

Learning Analytics Technology to Understand Learner Behavioral Engagement in MOOCs

Zhao, Yue

DOI

[10.4233/uuid:2aff1a7e-45eb-4d10-9944-8e06ef12b9fa](https://doi.org/10.4233/uuid:2aff1a7e-45eb-4d10-9944-8e06ef12b9fa)

Publication date

2019

Document Version

Final published version

Citation (APA)

Zhao, Y. (2019). *Learning Analytics Technology to Understand Learner Behavioral Engagement in MOOCs*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:2aff1a7e-45eb-4d10-9944-8e06ef12b9fa>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



Learning Analytics Technology to Understand Learner Behavioral Engagement in MOOCs



Yue Zhao

Learning Analytics Technology to Understand Learner Behavioral Engagement in MOOCs

Yue Zhao

Learning Analytics Technology to Understand Learner Behavioral Engagement in MOOCs

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus prof.dr.ir. T.H.J.J. van der Hagen
Chair of the Board for Doctorates
to be defended publicly on
Tuesday, 9 April, 2019 at 12:30 o'clock

by **Yue ZHAO**

Master of Science in Computer Science and Technology
National University of Defense Technology, China
born in Zhengzhou, China.

This dissertation has been approved by the promotors.

Composition of the doctoral committee:

Rector Magnificus	chairperson
Prof.dr.ir. G.J.P.M. Houben	Delft University of Technology, promotor
Dr. C. Hauff	Delft University of Technology, copromotor
Dr. C. Lofi	Delft University of Technology, copromotor

Independent members:

Prof.dr. M.M. Specht	Delft University of Technology
Dr. A.E. Zaidman	Delft University of Technology
Prof.dr. V.G. Dimitrova	University of Leeds
Prof.dr. H. Drachsler	Goethe University Frankfurt am Main & Leibniz Research Institute DIPF.de
Prof.dr. E. Visser	Delft University of Technology, reserve member



SIKS Dissertation Series No. 2019-11 The research reported in this thesis has been carried out under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems. This work is supported by China Scholarship Council (CSC) under grant 201403170430.

Published and distributed by: Yue Zhao
E-mail: zhaoyue.89@gmail.com

ISBN: 978-94-028-1462-0

Keywords: Learning Analytics, MOOCs, Learner Engagement

Copyright © 2019 by Yue Zhao

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without written permission of the author.

Cover image: TU Delft campus by Marco van Middelkoop/Aerophoto-Schiphol.
Cover design by: Yue Zhao and Xing Li

Printed and bound in The Netherlands by Ipskamp Printing .

Acknowledgments

Today, I can still recall the first day I arrived in the Netherlands in September 2014. When I left the Schiphol airport before dawn, there was a little chill in the air. I was sitting in a taxi with my friends, watching the streetlights on the highway, and having an expectation for the future. When the taxi was driving, the sun gradually rose, and when we arrived at my studio, it is completely light already. After arrival, I took a short break and my new life started. Recalling the time of being a Ph.D. student for more than four years, I think it is like the journey of the first day. My heart was unsettled but my steps were steady. There were some people leading my way and some people going together with me, and the road was becoming more and more clearly. Therefore, at the time of the completion of my thesis, I would like to express my most sincere gratitude to many people.

First of all, I would like to express my highest gratitude to my promoter Geert-Jan Houben. Thank you for offering me an opportunity to be a member of the WIS group. Your guidance in our rho group meeting and 1-1 meeting helped me overcome the obstacles and find the right direction in my research career. I also thank you for the financial support with which I can participate in many conferences and summer schools in these years. Without your guidance, support, and encouragement, it is impossible for me to keep going through these years.

Secondly, I would like to thank my two co-promoters Claudia Hauff and Christoph Lofi. Thank you for bringing your critical thinking and high standards when I did research on each topic in my thesis. Thank you for spending plenty of time in the discussion and revision when we make each of my publication together. I am not a good storyteller, but you always keep encouraging me to express what in my mind.

It is my honor to have Marcus Specht, Andy Zaidman, Vania Dimitrova, Hendrik Drachsler and Eelco Visser as my committee members.

Next, I would like to thank my master supervisor, Yijie Wang. Thank you for leading me to start scientific research and encouraging me to apply

for a Ph.D. position. During these years, you have been concerned about my research progress and personal development.

I would also like to thank our current and former WISers. Thank Dan Davis, Guanliang Chen, Tarmo Robal, and Wing Nguyen for your help and collaboration. I learned a lot and really enjoyed to work with you. Thank Alessandro Bozzon, Asterios Katsifodimos, Nava Tintarev, Achilleas Psyllidis, Mohammad Khalil, Ke Tao, Jie Yang for your support and advice in my research. Thank Sepideh Mesbah, Shahin Sharifi, Sihang Qiu, Shabnam Najafian and Felipe Moraes for the wonderful time you shared with me in our Ph.D. group dinners. Thank Tamara Brusik, Roniet Sharabi, and Naomi Meyer for your kind help and nice arrangement every time when I participated in conferences and extended my visa in these years.

