Cloud and edge based data analytics for privacy-preserving multi-modal engagement monitoring in the classroom

Davy Preuveneers · Giuseppe Garofalo · Wouter Joosen

Received: date / Accepted: date

Abstract Learning management systems are service platforms that support the administration and delivery of training programs and educational courses. Pre-recorded, real-time or interactive lectures can be offered in blended, flipped or fully online classrooms. A key challenge with such service platforms is the adequate monitoring of engagement, as it is an early indicator for a student's learning achievements. Indeed, observing the behavior of the audience and keeping the participants engaged is not only a challenge in a face-to-face setting where students and teachers share the same physical learning environment, but definitely when students participate remotely. In this work, we present a hybrid cloud and edge-based service orchestration framework for multi-modal engagement analysis. We implemented and evaluated an edge-based browser solution for the analysis of different behavior modalities with cross-user aggregation through secure multiparty computation. Compared to contemporary online learning systems, the advantages of our hybrid cloud-edge based solution are twofold. It scales up with a growing number of students, and also mitigates privacy concerns in an era where the rise of analytics in online learning raises questions about the responsible use of data.

Keywords data analytics · multi-modal engagement monitoring · privacy · cloud and edge computing · browser.

1 Introduction

Student engagement is a topic that has sparked a fair amount of interest over the past decade within the e-learning research community as well as with higher education institutions. Previous research [6, 29] has shown that students who are engaged with their studies are more likely to be successful. However, engagement is a complex construct. Multiple theories have been proposed in the literature and compared by Kahu [28]. In general, engagement can be understood as a mix of behavioral, cognitive and emotional (or affective) factors. In this work, we propose a multi-modal solution for analyzing behavioral engagement in online and blended (i.e. combination of traditional and online) learning environments. Next to quantitative self-reports by means of student engagement questionnaires, behavioral engagement is the most frequent reported style of engagement. Furthermore, behavior is easier to measure compared to cognitive and emotional styles of engagement.

Maintaining a continuous awareness of the engagement level of students is not only a challenge in a face-to-face classroom setting, but even a bigger one in remote settings in which students participate in interactive online lectures. The research that we present in this paper tackles the following three challenges:

1. First, to capture behavior engagement, the system should process audiovisual, interaction and physiological data of the audience in near-real time. This requires a considerable amount of computational resources and network capacity which may not be readily available within or near the classroom.
2. Second, the effectiveness of different behavioral and interaction modalities to measure the engagement of an individual or a group of students is highly context-dependent. For example, gaze and head pose

Davy Preuveneers, Giuseppe Garofalo and Wouter Joosen
imec-DistriNet, KU Leuven, Celestijnenlaan 200A, B-3001
Heverlee, Belgium
E-mail: { davy.preuveneers, giuseppe.garofalo, wouter.joosen
}@cs.kuleuven.be
<https://distrinet.cs.kuleuven.be>

estimation may not always be a good indication whether students are paying attention as the teacher is walking around or the web camera may be ill-positioned.

3. Last but not least, the continuous tracking and centralized analysis of sensitive personal information may invade the privacy of remote students [36,34]. This remains an issue today¹ and is a growing concern. In a *The Washington Post* article of December 2018, Valerie Strauss [47] highlights why parents and students are protesting against an online learning program backed by Mark Zuckerberg and Facebook.

To address these three concerns, we present a decentralized multi-modal engagement analysis solution that benefits from the bring-your-own-device (BYOD) trend in the classroom by distributing the student engagement analysis towards the client devices (e.g. laptops of students). Not only do the students get insights into their own individual engagement levels, the values are also aggregated over all participants and reported to the teacher in near real-time so that the latter can act accordingly during interactive lectures, e.g. by launching a poll or quiz.

Our edge computing software framework analyses different kinds of student behaviors across different online learning contexts. It is implemented as a JavaScript-based web browser extension to simplify the deployment and upgrade cycle of our engagement analysis software framework across different operating systems and web browsers. Our decentralized framework federates the data analytics, machine and deep learning, not only to easily scale up with a growing number of students, but also to mitigate any privacy concerns due to the continuous student behavior tracking.

The main contributions of this research are the application of well-known and proven enabling technologies for decentralized data analytics, as well as best practices for preserving privacy, and this in a non-trivial educational context and real-world use case. The contributions can be summarized as follows:

1. A multi-modal engagement model that runs in the cloud and at the edge to easily scale with a growing number of participants
2. Privacy enhanced analysis of student engagement using federated behavior data processing and the application of secure multi-party computation
3. Performance and practical feasibility assessment using real-world data collected in the classroom

¹ <https://campustechnology.com/articles/2018/05/02/when-learning-analytics-violate-student-privacy.aspx>

While the distributed deployment of the different enabling technologies resembles a typical client-server architecture, the data can be processed and stored close to the location where it is used (e.g. the client), as well as in resource richer environments (e.g. the server). The deployment strategies – i.e. the decisions of where to process certain data – are mainly driven by privacy and computational complexity constraints. It is this flexibility to shift parts of the computations elsewhere and the ability of the student to decide on which data may be analyzed, that we believe is the main novelty of our hybrid cloud and edge based data analytics framework for engagement monitoring.

This research extends previous work reported in [41], and was carried out in the frame of the LECTURE+ project², which aims to research, design and evaluate a next-generation platform that improves decision support for teachers, room operators and students to enhance learner engagement in face-to-face and remote lectures. For this study, 25 students and various researchers participated in field trials during real-world lectures within the frame of this project during which browser-based engagement data was captured and analyzed. A detailed description of the data will be provided in a later section.

The remainder of this paper is structured as follows. In section 2 we review relevant related works on behavioral engagement monitoring. Our edge-based privacy-preserving multi-modal engagement monitoring solution and its building blocks are presented in section 3. In section 4, we discuss the experimental results with our framework from a performance and privacy perspective, and compare the outcome with the related work. Finally, section 5 summarizes the main results and gives concluding remarks.

2 Related work

As our research focuses on the application of edge computing in an online learning context, we mainly discuss relevant related work on engagement analysis to then investigate the feasibility for a scalable and real-time solution in an edge computing deployment, and implications from a privacy point of view.

