can views on FATE be changed? Quantitative analysis was employed in order to explore whether participants' perception of each individual construct for Scenarios changed after the FATE seminar. To examine whether participants' perception changed we conducted a Wilcoxon Signed Ranks Test. The comparison of the results between Scenario A in pre-seminar questionnaire with the corresponding Scenario B in post-seminar questionnaire, indicate significant statistical differences in the responses of the students for Agreement, Understanding, Fairness of the DM process and Trust. With selections after the seminar being considerably lower compared to prior. This shows that students' perception on those issues changed after they were educated on the FATE concepts. More specifically, more students selected lower scores in the Likert scale for Agreement with the decision of the system after the seminar ($z = -2.511$, $p = 0.012$), as well as Understanding of the process by which the decision was made ($z = -2.941$, $p = 0.003$). Similarly, students' responses on the Fairness of the DM process show that they perceived the decision making as less fair ($z = -2.424$, $p = 0.015$) and their Trust to the system's decision compared to a human also ($z = -2.064$, $p = 0.039$). We did not have any statistical significant differences in the students' responses regarding the other two scenarios.

4.1.1. Summary

Perception on FATE. Consistent with Holstein et al.[4], who looked into developers in the industry, students in our sample selected the Training Data and the Algorithm as the components they are most likely to cause unfairness in a system. They emphasised that potential biases and discrimination can exist in the training data and consequently will be learnt by the system. This shows that students understand the need for creating more diverse data sets to be used for training machine learning algorithms integrated in DM systems. Students also mentioned that the developers of a system might unintentionally be promoting their own biases, indicating that they are aware of the human influence in the process. However, CS degrees in their majority do not provide training that can educate students on how to remain neutral towards the algorithms they are developing. When asked about ways to make the system more transparent, it comes with no surprise that the majority of responses discussed providing explanations to the user. Explaining the system/algorithm was the most preferred strategy, although some responses opted (also) for explaining the output. More research is needed in order to understand how and in what way explanations can be used for providing transparency to the users.

Finally, when participants asked who needs to be held accountable, in case a system they develop behaves unfairly, the majority agreed that their team should be held accountable. This may be an indication that future developers understand their responsibility of delivering "fair behaving algorithms" to their users and the possible consequences in case the system they develop is behaving unfairly to some parts of the population. Since students in our sample indicated that they lack relevant training on these topics, we understand there is a need to provide training and resources that CS students will attend when they need to.

Changing views on FATE. Our finding that CS students lack knowledge on topics related to FATE, builds on previous work[4] and reflects the need for incorporating modules and training courses in the computing-related degrees. Our findings are aligned with previous work[5], which also reported evidence of statistically significant changes in perception and attitudes of students towards algorithmic fairness and transparency just after an hour of lecture and discussion. It is important for CS students – who are likely to develop such systems in the future to ensure they are aware of concepts related to FATE in algorithmic systems. They also need to be aware that the systems they are developing have an impact (positive or negative) to society.

Since we are expecting algorithmic systems to behave in a fair and just manner, we need to educate CS students on algorithmic FATE. They need to be aware of the possible ways that biases can be introduced in a system, ways of auditing their systems prior to release, and ways of making their systems overall more transparent to their users. In addition, CS students and future developers need to develop a sense of responsibility to the users of the systems they are developing and to society in general. CS degrees should be rationalized into incorporating

[4] Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé, Miro Dudik, and Hanna Wallach. 2019. Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need?. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–16.

[5] Emma Pierson. 2017. Demographics and discussion influence views on algorithmic fairness. arXiv:1712.09124 [cs.CY]

algorithmic FATE related courses. Although there are standalone seminars on these topics, courses such as Software Engineering could include modules that will provide students with the necessary knowledge on the above topics.

4.2. Follow-up Evaluation

This part of the evaluation took place approximately a year after the seminar. We wanted to understand whether the seminar had a long term effect on the students' view of algorithmic FATE. Seventeen (N = 17) students replied voluntarily to this questionnaire. 82.3% of the students who responded said that after completing the seminar they are able to better understand issues related to fairness and transparency in algorithmic systems, by selecting options 4 - 5 in the likert scale in the respective questions. Regarding understanding issues related to accountability in algorithmic systems, 76.5% replied positively (4 - 5 in 5 - point Likert scale) indicating that they are able to understand related issues better after the seminar. One participant selected option 2, indicating that they are not able to understand related issues better. Issues related to algorithmic accountability are under research by high profile organizations, governments and researchers, and this is not a straightforward topic for students to understand. However, we can see that the majority of the students gained something valuable from the seminar. It was important for us that students understood and kept in mind the dimensions of FATE in a system that they were going to develop. 82.4% of the participants replied affirmatively (selected options 4 - 5), indicating that they were considering FATE dimensions in the systems they were asked to develop after the seminar, while the majority (76.4%) were trying to work in a specific way in their attempt to create a 'more fair' system.

Equally important is for CS students and future developers to be able to identify and distinguish FATE issues when interacting with information systems (e.g. search engines, social media). 82.3% of the participants replied on the positive side (selected options 4 - 5 in the 5 - point scale) indicating that they feel confident they are able to identify such issues, while 64.7% replied affirmatively indicating they can distinguish when a system offers transparency in its algorithmic processes. Finally, 76.5% would recommend this seminar to other CS students.

5. Seminar Evaluation

We asked participants to evaluate the following using a five point likert scale (5-point Likert scale: 1, Strongly Disagree - 5, Strongly Agree):

- The structure of the seminar is clear and efficient
- The educational content and the material of the seminar are well understood and accessible
- The goals and objectives of the seminar are clear and understandable
- The seminar is potentially useful for all computing students
- The seminar met my expectations

All participants selected options in the high-end of the scale (4-5) in all statements, indicating that the seminar was successfully developed, structured and delivered to their expectations.

6. Conclusion

The primary goal of this deliverable is to provide a holistic evaluation of the seminar provided by CyCAT to Computer Science students. In addition, the evaluation studies examined how Computer Science students perceive algorithmic FATE and whether their views and attitudes towards FATE can change after attending a relevant seminar. In the completion of the seminar, participants felt more knowledgeable on FATE topics; they became more likely to consider elements of fairness in their system and believed the team developing a system should be held accountable in case the system behaves unfairly. Finally, short seminars can make a difference in the attitude of students towards FATE in algorithmic decision making in the short but also in the long term. CyCAT will continue to work towards this direction, educating future developers on algorithmic FATE and creating awareness on the impact that such systems can potentially have in society.