I have met a lot of friends in these years, and I am really grateful for the life we shared together. First, I would like to thank my former and current housemate Xiang Fu, Jian Fang, Tiantian Du, Jiapeng Yin, Yuan Li, Shengzhi Xu, and Yinzheng Qu for most of my time in these years I shared with you and various daily activities we did together. Second, I would like to thank Yu Xin, Zhijie Ren, Shanshan Ren, Renfei Bu, Xiaohui Wang, Rong Zhang, Xin Guo, Jintao Yu and Zhaokun Guo for sharing dinner and board games with me on many nights. I am especially grateful to Zixuan Zheng and Lingling Lao for their handmade noodles on many Saturdays. Then, I would like to thank my adventure buddies Qiang Liu, Xin Du, Rui Jian, Oscar Li, Xiao Gong, Hai Zhu, Longjian Piao, Mengmeng Ye, Meng Meng, Xiang Teng, Yancong Lin, Tao Lv, Yang Qu, and Mingjuan Zhao for the good time we used to ski, climb or swim together. I would also like to thank Baozhou Zhu, Bo Zhao, Changlin Chen, Guangming Li, Hai Gong, He Wang, Jiangpeng Zhang, Jie Shen, Jie Zhao, Jinyu Tang, Laobing Zhang, Lizhou Wu, Minghe Shan, Nannan Chen, Pengling Wang, Qianqian Zhu, Shilong Sun, Shuai Yuan, Siqi Shen, Vincent Gong, Wenjie Pei, Xiangrong Wang, Xinyi Li, Yande Jiang, Yangyueye Cao, Yazhou Yang, Yingying Bao, Yong Guo, Yue Chen, Yunlong Li, Yutian Yao, Yuxiao Hou, Yuxin Wang, Zhe Hou, and Zhi Hong for your help and friendship.

I owe a big thank to my wife, Xing Li. Because of academic and professional pursuits, we have experienced long-term life in different places, but our relationship was getting better and better during this time. In these years, in addition to completing your master's degree, you have spent a lot of time traveling between China and the Netherlands for me. Thank you for the journey you have accompanied me. This journey will continue and we will go together in the future. I also want to thank my father-in-law and my

mother-in-law, Xueying Li and Xiaohuan Yang. Every time we go back to your home, you always take care of me better than Xing.

Last but not least, I would like to express my sincerest gratitude to my parents, Rongcai Zhao and Zhaolan Zhang. Thank you for giving me full love, trust, understanding and tolerance, and always providing me assistance as far as you can.

The text is sometimes so pale that I cannot express all my gratitude. Finally, I would like to use my most sincere heart to thank all my friends, relatives and teachers who have not been mentioned before. Thank you for your support, encouragement, care, and companionship. I am so lucky that I can meet all of you.

Yue Zhao
March 2019
Delft, the Netherlands

Contents

1	Introduction	1
1.1	Learner Engagement	4
1.2	Research Questions	6
1.2.1	Long-Term Behavioral Engagement	6
1.2.2	Mid-Term Behavioral Engagement	7
1.2.3	Short-Term Behavioral Engagement	8
1.3	Contributions	9
2	The Change of Learner Behavior after Certificate Achieving	11
2.1	Introduction	13
2.2	Background	14
2.3	MOOC Datasets	16
2.4	Methodology	19
2.4.1	Concept Definitions	19
2.4.2	From Concepts to Analyses	20
2.4.3	Definition of Behavior Patterns	21
2.5	Results	23
2.5.1	Pre/Post-Passing Behaviors	23
2.5.2	Learners' Core Behavior Patterns	31
2.6	Conclusions	33
3	Mobile vs. Stationary Learning	35
3.1	Introduction	36
3.2	Background	37
3.2.1	Multitasking and Divided Attention	37
3.2.2	Mobile Learning	38
3.3	Study Design	40

3.3.1	Learning Situations	40
3.3.2	Learning Materials	40
3.3.3	Environmental Conditions	41
3.3.4	User Study Steps	42
3.3.5	Metrics	43
3.3.6	Study Participants	44
3.4	Results	44
3.4.1	Learning Gain, Efficiency and Interactions	44
3.4.2	Learning and Perceived Workload	47
3.5	Conclusions	49
4	Eye-Tracking Based Inattention Detection	53
4.1	Introduction	55
4.2	Background	58
4.3	Methodology	59
4.3.1	Study Setup	60
4.3.2	Inattention Detection with Gaze Features	62
4.3.3	Research Questions	64
4.4	Results	65
4.4.1	Exploratory Analysis of Inattention Reports	65
4.4.2	Inattention Detection	65
4.5	Conclusions	70
5	Face-Tracking Based Inattention Detection	73
5.1	Introduction	74
5.2	Background	75
5.3	Eye/Face-Tracking Frameworks	77
5.3.1	WebGazer.js	78
5.3.2	Tracking.js	78
5.3.3	Detecting Face-Miss Events	79
5.4	Methodology	81
5.4.1	Study Setup	83
5.4.2	Detection Accuracy	84
5.5	Results	86
5.5.1	Accuracy	86
5.5.2	Reaction Times	86
5.5.3	Background as an Influencing Factor	87