2.1 Learning analytics for engagement analysis

The domain of higher education [16] is exploring big data and learning analytics technology to (1) identify

² <https://www.imec-int.com/en/what-we-offer/research-portfolio/lecture>

useful data, (2) transform that data to identify useful patterns and deviations from patterns in real-time for educational decision making. Such learning analytics systems can engage students to create new insights, investigate and compare models on learning factors, and detect and appreciate insights generated by their peers. We refer to the work by Henrie et al. [25] in which the authors review relevant related work on the measuring of student engagement. They argue that consensus is needed for the definition and operationalization of student engagement. While most technology enhanced learning research relies on the self-reporting of engagement, they conclude that physiological and systems data offer an alternative method to measuring engagement, but that more research is necessary to measure their value.

Rather than using quantitative self-reporting with student engagement questionnaires, Thomas et al. [48] relied on facial expressions, head pose and eye gaze using computer vision techniques, as well as machine learning methods to analyze the engagement of students in a class room setting. Their method relies on the analysis of 10 second video clips recorded in a 1920x1080 resolution at 25 frames per second. From these videos, they extracted a 30-dimensional feature vector of which 27 were statistically significant. They tested Supported Vector Machine (SVM) and Logistic Regression machine learning methods to classify the engagement level. They were able to improve the accuracy with up to 10% against a baseline classifier (method which classifies everything to the majority class). Contrary to our work, the evaluation of the individual video clips was carried out offline, rather than in a real-time manner. Also, the high-end video equipment may not be readily available to students.

Engagement in a collaborative learning setting supported by the Slack communication tool was explored by Zhang et al. [53]. They investigated to what extent the openness and visibility of each group member's contribution on Slack can promote students' engagement. Their empirical study shows that mutual trust, social influence and reward valence have positive influence on teamwork engagement. Furthermore, this engagement has positive effects on personal success in terms of learning and work satisfaction. Our solution analyzes interaction events in collaboration hubs as well. However, it is not tied to any specific online communication platform, as long as students interact through a browser.

More recently, Gray et al. [23] evaluated several candidate machine learning methods based on a variety of student attendance and engagement information to produce predictions of student outcome. They demonstrate how students can be identified as early as week 3

with approximately 97% accuracy. Our contribution depends on machine learning methods as well, with deep learning methods for facial expressions being the most sophisticated. Our contribution focuses on the technical challenges in terms of performance to monitor engagement in near-real time. Rather than using the data to predict student outcomes, our goal is to be able to adapt the lecture to retain student engagement.

2.2 Data analytics with edge and fog computing

With the proliferation of Internet of Things devices, edge [46,44] and fog computing [11,10] have emerged as viable paradigms to process many data streams and push the communication, computation, control and storage resources from the cloud to the edge of the network. Cao et al. [12] proposed an edge computing platform to uncover meaningful patterns from real-time transit data streams and generate actionable insights. Their platform supports mobile edge nodes and adopts descriptive analytics for monitoring bus routes.

Similar work was proposed by He et al [24]. They present a multi-tier fog computing platform based on Raspberry Pi devices that consists of ad-hoc fogs and dedicated fogs with opportunistic and dedicated computing resources. The platform offers functionalities for QoS aware admission control, offloading and resource allocation schemes. The application area of the platform is smart city analytics services, and experimental results demonstrate that fogs can largely improve the performance of such services compared to when they are deployed in the cloud.

A detailed discussion of the many platforms presented in the literature is beyond the scope of this work. We refer interested readers to the following application overviews and surveys [5,52,3].

2.3 Data analytics and privacy

With the collection and processing of large amounts of data come severe security and privacy challenges. Roman et al. [43] analyze the security threats, challenges, and mechanisms inherent to the edge computing paradigm.

To address privacy concerns with data analytics, various theoretical frameworks have been proposed – including k-anonymity, l-diversity, t-closeness [33], and differential privacy [19] – to mitigate the leakage of potentially sensitive information during data analytics. In [38], Mohan et al. presented and evaluated GUPT, a system for privacy preserving data analysis for applications not designed with privacy in mind [15]. GUPT

enables differential privacy while minimizing the loss in output accuracy, which makes it suitable for a wide variety of data analysis programs while providing both utility and privacy. More recently, Abadi et al. [2] illustrated that it is feasible to train deep neural networks with differential privacy, making sure that the learned models do not expose private information from the datasets. Their experimental results show that it is feasible with an acceptable overhead in software complexity, training efficiency, and model quality.

Orthogonal to these approaches are the privacy enhancing technologies (PETs) that exploit cryptographic protocols – such as secure multi-party computation [40, 26, 14] or homomorphic encryption [18] – to gain insights without having access to the complete datasets or to outsource data analytics to the cloud. An illustration of the feasibility of secure multi-party computation is the work by Ma et al. [35]. They presented a framework for privacy-preserving multi-party deep learning in cloud computing. The training data is shared across many parties, and the privacy of the local dataset as well as the learned model are protected against the cloud server. An example of the second approach is described in [20], in which El-Yahyaoui et al. leverage fully homomorphic encryption to search over encrypted data in a cloud context. The following recent surveys [8, 4, 45] provide more background on these cryptographic protocols and their applications.

2.4 Bridging the gap

The gap that we aim to address with this work is not a further improvement of the accuracy for one of the individual classification methods, nor designing novel cryptographic protocols, but rather providing an integrated framework that offers the same multi-modal engagement monitoring capabilities as in the related work in a unified manner. Our goal is to offer a practical solution that is easy to deploy for which the feasibility of near-real time analysis is demonstrated in a realistic online or blended learning environment, while at the same time addressing performance, scalability and privacy challenges that were not considered as key concerns in the related work.

3 An edge-based framework for privacy-preserving engagement analysis of multi-modal behaviors

This section first elaborates on a motivating scenario from which the requirements regarding multi-modal engagement analytics are elicited. We then continue with

the identification and implementation of the individual building blocks of our edge-based framework for privacy-preserving engagement analysis.

3.1 A blended learning motivating scenario

As motivating use case, we consider a blended learning [22] scenario as a combination of face-to-face education with distance learning, where the latter is offered via an online learning environment in which the student can access digital material, web lectures and other external digital tools and resources. Contemporary video streaming solutions nowadays allow remote students to follow courses in a similar manner as in a face-to-face setting. The computer assisted learning platform allows both remotely and face-to-face participating students to engage in interactive quizzes, polls and assignments launched by the teacher. Online tools and resources offered by third parties can further augment the learning experience. These may include communication tools, such as Slack [53] for collaboration on assignments and exchanging ideas, simulations for a physics course, online translation tools for language courses, etc. In such a scenario, measuring the engagement of a student is not straightforward for two reasons:

1. The learning context continuously changes, and a good engagement indicator in one learning context may not be adequate in another context.
2. Online content and resources of third parties cannot be modified to capture relevant interaction events that are meaningful to measure engagement.

Regarding the first observation, we not only need to distinguish between on-task and off-task behavior [7, 13], but also identify the appropriate modality for measuring engagement during on-task behavior. For example, the webcam [48] may detect whether a remote student is paying attention and looking at the screen during a video stream of the teacher. However, the webcam is not a good engagement indicator when the student is participating face-to-face, or watching irrelevant YouTube videos while the lecture is streaming in another browser tab. It is also not a useful indicator during an interactive quiz, as a webcam can only monitor facial expressions but not the interaction with the quiz itself. Mouse and keystroke dynamics may provide insights into engagement during an interactive quiz or some online simulation tools, but not when the student is chatting with friends on Facebook.

The second observation is particularly a concern for courses that rely on online resources (e.g. collaboration tools, interactive simulations, scientific articles) offered by third parties. An online learning platform can be

Fig. 1 Technology-enhanced multi-modal engagement monitoring of students in blended learning environments

augmented with logging functionality to not only audit the learning progress of the student, but also the participation in, for example, polls and quizzes. However, this is not easily achieved with third party resources if they do not offer any integration capabilities to monitor student interactions in a fine-grained manner. Furthermore, if such functionality would be available, care must be taken to not invade a student's privacy when such online resources are used both during and outside lectures (e.g. Youtube videos or Slack communication).

3.2 Edge-based multi-modal engagement analytics

To support teachers with retaining the engagement of their students, our engagement analytics framework (as depicted in Fig. 1) must be able to analyze in real-time heterogeneous static silos of data (e.g. user profile data of students) and dynamic data streams (e.g. behavioral and physiological measures, click streams, evolving col-

laborations with peers). For multi-modal learning and engagement analytics, a key challenge is that different end-points are responsible for collecting and pre-processing the raw engagement data (e.g. audiovisual data from cameras and click streams from interactions with online applications and content). That is why our edge computing solution leverages the standardized extension capabilities of the web browser to analyze the behavior of the student:

1. Keyboard usage
2. Mouse motion events and clickstream data
3. Web camera
4. Browser tab activations, updates and removals, and website snapshots
5. Browser information and window focus listener

The main benefits of our browser extension is that it (1) can analyze on- and off-task behavior on the client via the browser of the student across different operating systems and web browsers, (2) reduces the cloud-

Fig. 2 Browser extension for multi-modal and cross-site behavior engagement analysis

based workload by shifting computations towards the client, (3) measure the level of engagement in an application or website agnostic manner which simplifies integration with third party resources, and (4) avoid the transmission of privacy-sensitive information to a centralized virtual server for analysis.

All business logic to monitor and analyze engagement is implemented in cross-platform JavaScript code. It runs in the background as part of the browser extension (see Fig. 2), and does not interfere directly with visited websites (i.e. in a way that ad blockers do). Our browser extension maintains a list of URLs of websites that are considered on-task for each course. After installation of the browser extension, the student can configure and grant consent to which events can be captured and analyzed. Our browser extension has been tested in recent versions of the Google Chrome, Mozilla Firefox and Opera browsers. More details about the browser extension are available at <https://people.cs.kuleuven.be/~davy.preuveneers/lecture+/>.

3.2.1 User interaction through keyboard usage, mouse motion events and clickstream analysis

Both the keyboard usage and mouse interactions give an explicit indication of user activity. For privacy reasons, we discard the actual keys pressed, as it could otherwise incidentally capture sensitive information such

Fig. 3 Calibration of face orientation based on 5 landmarks (nose, eyes and ears)

as private communications or passwords. Mouse motion and click events only contain x,y coordinates. Both types of events do not hold any website specific details (e.g. buttons or input fields). The derived interaction features are aggregated over 10 second tumbling (i.e. non-overlapping) window intervals and interpreted in context with respect to the active URL.

Fig. 4 Cloud-based aggregation of student engagement streams on top of the Spring Cloud Data Flow framework

3.2.2 Head pose estimation via the webcam

If available, the webcam is used for real-time head pose estimation. Our browser extension uses a deep learning model similar to [49] for facial landmark detection. To evaluate the model within the web browser, it was re-implemented with TensorFlow.js³, a low-level JavaScript implementation of the TensorFlow deep learning library [1]. Our browser extension only uses a subset of the landmarks for presence detection and eye-gaze estimation. Our solution uses 5 landmarks (nose, eyes and ears) after calibration to identify whether a person is looking at the screen or not (see Fig. 3). From a practical point of view, note that webcams usually do not offer the same high-end capturing capabilities as the equipment used in [48].

3.2.3 Browser tab activity for context-dependent interpretation of engagement indicators

The browser extension keeps track of whether the web browser or another application is running in the foreground. Furthermore, it monitors which browser tab is currently active and which website/URL is being visited. This way, the active browser tab can indicate whether the student is interacting with (1) the main online learning platform (e.g. slides, assignment, poll, quiz), (2) an external resource associated with the course (e.g. a collaboration hub such as Slack) or (3) a website considered as off-task (e.g. social media websites).

The tab snapshots (i.e. a snapshot image of the website) are used to compute a perceptual hash [51] of the website. The advantage of these hashes is that the Hamming distance between two hashes indicates how far visually apart two images are, irrespective of their resolutions. If only a few bits of the hash differ, it is a

³ <https://js.tensorflow.org/>

Fig. 5 Visualization of performance and throughput through online dashboard

good indication that the two images are identical. Our browser extension relies on perceptual hashes when the URL of a website is not a good discriminator (e.g. in the case of dynamic web content). This way, the websites that individual students visit can be automatically compared for similarity without disclosing the actual current content of the website. This further improves the privacy of the engagement monitor.

3.3 Context-dependent aggregation of engagement indicators

The teacher can influence the importance of each of the engagement indicators during the lecture. By launching a quiz or poll, students are expected to interact with the online learning platform. During these brief periods, the explicit user interaction events have a higher weight compared to the webcam analytics. The assumption is that in order to interact with the poll or quiz, the student has to look at the screen anyhow. When the teacher is presenting the content of the course or drawing on the blackboard, the expected level of interaction by the students will be limited (e.g. raising hand to ask a question). In that case, the webcam will be the

only source of information to confirm the student is engaged, unless off-task behavior – such as interaction on social media – already provides strong indication that the student is not engaged in the lecture.