5.6	Conclusions	87
6	Near Real-Time Inattention Detection Widget	89
6.1	Introduction	90
6.2	Related Work	91
6.2.1	Attention Loss in Learning	91
6.2.2	Automatic Detection of Attention Loss	92
6.2.3	MOOC Interventions	93
6.3	<i>IntelliEye</i>	94
6.3.1	Architecture	94
6.3.2	User Interface	100
6.4	MOOC Setting	100
6.5	Empirical Evaluation	103
6.5.1	Technological Capabilities	103
6.5.2	Acceptance of <i>IntelliEye</i>	105
6.5.3	Impact of <i>IntelliEye</i>	107
6.6	Conclusions	109
7	Conclusion	111
7.1	Summary of Contributions	112
7.1.1	Long-Term Behavioral Engagement	112
7.1.2	Mid-Term Behavioral Engagement	113
7.1.3	Short-Term Behavioral Engagement	114
7.2	Future Work	115
7.2.1	Personal Analytics in MOOCs	115
7.2.2	Adaptive Learning in MOOCs	117
7.2.3	Multimodal Learning Analytics in MOOCs	118
	Bibliography	121
	List of Figures	137
	List of Tables	141
	Summary	145
	Samenvatting	147

Curriculum Vitae

151

Chapter 1

Introduction

Massive open online courses (MOOCs) have attracted extensive attention from learners, educators, and education institutes since 2012 (known as the year of the MOOC [85]) [28, 18, 32, 87, 114, 75]. MOOCs have thus become one of the most prominent examples of technology-enhanced learning. Based on the definition in [87, 75], MOOCs refer to online courses which can be accessed by a massive number of learners with internet connections anytime and anywhere, and there are no entry qualifications and charges for the access of all course materials¹. UNESCO [87] treats MOOCs as an important tool to achieve the 4th Sustainable Development Goal (SDG 4), *Ensure inclusive and equitable quality education and promote lifelong learning opportunities for all*, as set by the United Nations in 2015. Class Central [114], whose annual reports on MOOCs are referenced by both UNESCO [87] and OECD [75], reports that by the end of 2017, there were around 9,400 MOOCs provided by more than 800 universities and companies online, which had attracted 81 million learners—it is equivalent to around 40% of the total number of students in higher education institutions around the world in 2014 [133].

However, a low completion rate—compared to the large number of learners registered in MOOCs, only a small percentage of them got scores higher than or equal to the course requirements—is a ubiquitous and severe problem in MOOCs. In [54], it is reported that the median completion rate was 6.5% among 39 MOOCs across different premier MOOC platforms (e.g. Coursera²,

¹Nowadays, for the sustainability of MOOCs, some MOOC platforms (e.g. Coursera, edX, Udacity, and FutureLearn) charge learners for the enrolment of some MOOCs for professional development and provide fee-based certification for the course completion [75].

²<https://www.coursera.org/>

edX³, and Udacity⁴) from 2011 to 2013. Inspired by this study, we investigate 32 popular MOOCs provided by Delft University of Technology on edX from 2013 to 2017. There are 50 runs in total, as some MOOCs ran repeatedly in these years. In each run, at least 1,000 learners were registered and the average number of registered learners was 14,841. As shown in Figure 1.1, we find that the completion rates of most MOOCs were lower than 5%.

In this thesis, we investigate *learner engagement*, since it is highly related to the completion rates of MOOCs [28, 32]. For learners, engagement is considered as a necessary prerequisite for effective learning in MOOCs [41]. For MOOC providers, maintaining and cultivating learner engagement help them to make their impact broadly [95]. In traditional classroom contexts, experienced educators can observe learner engagement and keep learner engagement by adjusting course content and the way they teach. However, due to the properties of this MOOC technology and its nature of asynchronous interactions between educators and a large number of learners, educators cannot observe learner engagement in MOOC learning the same way they usually do in traditional classrooms. Without support from educators, current MOOC contexts, which revolve around a large number of videos and automatically graded questions, require learners to be skilled in *self-regulated learning* [92] (e.g. to plan their learning, monitor their learning progress, or keep their focus during learning by themselves). Many learners lack such skills and cannot keep their engagement across a course, even in a single learning session, which leads to high dropout rates of MOOCs.