The browser extension computes an aggregated score between -1.0 and 1.0 for each 10 second interval. Every second, it checks for off-task and on-task behavior, and computes the engagement score as the cumulative sum as follows:

$$\text{score} = \sum_{t=1}^{10} w_t$$

$$\text{with } w_t = \begin{cases} -0.1, & \text{off-task behavior} \\ 0.0, & \text{no engagement events} \\ 0.05, & \text{webcam engaged} \\ 0.1, & \text{webcam+interaction engaged} \end{cases}$$

This simple cumulative weighted sum is fairly easy to extend with new engagement indicators (e.g. physiological data). To analyze or visualize longer term engagement trends, the above score is further aggregated in a 5 minute sliding interval (i.e. the normalized sum of the last 30 scores). The actual values for the duration of the tumbling interval ($\Delta t = 10$ seconds), the

sliding interval ($\Delta t = 5$ minutes), and the values w_t for (dis)engaged behavior can be customized, though should be consistent across all students.

The browser extension provides a visualization of the above values, as well as a slider for a user to self-report his engagement level. This value is currently logged but not taken into consideration for cross-student analysis, as it is not trivial to calibrate the self-reported engagement levels.

3.4 Server-side aggregation of engagement streams with different data processing pipelines

The browser-based multi-modal engagement analysis is complemented with a server-side backend (see Fig. 4) that aggregates the high-level engagement scores of the individual students. It offers a WebSocket interface to easily integrate the high-level aggregated event stream into a dedicated visualization dashboard for the teacher. When overall engagement drops, it may be time for the teacher to launch another quiz or poll, and the system make it also possible to evaluate the effect of this intervention on students' engagement.

Furthermore, the backend implements the same multi-modal engagement analysis functionality available in the browser. This allows us to systematically compare the performance gains of analyzing the raw data streams on the server versus shifting (part of) the processing workload towards the student's browser. The performance and throughput of the event stream processing is visualized by means of an online dashboard, as depicted in Fig. 5.

3.5 Secure cross-user aggregation of engagement scores

Our edge-based framework further improves the privacy of its users by aggregating the engagement scores of the individual students in a privacy preserving manner so that a teacher only has an aggregated engagement score for the entire audience. This is achieved by summing the engagement scores by means of secure multi-party computation. Secure multi-party computation [9] allows computations on encrypted values. This way it is possible to compute an average engagement score without any individual group member revealing her personal engagement score to the others. We make use of the JIFF JavaScript library⁴ to compute the overall engagement level.

⁴ <https://multiparty.org/>, <https://github.com/multiparty/jiff>

4 Evaluation

This section reports on various qualitative and quantitative evaluation metrics of our solution, identifying performance and impact trade-offs. We measure the overhead of the engagement analysis at the client-side, and carry out a systematic performance analysis of centralized versus edge-based engagement analysis on the server – hosting an Intel Core i7-7700 CPU @ 3.60GHz with 32GB of memory – for a growing number of students.

4.1 Performance impact on server for centralized and edge-based engagement analysis configurations

We collected real-world interaction traces of student behavior during various interactive lectures⁵. The data being collected by the web browser extension consists of the following types of information and events:

- **Personal details:** The first and last name of the individual, as well as the email address.
- **Consent:** Explicit consent from the individual to collect and process information, as well as the selection of events that may be processed.
- **Device configuration:** Name and version of the web browser, as well as the operating system on top of which it is running.
- **Self-reported engagement events:** The individual can provide a score, from -2 to 2, indicating the level of (dis)engagement.
- **Browser focus events:** Events indicated whether the web browser gains from or loses focus to another application running on the individual's device.
- **Browser tab events:** Events indicating when a browser tab was created, deleted, activated, or updated when visiting a new website.
- **Keyboard events:** The amount of keystrokes in a given fixed time interval (e.g. every 10 seconds).
- **Mouse events:** The amount of mouse clicks and mouse motions in a given fixed time interval (e.g. every 10 seconds).
- **Head pose events:** Events every second indicating whether the individual is looking at the screen, looking away, or is not present (see Fig. 3).

Due to privacy reasons, we only capture the amount, but not the details of keyboard and mouse events. So, it is not possible to reconstruct what was typed, or where on the website the individual clicked. We correlate these events with the URL of the website. Again,

⁵ The data collection was approved on April 2018 by the Social and Societal Ethics Committee of the university with case number G-2018 04 1206

Fig. 6 Server-side performance impact in cloud

Fig. 7 Client-side performance impact (average CPU load over 1 minute) on edge device, measured through task manager of Google Chrome browser

due to privacy reasons, we check whether the website is white-listed as *on-task* (as an indication the website is relevant for the ongoing lecture), and classify it into 5 different categories: *on-task*, *news site*, *social network*, *search engine*, or *other*. After classification, the original URL is discarded.

From the 25 students that consented to have their data collected, we filtered the top 3 that created the most events, and used those traces to create a baseline for a “worst case scenario” from a performance point of view. These traces were replicated and replayed to simulate from 10 up to 500 concurrent students to analyze the scalability of the back-end of our solution. Fig. 6 compares the performance impact in terms of CPU load on the server for both an edge-based and a centralized data processing pipeline. For the centralized variant, we distinguish two variations where the deep learning model for head pose detection is either run on the CPU or delegated to a CUDA card (more

specifically, an NVIDIA Titan V graphics card). It is clear that the edge-based configuration scales to many more users (the maximum capacity of the server was not reached at 500 concurrent students). Furthermore, by sending the raw data to the server in the centralized scenario, the network bandwidth usage is more than 50 times higher due to the need to stream webcam images to the server. The additional overhead is less significant if the webcam video of each student is streamed to all participants anyhow (e.g. via the WebRTC protocol). However, in such a scenario, the bottleneck in terms of concurrent users is the network capacity on the server.