Learning analytics technology has been used by educators and researchers to not only observe learners in MOOC learning but also provide feedback about their learning progress based on large-scale data generated from learner interactions on MOOC platforms [102]. For example, to explore learners' registration and enrollment for the first-year MOOCs in HarvardX and MITx, Ho et al. [49] analyze large-scale data about learners' certification, demographic information (e.g. gender, age, or academic degree), geographic information, and click activities which are collected from 597,692 users in 17 MOOCs. To understand learner interactions with different course components (e.g. quiz question, lab, video lecture, tutorial, book, discussion, or wiki), Breslow et al. [12] dig into the first MOOC "Circuits and Electronics" on edX based on large-scale data with 230 million learner interactions from about 155,000 learners. To motivate learners, Davis et al. [27] build a feedback system for MOOC learners based on learning analytics technology

³<https://www.edx.org/>

⁴<https://www.udacity.com/>

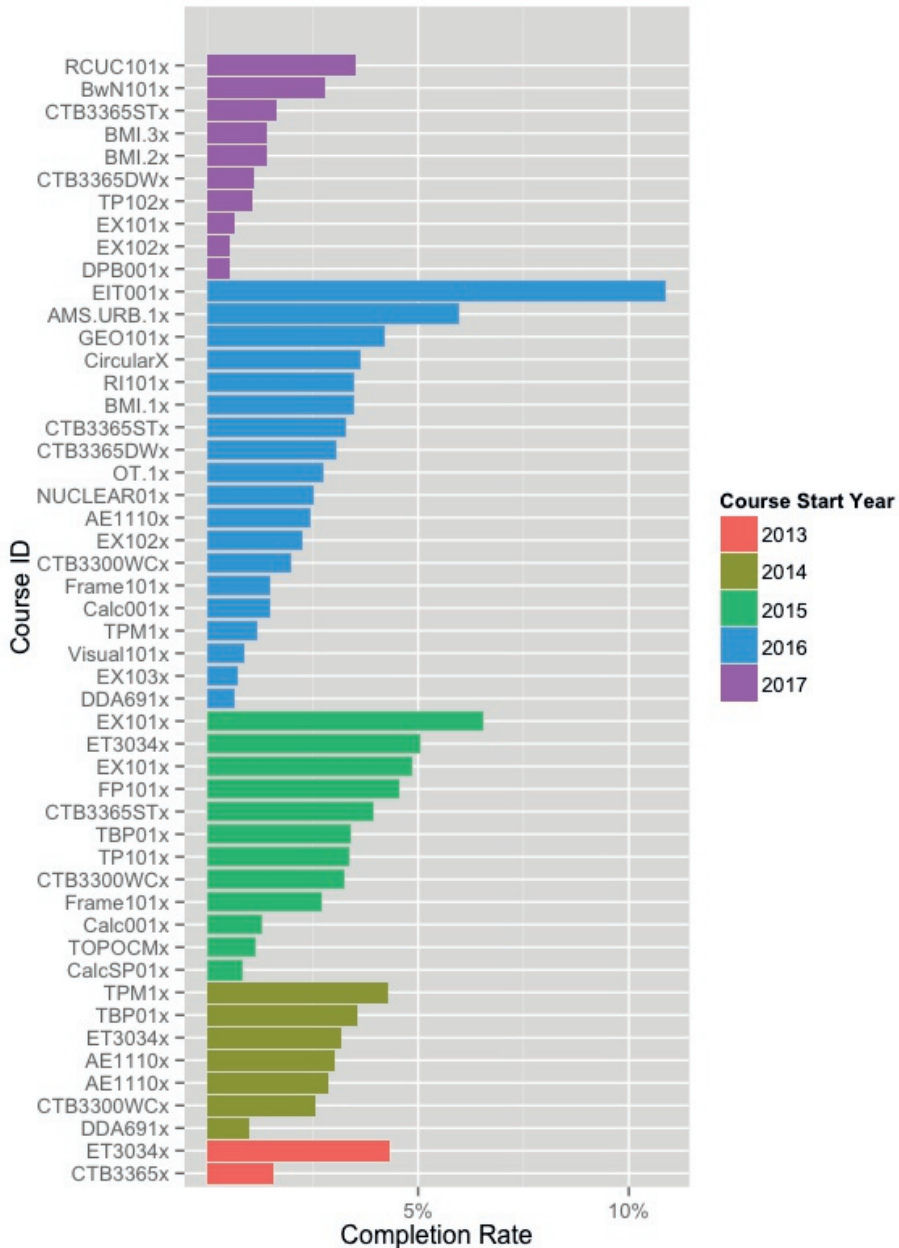


Figure 1.1: The completion rates of MOOCs provided by Delft University of Technology on edX from 2013 to 2017. Some MOOCs ran repeatedly in these years.

employed on their trace-data. Their system compares the learning progress of MOOC learners with previous learners who completed the course in the past and provides weekly feedback for learners. In this thesis, we also make use of learning analytics technology employed on data generated by learners to specifically consider learner engagement in MOOC learning.

1.1 Learner Engagement

In this thesis, we focus on learner engagement in MOOC learning, since it is commonly presumed to be essential to the success of learning [20, 1, 33, 52, 35, 41]. However, regarding the definition of learner engagement, there is little consensus among researchers in previous studies [52, 35]. In this section, we first clarify which kind of learner engagement in MOOC learning we investigate in our study.

As pointed out by Fredrick et al. [35], *the attempt to conceptualize and examine portions of the literature under the label "engagement" is potentially problematic; it can result in a proliferation of constructs, definitions, and measures of concepts that differ slightly, thereby doing little to improve conceptual clarity.* Therefore, three dimensions are usually used for understanding learner multidimensional engagement in both traditional classrooms [52, 35] and in MOOCs [48]:

Behavioral engagement: refers to the participation of learners in learning. Previous studies in traditional classrooms investigate behavioral engagement based on learner attendance in the course [121, 20, 33], learner performance of course assignments [20, 1, 33], or learner attention to the course [121]. Most studies about learner engagement in MOOCs are about behavioral engagement; learning analytics technology with learner trace data is used in these studies [56, 41, 19, 95, 90, 142]

Emotional engagement: refers to the feeling learners have about learning. Previous studies in traditional classrooms investigate learner emotional engagement based on learners' reports about their boredom, happiness, anxiety, or anger in the classroom and in the school [121, 20]. In MOOCs, previous studies investigate emotional engagement based on learners' facial expression during learning [93], their posts on forums [138] or in-person interviews after learning [24]

Cognitive engagement: refers to cognitive strategies that learners employ in learning. Previous studies in traditional classrooms investigate learner

cognitive engagement based on learners' memorization of course content [134], strategies used for regulating themselves during learning (e.g. monitoring and regulating cognition) [70, 106, 47], or the task-specific thinking during learning [47]. In MOOCs, cognitive engagement of learners is studied based on the content of their posts on forums [139, 136].

In this thesis, we focus on behavioral engagement of MOOC learners based on learning analytics technology. Many activities related to emotional engagement and cognitive engagement of learners (e.g. hesitating to study course content, or making plans for their learning) happen outside of MOOC platforms, while most activities related to their behavioral engagement are recorded by the technology of MOOC platforms in the form of a large amount of trace data which can be studied comprehensively and at different scales.

Behavioral engagement can be studied on different time scales in MOOCs. For example, Kizilcec et al. [56] study learner persistence in a complete course and group learners based on their engagement in each assessment period, while Guo et al. [41] investigate learner engagement in a video lecture based on learners' video watching time and whether they attempted follow-up problems in the same learning session. As mentioned by Fredrick et al. [35], *Engagement can vary in intensity and duration; it can be short-term and situation specific or long-term and stable*. However, Fredrick et al. do not clearly evaluate learner engagement on different time scales. In this thesis, we evaluate behavioral engagement of MOOC learners on three different time scales:

Long-term behavioral engagement: refers to learner behavioral engagement throughout a course. The large-scale trace data collected from a large number of learners provides an opportunity to measure behavioral engagement of learners based on their long-term interactions with course materials. Therefore, *how to explore the engagement of MOOC learners throughout a course based on learning analytics technology with large-scale trace data* is the main challenge.

Mid-term behavioral engagement: refers to learner behavioral engagement in a learning session. A learning session usually lasts from several minutes to several hours in which learners have continuous interactions with course materials. Studies on behavioral engagement of learners in a learning session can be more fine-grained and controllable than across a course. Regarding that learner behavioral engagement might be affected by different factors (e.g. prior knowledge to course contents, their preference to course

contents, or environments where they study), our main challenge to mid-term behavioral engagement is *how to measure the impact of different factors on behavioral engagement of MOOC learners in a learning session.*

Short-term behavioral engagement: refers to learner behavioral engagement in a short period of time (≤ 30 seconds in our studies). If behavioral engagement of MOOC learners can be measured on a large scale and in real-time, interventions can be provided once learners are disengaged. To measure short-term behavioral engagement of learners, the main challenge is *how to track learner behavioral engagement in a real-time and scalable way.*

1.2 Research Questions

In this section, we present our research questions about behavioral engagement of MOOC learners with different time scales.

1.2.1 Long-Term Behavioral Engagement

To study long-term behavioral engagement of MOOC learners, previous studies [56, 19, 95] mainly focus on high dropout rates of learners and cluster learners into different engagement patterns based on analyzing their activities on videos and questions in each week. In our investigations, we are interested in the change of long-term behavioral engagement of learners across a course. Previous studies in traditional classroom contexts reveal that learners are strategic and tend to spend most of their efforts on course content that (they believe) is being assessed [105, 38]. Based on the current MOOC assessment setting, in many MOOCs learners can pass the course when their scores reach the threshold (50 – 70% typically) and before accessing the last part of course contents. Since the remaining contents after passing have no contribution to the certification, this “passing” event gives an opportunity to observe the change of long-term behavioral engagement of learners before and after passing the course. If learners show different behavior patterns before and after passing, educators and course providers should consider this change in the course design and the grading schema of MOOCs. Otherwise, the course contents in the last part of a MOOC might be in vain for learners. Therefore, our **Research Questions** to investigate long-term behavioral engagement of MOOC learners are:

RQ 1.1: Do MOOC learners behave differently after clinching a passing grade?