4.2 Performance impact on edge device for client-side engagement analysis

Given that the deep learning component for head pose estimation and eye-gaze analysis is computationally the most intensive, we tested the practical feasibility of this

component with Google Chrome 69 running on different Dell laptops, the oldest being a Dell Latitude E6330 of more than 5 years old and the newest being a Dell Latitude 7480 of more than 1 year old:

- *5 year old laptop*: Dell Latitude E6330 with Intel Core i5-3320M CPU @ 2.60GHz, 8 GB memory
- *1 year old laptop*: Dell Latitude 7480 with Intel Core i7-7600U CPU @ 2.80GHz, 32 GB memory

For a 640x480 webcam configuration, the maximum framerate including face landmark detection varied between 12 and 22 frames per second, but for engagement analysis one frame per second suffices. Neither of the laptops relied on CUDA hardware acceleration to evaluate the deep learning models. Fig. 7 provides a breakdown of the computational overhead of the browser extension into the individual engagement indicators. These values were obtained through the built-in task manager of Google Chrome, which allows to measure the memory footprint, CPU and network usage for each browser tab and extension. It is clear that the JavaScript based Tensorflow model for head pose estimation has the biggest overhead. TensorFlow.js can use a GPU to accelerate math operations. So if available on the device, we expect the performance overhead on the client to be significantly less.

4.3 Impact and limitations for student engagement

Complementary to this work and within the frame of the same ICON LECTURE+ research project, Raes et al. [42] previously investigated the effect on students' engagement as a result of technology-enhanced quizzes launched throughout the online education platform. The interactions of the students with the education platform are also monitored by our browser-based solution.

The authors identified a positive effect from the quizzes on the students' motivation but that further research is necessary to validate the impact of different kinds and the timing of quizzes. The authors also acknowledge that hybrid virtual classrooms are a promising educational environment for flexible learning, but that they remain challenging to teach in and to learn in as a remote participant.

One shortcoming identified in the web browser extension is the currently limited support for the notion of synchrony as an indicator for engagement. Kawamura et al. [30] recently investigated engagement based on the synchrony of head movements in e-learning scenarios. Changes in appearance, such as looking away, may be an indicator that a student is disengaged. However, when all students look away at the same time and in the

same direction, this notion of synchrony may be an indication that the teacher is walking around. In an online environment, students may be looking away when working on a pen-and-paper assignment. In both situations, the students are in fact engaged, but an individual assessment may lead to contradictory results. Indeed, synchrony is an engagement indicator that operates on a group level, rather than individually. As head pose and movements are currently analyzed within the browser for privacy reasons, novel techniques are needed to use synchrony as a collective engagement indicator without jeopardizing the individual's privacy.

4.4 Security analysis and threat model

The web browser extension collects and analyzes a variety of multi-modal event data and can therefore be considered as fairly invasive. An immediate concern is whether the web browser extension may be abused and therefore imposes a security risk. To analyze this risk, we conceive a threat model with either an *honest but curious* party or a *malicious* party. The former party only aims to intercept possibly sensitive information, whereas the latter also manipulates the data. The parties to be considered in the given scenario are (1) the student, (2) the teacher, and (3) a curious or malicious website.

The web browser extension interacts with any website to intercept useful events, but the website cannot directly interact with the extension to intercept and exfiltrate sensitive data. It may, however, use JavaScript code to emulate mouse click or keyboard events that may be mistakenly perceived by the web browser extension as genuine student interaction. However, only white-listed URLs at which such events are collected contribute positively to a student's engagement level, whereas all other websites have a negative influence. As such, a malicious website would not have a significant impact on the engagement level as the website would not be white-listed.

The teacher cannot interact with the software running at the client-side in the student's web browser. However, anyone with the necessary privileges to interact with the back-end infrastructure, may be able to manipulate the aggregated data. In theory, such a teacher may modify the engagement results, but such interventions would beat the purpose of the framework and that is to help the teacher keeping the students engaged. The goal of the framework is not to evaluate the lecturing skills of the teacher.

From a fraud point of view, the student is still able to cheat, as the software runs on his own device, and the code of the browser extension could be modified to

send fabricated data. As there is no solution to make the client-side code tamper-proof, this is a risk that is difficult to mitigate without the proper hardware support, including Trusted Execution Environments (TEE) such as the Intel SGX. Previous work by Krawiecka et al. [32] proposed a solution, called SafeKeeper, to securely store web passwords using TEEs. A similar solution would be required to make the engagement analysis tamper-proof, whereby the client-side code that analyzes engagement would need to run within the TEE. However, even in that scenario, there is still an opportunity for a student to cheat using an additional web browser extension that emulates fake interaction events for the white-listed websites.

4.5 Privacy impact analysis

Compared to online learning analytics solutions that centralize and process all raw data in the cloud, our edge-based system provides the following benefits:

- The user remains in control of which data is collected, and must provide consent before any data is captured.
- The individual engagement score is computed on the user’s client. No sensitive data (e.g. keystrokes, webcam or website snapshots) are sent to the cloud.
- The cross-user engagement score is aggregated by means of secure multi-party computation so that individual scores are not revealed to others.

While efficiently preserving students’ privacy, our approach merely reduces engagement monitoring to a single aggregated value per student or for a group of students. Yet, student engagement has been regarded as a complex analysis built upon multiple factors. Given that teachers hold a prominent position in shaping levels of engagement to improve student outcomes [21], leveraging richer data might boost their ability to provide personalized feedback to their students. In an online scenario, this translates into exposing possibly sensitive student data for the sake of conveying richer information and creating a better characterization of student engagement. For example, automatic engagement monitoring from facial expressions has been regarded as a potential future asset in remote learning environments [39,50]. Whitehall et al. [50] used state-of-the-art machine learning algorithms employed in emotion recognition to classify the levels of user engagement, finding a positive correlation between the automatic recognition and students outcome after receiving feedback.

Enhancing our current solution with emotion recognition requires the student to collaborate for training

the model, which can be done in two ways. The first one is decentralized training of a common machine learning-based engagement recognition model, which allows the students to keep their data local while contributing to the process by only sharing incremental training updates or model parameters. Federated learning embodies this concept and enhances privacy by protecting sensitive data [31]. However, it has been proven vulnerable to data leakage under certain attack scenarios [37]. Additionally, exploiting federated learning prevents from central pre-processing, such as teacher-aided labelling of video frames, ultimately affecting the classification outcome. The second solution consists of sharing video frames with a central node which is in charge of training the engagement recognition model based on all students’ input. However, sharing sensitive raw input might lead to various privacy issues: (1) a face image represents a biometric trait from which a number of personal details can be deduced, a harbinger of function creep by curious third-parties [17]; (2) sharing sensitive data might leak through an untrusted communication channel or due to mismanagement by the service provider. By anonymizing data upon sharing, we can protect against the aforementioned threats. Traditional techniques, like differential privacy, provide us with formal guarantees w.r.t. private data release at the expense of a drop in performance [2]. Conversely, context-aware techniques involving deep feature extraction may allow for a better balance between privacy and utility [27]. By learning a compressed representation of the original input, context-aware privacy models may contribute to overcome network bandwidth constraints, which is especially advantageous for client devices that lack the resources to perform sophisticated machine learning tasks.

5 Conclusion and future work

In this work, we presented an edge-based multi-modal engagement solution that runs as an extension with contemporary web browsers. It supports the monitoring of interaction data with the online learning platform as well third party resources, and can analyze head pose estimation using a deep learning model implemented and evaluated in JavaScript. The added value of our contribution is that the multi-modal engagement analysis is off-loaded towards the students’ browsers, allowing our framework to easily scale up with a growing number of students. At the same time, it mitigates any privacy concerns that students may have due to continuous tracking of their on-task and off-task behavior. One limitation of our solution is that interaction with

non-browser tools (e.g. a native word processing or presentation application) used in the frame of a course are ignored to monitor engagement.

As future work, we will enhance autotuning of the configuration of the browser extension to automatically reduce the performance overhead on the client below a certain threshold. Also, the current solution allows for students to cheat by manipulating the client-side data collection process. As future work, we will investigate the feasibility for temporary replication of the engagement analysis on both the clients and server to detect adversarial behavior.

Acknowledgment

This research is partially funded by the Research Fund KU Leuven, and by imec through ICON LECTURE+. LECTURE+ is a project realized in collaboration with imec, with Barco, Televic Education and Limecraft as project partners and with project support from VLAIO.

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan V GPU used for this research.

References

1. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., Zheng, X.: Tensorflow: A system for large-scale machine learning. In: Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation. pp. 265–283. OSDI'16, USENIX Association, Berkeley, CA, USA (2016), <http://dl.acm.org/citation.cfm?id=3026877.3026899>
2. Abadi, M., Chu, A., Goodfellow, I., McMahan, H.B., Mironov, I., Talwar, K., Zhang, L.: Deep learning with differential privacy. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 308–318. ACM (2016)
3. Abbas, N., Zhang, Y., Taherkordi, A., Skeie, T.: Mobile edge computing: A survey. *IEEE Internet of Things Journal* **5**(1), 450–465 (2018)
4. Acar, A., Aksu, H., Uluagac, A.S., Conti, M.: A survey on homomorphic encryption schemes: Theory and implementation. *ACM Comput. Surv.* **51**(4), 79:1–79:35 (Jul 2018). <https://doi.org/10.1145/3214303>, <http://doi.acm.org/10.1145/3214303>
5. de Assuncao, M.D., da Silva Veith, A., Buyya, R.: Distributed data stream processing and edge computing: A survey on resource elasticity and future directions. *Journal of Network and Computer Applications* **103**, 1–17 (2018)
6. Atherton, M., Shah, M., Vazquez, J., Griffiths, Z., Jackson, B., Burgess, C.: Using learning analytics to assess student engagement and academic outcomes in open access enabling programmes. *Open Learning: The Journal of Open, Distance and e-Learning* **32**(2), 119–136 (2017). <https://doi.org/10.1080/02680513.2017.1309646>, <https://doi.org/10.1080/02680513.2017.1309646>
7. Baker, R.S.: Modeling and understanding students' off-task behavior in intelligent tutoring systems. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 1059–1068. CHI '07, ACM, New York, NY, USA (2007). <https://doi.org/10.1145/1240624.1240785>, <http://doi.acm.org/10.1145/1240624.1240785>
8. Beimel, A.: Secret-sharing schemes: a survey. In: International Conference on Coding and Cryptology. pp. 11–46. Springer (2011)
9. Ben-David, A., Nisan, N., Pinkas, B.: Fairplaymp: A system for secure multi-party computation. In: Proceedings of the 15th ACM Conference on Computer and Communications Security. pp. 257–266. CCS '08, ACM, New York, NY, USA (2008). <https://doi.org/10.1145/1455770.1455804>, <http://doi.acm.org/10.1145/1455770.1455804>
10. Bonomi, F., Milito, R., Natarajan, P., Zhu, J.: Fog computing: A platform for internet of things and analytics. In: Big data and internet of things: A roadmap for smart environments, pp. 169–186. Springer (2014)
11. Bonomi, F., Milito, R., Zhu, J., Addepalli, S.: Fog computing and its role in the internet of things. In: Proceedings of the First Edition of the MCC Workshop on Mobile Cloud Computing. pp. 13–16. MCC '12, ACM, New York, NY, USA (2012). <https://doi.org/10.1145/2342509.2342513>, <http://doi.acm.org/10.1145/2342509.2342513>
12. Cao, H., Wachowicz, M., Cha, S.: Developing an edge computing platform for real-time descriptive analytics. In: 2017 IEEE International Conference on Big Data (Big Data). pp. 4546–4554 (Dec 2017). <https://doi.org/10.1109/BigData.2017.8258497>
13. Cetintas, S., Si, L., Xin, Y.P., Hord, C., Zhang, D.: Learning to identify students' off-task behavior in intelligent tutoring systems. In: Proceedings of the 2009 Conference on Artificial Intelligence in Education: Building Learning Systems That Care: From Knowledge Representation to Affective Modelling. pp. 701–703. IOS Press, Amsterdam, The Netherlands (2009), <http://dl.acm.org/citation.cfm?id=1659450.1659578>
14. Cramer, R., Damgrd, I.B., Nielsen, J.B.: Secure Multi-party Computation and Secret Sharing. Cambridge University Press, New York, NY, USA, 1st edn. (2015)
15. Danezis, G., Domingo-Ferrer, J., Hansen, M., Hoepman, J., Métayer, D.L., Tirtea, R., Schiffner, S.: Privacy and data protection by design - from policy to engineering. *CoRR abs/1501.03726* (2015), <http://arxiv.org/abs/1501.03726>
16. Daniel, B.K.: Big data and learning analytics in higher education. Springer (2016)
17. Dantcheva, A., Elia, P., Ross, A.: What else does your biometric data reveal? a survey on soft biometrics. *IEEE Transactions on Information Forensics and Security* **11**(3), 441–467 (March 2016). <https://doi.org/10.1109/TIFS.2015.2480381>
18. van Dijk, M., Gentry, C., Halevi, S., Vaikuntanathan, V.: Fully homomorphic encryption over the integers. In: Gilbert, H. (ed.) *Advances in Cryptology – EUROCRYPT 2010*. pp. 24–43. Springer Berlin Heidelberg, Berlin, Heidelberg (2010)
19. Dwork, C.: Differential privacy. *Encyclopedia of Cryptography and Security* pp. 338–340 (2011)
20. El-Yahyaoui, A., El Kettani, M.D.E.C.: Fully homomorphic encryption: Searching over encrypted cloud