RQ 1.2: What are the core behavior patterns of MOOC learners before and after passing, and how can learners be classified?

1.2.2 Mid-Term Behavioral Engagement

To investigate mid-term behavioral engagement of MOOC learners, we specifically focus on their participation in a learning session with mobile devices. With the increasing popularity of mobile technology, smartphone ownership has already surpassed the ownership of desktop and laptop computers. In 2015, about 86% of Americans in ages 18-29 owned smartphones while 78% of adults under 30 owned a laptop or desktop computer (which was 88% in 2010) [3]. MOOC learning can be conducted on mobile devices on many well-known MOOC platforms (e.g. edX, Coursera, and Udacity) by 2015 [66].

Mobile learning provides a scenario in which learners study MOOCs on-the-go and cannot be fully engaged in a learning session. MOOC learners usually study course materials while being stationary in a comfortable environment (e.g., sitting in the office or at home) where they can be fully engaged in learning (*stationary learning*). However, when learners learn on-the-go with mobile devices, they have to use smaller screens in various and possibly changing environments (*learning on-the-go*). It leads to an increase of interruptions and distractions [117], cognitive load [129, 16, 29], and frustration [23]

Existing studies on mobile learning in MOOCs mainly focus on the design and delivery of course content for mobile devices [97, 66] and the learning experience on mobile devices [90, 142, 23] which is typically studied in the lab, rather than real-world environments. Thus, little is known by educators and researchers about how divided engagement and real-world environments affect MOOC learning on-the-go compared to stationary learning. To observe the impact of mobile learning on learner engagement in a learning session, learning analytics technology on trace data can be used to measure behavioral engagement based on learner performance and interactions. Therefore, our **Research Questions** for mid-term behavioral engagement of MOOC learners in mobile learning are:

RQ 2.1: To what extent does learning on-the-go (compared to stationary learning on a mobile device) affect MOOC learners' learning gain, learning efficiency and interactions with the course content?

RQ 2.2: How do learners perceive their workload (physical as well as mental) in the stationary and learning on-the-go conditions and how does it relate to their learning performance and interactions?

1.2.3 Short-Term Behavioral Engagement

To study behavioral engagement of MOOC learners in a short time, we specifically focus on learners' attention during video watching. Many of today's MOOCs are centered around video lectures, and learners lose their attention frequently during video watching without realizing it [98, 128, 63]. Due to the use of digital display devices, there are a significant group of learners with "heavy media multitasking" behaviors. It is hard for them to focus on video watching while learning [63].

If the loss of attention within video watching can be detected automatically and in real-time, interventions can be provided to MOOC learners once they are being disengaged. To detect MOOC learners' inattention during video watching, we require an approach that is *scalable* (it can be deployed to thousands of learners), *near real-time* (inattention is detected as soon as it occurs), *unobtrusive* (learners are not distracted by the detection procedure) and *autonomous*. An ideal method is to track learners' inattention with an ordinary webcam. Therefore, in our study, our **Research Questions** for short-term behavioral engagement of MOOC learners are:

RQ 3.1: How often do MOOC learners experience inattention within video watching?

RQ 3.2: How well do our webcam-based inattention detection methods perform?

RQ 3.3: To what extent is MOOC learners' hardware capable to enable the webcam-based inattention detection?

RQ 3.4: To what extent do MOOC learners accept our inattention detection technology that is designed to aid their learning but at the same time is likely to be perceived as privacy-invading (even though it is not)?

RQ 3.5: What impact does the webcam-based inattention detection have on learners' behaviors and to what extent does it affect learners' video watching behaviors?

1.3 Contributions

To answer **RQ 1.1** and **RQ 1.2**, in **Chapter 2** we present a data-driven approach for understanding long-term behavior patterns of MOOC learners based on large-scale trace data. We analyze trace data from 4,000 MOOC passers in four different MOOCs. A number of pre-passing and post-passing behavior patterns are defined in our study and we find the majority of learners to fall into a narrow band of behaviors independent of the specific MOOC under investigation. We also find that a certain subset of learners heavily reduced their engagement in question answering after clinching a passing grade. These findings suggest course designers and educators to refine their course structure and grading schema which require learners to display mastery of an entire course subject before earning a certificate. To our knowledge, this analysis has been the first to focus on the event of passing and the impact of this event on behavioral engagement of MOOC learners. This study is published in the *ACM Conference on User Modeling, Adaptation and Personalization* [145].