- data. In: Proceedings of the 2Nd International Conference on Big Data, Cloud and Applications. pp. 10:1–10:5. BDCA'17, ACM, New York, NY, USA (2017). <https://doi.org/10.1145/3090354.3090364>, <http://doi.acm.org/10.1145/3090354.3090364>
21. Fredricks, J.A., Blumenfeld, P.C., Paris, A.H.: School engagement: Potential of the concept, state of the evidence. *Review of Educational Research* **74**(1), 59–109 (2004). <https://doi.org/10.3102/00346543074001059>, <https://doi.org/10.3102/00346543074001059>
 22. Gong, L., Liu, Y., Zhao, W.: Using learning analytics to promote student engagement and achievement in blended learning: An empirical study. In: Proceedings of the 2Nd International Conference on E-Education, E-Business and E-Technology. pp. 19–24. ICEBT 2018, ACM, New York, NY, USA (2018). <https://doi.org/10.1145/3241748.3241760>, <http://doi.acm.org/10.1145/3241748.3241760>
 23. Gray, C.C., Perkins, D.: Utilizing early engagement and machine learning to predict student outcomes. *Computers & Education* **131**, 22 – 32 (2019). <https://doi.org/10.1016/j.compedu.2018.12.006>, <http://www.sciencedirect.com/science/article/pii/S0360131518303191>
 24. He, J., Wei, J., Chen, K., Tang, Z., Zhou, Y., Zhang, Y.: Multitier fog computing with large-scale iot data analytics for smart cities. *IEEE Internet of Things Journal* **5**(2), 677–686 (2018)
 25. Henrie, C.R., Halverson, L.R., Graham, C.R.: Measuring student engagement in technology-mediated learning: A review. *Computers & Education* **90**, 36 – 53 (2015). <https://doi.org/10.1016/j.compedu.2015.09.005>, <http://www.sciencedirect.com/science/article/pii/S0360131515300427>
 26. Holzer, A., Franz, M., Katzenbeisser, S., Veith, H.: Secure two-party computations in ansi c. In: Proceedings of the 2012 ACM Conference on Computer and Communications Security. pp. 772–783. CCS '12, ACM, New York, NY, USA (2012). <https://doi.org/10.1145/2382196.2382278>, <http://doi.acm.org/10.1145/2382196.2382278>
 27. Hukkelås, H., Mester, R., Lindseth, F.: DeepPrivacy: A Generative Adversarial Network for Face Anonymization. arXiv e-prints arXiv:1909.04538 (Sep 2019)
 28. Kahu, E.R.: Framing student engagement in higher education. *Studies in Higher Education* **38**(5), 758–773 (2013). <https://doi.org/10.1080/03075079.2011.598505>, <https://doi.org/10.1080/03075079.2011.598505>
 29. Kahu, E.R., Nelson, K.: Student engagement in the educational interface: understanding the mechanisms of student success. *Higher Education Research & Development* **37**(1), 58–71 (2018). <https://doi.org/10.1080/07294360.2017.1344197>, <https://doi.org/10.1080/07294360.2017.1344197>
 30. Kawamura, R., Toyoda, Y., Niinuma, K.: Engagement estimation based on synchrony of head movements: Application to actual e-learning scenarios. In: Proceedings of the 24th International Conference on Intelligent User Interfaces: Companion. pp. 25–26. IUI '19, ACM, New York, NY, USA (2019). <https://doi.org/10.1145/3308557.3308660>, <http://doi.acm.org/10.1145/3308557.3308660>
 31. Konečný, J., McMahan, H.B., Yu, F.X., Richtárik, P., Suresh, A.T., Bacon, D.: Federated learning: Strategies for improving communication efficiency. *CoRR abs/1610.05492* (2016), <http://arxiv.org/abs/1610.05492>
 32. Krawiecka, K., Kurnikov, A., Paverd, A., Mannan, M., Asokan, N.: Safekeeper: Protecting web passwords using trusted execution environments. In: Proceedings of the 2018 World Wide Web Conference. pp. 349–358. WWW '18, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland (2018). <https://doi.org/10.1145/3178876.3186101>, <https://doi.org/10.1145/3178876.3186101>
 33. Li, N., Li, T., Venkatasubramanian, S.: t-closeness: Privacy beyond k-anonymity and l-diversity. In: 2007 IEEE 23rd International Conference on Data Engineering. pp. 106–115. IEEE (2007)
 34. Lorenz, B., Sousa, S., Tomberg, V.: Privacy awareness of students and its impact on online learning participation – a case study. In: Ley, T., Ruohonen, M., Laanpere, M., Tatnall, A. (eds.) *Open and Social Technologies for Networked Learning*. pp. 189–192. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
 35. Ma, X., Zhang, F., Chen, X., Shen, J.: Privacy preserving multi-party computation delegation for deep learning in cloud computing. *Information Sciences* **459**, 103–116 (2018)
 36. May, M., George, S.: Using students' tracking data in e-learning: Are we always aware of security and privacy concerns? In: 2011 IEEE 3rd International Conference on Communication Software and Networks. pp. 10–14 (May 2011). <https://doi.org/10.1109/ICCSN.2011.6013764>
 37. Melis, L., Song, C., Cristofaro, E.D., Shmatikov, V.: Inference attacks against collaborative learning. *CoRR abs/1805.04049* (2018), <http://arxiv.org/abs/1805.04049>
 38. Mohan, P., Thakurta, A., Shi, E., Song, D., Culler, D.: Gupt: Privacy preserving data analysis made easy. In: Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data. pp. 349–360. SIGMOD '12, ACM, New York, NY, USA (2012). <https://doi.org/10.1145/2213836.2213876>, <http://doi.acm.org/10.1145/2213836.2213876>
 39. Monkaresi, H., Bosch, N., Calvo, R.A., D'Mello, S.K.: Automated detection of engagement using video-based estimation of facial expressions and heart rate. *IEEE Transactions on Affective Computing* **8**(1), 15–28 (Jan 2017). <https://doi.org/10.1109/TAFFC.2016.2515084>
 40. Prabhakaran, M., Sahai, A.: *Secure Multi-Party Computation*. IOS Press, Amsterdam, The Netherlands, The Netherlands (2013)
 41. Preuveneers, D., Joosen, W.: Edge-based and privacy-preserving multi-modal monitoring of student engagement in online learning environments. In: Proceedings of the IEEE International Conference on Edge Computing (IEEE EDGE 2019). pp. 1–3. IEEE (July 2019)
 42. Raes, A., Vanneste, P., Pieters, M., Windey, I., Noortgate, W.V.D., Depaepe, F.: Learning and instruction in the hybrid virtual classroom: An investigation of students' engagement and the effect of quizzes. *Computers & Education* **143**, 103682 (2020). <https://doi.org/10.1016/j.compedu.2019.103682>, <http://www.sciencedirect.com/science/article/pii/S0360131519302350>
 43. Roman, R., Lopez, J., Mambo, M.: Mobile edge computing, fog et al.: A survey and analysis of security threats and challenges. *Future Generation Computer Systems* **78**, 680–698 (2018)
 44. Satyanarayanan, M.: The emergence of edge computing. *Computer* **50**(1), 30–39 (2017)

45. Shan, Z., Ren, K., Blanton, M., Wang, C.: Practical secure computation outsourcing: a survey. *ACM Computing Surveys (CSUR)* **51**(2), 31 (2018)
46. Shi, W., Cao, J., Zhang, Q., Li, Y., Xu, L.: Edge computing: Vision and challenges. *IEEE Internet of Things Journal* **3**(5), 637–646 (2016)
47. Strauss, V.: Why parents and students are protesting an online learning program backed by mark zuckerberg and facebook. [https://www.washingtonpost.com/education/2018/12/20/why-parents-students-are-protesting-an-online-learning-program-backed-by-mark-zuckerberg-facebook/\(December 2018\)](https://www.washingtonpost.com/education/2018/12/20/why-parents-students-are-protesting-an-online-learning-program-backed-by-mark-zuckerberg-facebook/(December%202018))
48. Thomas, C., Jayagopi, D.B.: Predicting student engagement in classrooms using facial behavioral cues. In: *Proceedings of the 1st ACM SIGCHI International Workshop on Multimodal Interaction for Education*. pp. 33–40. MIE 2017, ACM, New York, NY, USA (2017). <https://doi.org/10.1145/3139513.3139514>, <http://doi.acm.org/10.1145/3139513.3139514>
49. Toshev, A., Szegedy, C.: DeepPose: Human pose estimation via deep neural networks. In: *2014 IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1653–1660 (June 2014). <https://doi.org/10.1109/CVPR.2014.214>
50. Whitehill, J., Serpell, Z., Lin, Y., Foster, A., Movellan, J.R.: The faces of engagement: Automatic recognition of student engagement from facial expressions. *IEEE Transactions on Affective Computing* **5**(1), 86–98 (Jan 2014). <https://doi.org/10.1109/TAFFC.2014.2316163>
51. Yang, B., Gu, F., Niu, X.: Block mean value based image perceptual hashing. In: *Proceedings of the 2006 International Conference on Intelligent Information Hiding and Multimedia*. pp. 167–172. IHH-MSP '06, IEEE Computer Society, Washington, DC, USA (2006). <https://doi.org/10.1109/IHH-MSP.2006.66>, <http://dx.doi.org/10.1109/IHH-MSP.2006.66>
52. Yi, S., Hao, Z., Qin, Z., Li, Q.: Fog computing: Platform and applications. In: *2015 Third IEEE Workshop on Hot Topics in Web Systems and Technologies (HotWeb)*. pp. 73–78. IEEE (2015)
53. Zhang, X., Meng, Y., de Pablos, P.O., Sun, Y.: Learning analytics in collaborative learning supported by slack: From the perspective of engagement. *Computers in Human Behavior* (2017). <https://doi.org/10.1016/j.chb.2017.08.012>, <http://www.sciencedirect.com/science/article/pii/S0747563217304788>

researchers and projects on these topics with particular focus on applications in the area of the Internet of Things, pervasive and ubiquitous computing, e-health, finance, media and Industry 4.0. He manages various projects with industry on user friendly and privacy-aware federated authentication, context-aware behavioral authentication, enhanced authorization and access control in e-health, batch and streaming computation and federated data analytics for enterprise-level context-aware applications, and the systematic scalability analysis of complex enterprise applications.

Giuseppe Garofalo is a cyber security researcher and PhD student at the imec-DistriNet research group of the Department of Computer Science, KU Leuven, Belgium. He obtained his M.Sc. in Computer Engineering at the Politecnico di Torino, Italy. In the fall of 2018, he started pursuing a PhD in Computer Science on the topic of Cyber Security. His field of interest is the use of machine learning techniques for cyber security. He works on the mitigation of privacy leaks of user behavior in deep neural network applications, and adversarial machine learning for biometric and behavioral authentication.

Prof. dr. Wouter Joosen is a full professor at the Department of Computer Science of KU Leuven, Belgium. He heads the imec-DistriNet Research Group and is leading multiple research projects related to cybersecurity in the context of the web, cloud computing and data management environments. He has been a full Professor in Software Engineering (part time) at the Maersk Mc Kinney institute for production technology (Denmark). Wouter co-founded several spin off companies of KU Leuven, including Ubizen, Luciad, Inmanta and Elimity. His research interests include software architectures for secure distributed systems; critical security services such as authentication, authorization, and audit; security analytics; security and privacy by design, development and deployment support for secure software.

Dr. Davy Preuveneers is a Research Manager of imec-DistriNet, the Distributed and Secure Software research group of the Department of Computer Science of KU Leuven, Belgium. His research interests are in the field federated identity and access management, behavioral authentication, data mining and machine learning for security and privacy, and scalable distributed systems and event-based middleware for smart mobile and embedded applications. He leads a team of