To answer **RQ 2.1** and **RQ 2.2**, in **Chapter 3** we analyze learners' trace data and their data collected from questionnaires. A study is conducted with 36 learners based on their 30-minute mobile MOOC learning while sitting in the lab and walking in the real-world environment (not in the lab). We find that the necessity to multitask and divide attention while learning on-the-go on mobile devices, as well as changing environmental conditions contributed to lowered learning performance (7% less) from MOOC videos. We also find that learners spent a different amount of time on video watching between sitting in the lab and walking with learning. This study is published in the *European Conference on Technology-Enhanced Learning* and the *ACM Conference on User Modeling, Adaptation and Personalization* [149, 148].

To answer **RQ 3.1** and **RQ 3.2**, we first design a user study with eye tracking (in **Chapter 4**). We conduct the lab study with 13 participants to collect their inattention report and full set of gaze data from both the webcam and the professional eye-tracker. This study is the first precursor study for real-time webcam-based attention tracking in MOOCs, which indicates that a large-scale application of the webcam-based inattention detection in MOOCs is indeed possible. This study is published in the *European Conference on Technology-Enhanced Learning* [147].

Since the methods with eye tracking tend to have a high detection lag, can be inaccurate, and are complex to design and maintain, we propose another method with face tracking to answer **RQ 3.2** (in **Chapter 5**). We conduct an

extensive study with 20 participants involving two open-source browser-based software frameworks for gaze and face detection. As our second precursor study for real-time webcam-based attention tracking in MOOCs, a benchmark suite of 50 typical MOOC learner activities related to their attention and the loss of attention is compiled. Our evaluation on this benchmark suite reveals that the face-tracking method shows significantly higher performance for nearly all benchmark tasks than eye tracking. Moreover, the observed detection delay of the face-tracking method is below 2 seconds, making it manageable for the near real-time detection in MOOCs. This study is published in the *ACM Conference on Intelligent User Interfaces* [100].

In **Chapter 6**, to answer **RQ 3.3**, **RQ 3.4** and **RQ 3.5**, we implement IntelliEye, a near real-time webcam-based attention tracking widget which is privacy-aware⁵ and scalable. IntelliEye was deployed in a real MOOC across a period of 74 days. We find that most learners (78%) used hardware and software setups which were capable to support such widgets, making the widespread adoption of our approach realistic from a technological point of view. Around 32% of learners with capable setups were willing to allow the use of webcam-based attention tracking techniques. Among the learners using IntelliEye, we observe (i) high levels of inattention and (ii) an adaptation of learners' behavior towards the attention tracking technology. This study is published in the *ACM Conference on Hypertext and Social Media* [99].

⁵There is no image/video from learners' webcams transmitted over the network.

Chapter 2

The Change of Learner Behavior after Certificate Achieving

In this chapter, we introduce our study on long-term behavioral engagement of MOOC learners across a course. In this study, long-term behavioral engagement refers to the participation in a complete course from the first part of course contents to the final part. This study is intended to serve as a foundation for designing systems which allow tracking some aspects of long-term behavioral engagement of MOOC learners in a scalable, unobtrusive, and generalizable fashion. Based on the tracking of long-term behavioral engagement of current learners, educators and course designers can adjust course contents to maintain the engagement of subsequent learners. To this end, we intend to make use of large-scale trace data of detailed learner activities in real-world MOOCs. Such trace data is already recorded in daily logs of MOOC platforms. To make the method scalable, unobtrusive, and generalizable, we do not consider further data sources which incur additional overheads like questionnaires (e.g. our study in Chapter 3), or extra sensing technologies (e.g. our study in Chapter 6)

Due to the nature of trace data, we focus on aspects of long-term behavioral engagement which are linked to the active participation and interactions with course contents like video lectures and quiz questions. We are particularly interested in changes of this participation over time for different learner populations. In traditional classrooms, there is an observation that most learners tend to selectively neglect course contents that (they think) are not

being assessed. In many MOOCs, learners can pass the course when their scores reach the course requirements and before the course ends based on the current MOOC assessment setting. It means that the remaining course contents have no contribution to the certification of learners who already passed the course (named as *passers*). Therefore, we explore how “passing” impacts MOOC learners: *do learners alter their behavior after this point?* And if so how? While in traditional classroom contexts the role of assessment and its influence on learning behavior has been well-established, we provide answers to these questions in the context of MOOCs, providing valuable insights which can be used to design better courses in the future.

This chapter is published as “Certificate Achievement Unlocked: How does MOOC learners’ behavior change?” [145], by Yue Zhao, Dan Davis, Guanliang Chen, Christoph Lofi, Claudia Hauff and Geert-Jan Houben, in Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization, pp. 83-88. ACM, 2017.

2.1 Introduction

For decades, researchers in the learning sciences have explored how the assessment of learning shapes learning strategies and behaviors of learners in the classroom [124, 71, 67]. One commonly encountered phenomenon, especially in higher education, is learners' adaptation of their learning strategies to the specific assessment tools: while some assessment choices such as multiple-choice questions are driving learners towards *surface learning* strategies (that is, learners aim to maximize recall of the material) other assessment types including essay writing are more likely to lead to *deep learning*, meaning learning that focuses on understanding [111]. Despite this knowledge, many MOOCs today rely to a large extent on a continuously distributed set of multiple choice questions for assessment, due to their inherent scalability (through auto-grading) to very large groups of learners. To illustrate this issue, we manually inspect all 46 university-level computer-science MOOCs² offered on the edX platform in October 2016 according to their assessment type as shown in Table 2.1: 73% rely on multiple-choice questions conjointly with some other assessment technique, while 24% exclusively use only multiple-choice assessment without additional evaluation techniques. Only one course abstains from using any kind of multiple choice assessment.

Assessment is a concept closely related to *learner effort* as learners tend to spend most of their learning efforts on course concepts that (they know) are being assessed [105, 38]. Educational researchers have long advocated for the even distribution of learner effort across topics and course weeks [37]. Once again, MOOCs tend not to follow this basic guideline as shown in Table 2.1: most MOOCs (31 out of 46 to be exact) can be passed after reaching less than 60% of the total score before the end of courses.

Classroom-based learning bears only a passing resemblance to MOOC learning for a variety of reasons including the scale, the heterogeneity of the learner group [40] with respect to age, educational and cultural background as well as the issues of isolation and remoteness that learners face [42]. It is thus an open question, whether the classroom-based findings of assessment and their influence on learning behaviors hold in MOOCs. In this chapter, we answer this question by empirically exploring *to what extent* MOOC learners' behaviors are impacted by one particular assessment event: the *course passing* event (i.e. the moment the learner accumulate sufficient scores to receive a certificate), which—depending on a MOOC's design—may potentially occur as early as half-way through the course. Furthermore, we generalize our

²We choose this category as it is popular on the edX platform.

findings into core learner behavior patterns, and provide an effective technique for classifying learners with respect to those patterns. In summary, we address two research questions in this chapter:

RQ 1.1: Do MOOC learners behave differently after clinching a passing grade?

RQ 1.2: What are the core behavior patterns of MOOC learners before and after passing, and how can learners be classified?

To this end, we analyze the log traces (our *observable* events from which to infer learning behavior) of 4,000 MOOC learners in four different edX MOOCs that earn a course certificate. Besides the scientific curiosity that underlie these questions we also believe the outcomes of this study will significantly further the discussion on MOOC course designs: Understanding and modeling learner behaviors is a prerequisite for designing MOOCs with adaptive features.

Table 2.1: Overview of the summative assessment type(s) and average passing threshold τ_{pass} of all 46 computer science & programming MOOCs (in English, geared at undergraduates) open for enrolment on the edX platform on October 15, 2016. Assessments types include multiple choice (**MC**), fill-in-the-blank (**FB**), code submissions (**CS**), peer reviews (**PR**) and discussions (**D**). The column **Early Passing** shows the number of courses learners can pass before the final assessment is released.

Assessment Type(s)	#MOOCs	#Early Passing	Avg. τ_{pass}
MC+FB	13	12	50.0%
MC	11	7	59.1%
MC+FB+CS	11	9	52.3%
MC+FB+PR	4	3	57.5%
MC+FB+CS+PR	3	3	63.3%
MC+PR	1	1	70.0%
CS	1	0	65.0%
MC+CS	1	1	50.0%
MC+FB+D	1	1	50.0%

2.2 Background

The impact of assessment on learners' learning and behaviors has long been a topic of research in the education literature [124, 71, 67]. Such studies emphasize the role of assessment as an influence on the learning process, specifically on the manner by which learners elect to engage with ensuing course content. As pointed by Gibbs and Simpson, assessment has "an overwhelming influence on what, how and how much students study" [39].

The impact of assessment on learning behavior manifests itself in a multitude of ways. Newble and Jaeger [78] report that the changes in exam type (rote memorization-based versus application of conceptual knowledge) in a medical school led to changes of learners' exam preparation. The most notable change was in their choice of study location; rote memorization-based exams drove learners to spend a disproportionate amount of time in the library, whereas the concept application-focused exams led learners to prepare and study in hands on environments such as laboratories. Natriello and Dornbusch [77] finds that assessments with higher standards for mastery lead to learners exerting more effort towards the course. Sambell and McDowell [108] report that learners build their own curriculum based on their experience and the types of assessments. A case study by Cooks [21] shows that the change from fixed assessment to flexible assessment (where learners could each pick their own grading scheme for the course) affected not only learners' behaviors but also their emotions in the way they approached exams.

Other work has found that learners engage with assessed course content differently than they do with unassessed content (e.g. the dreaded "Will this be on the test?" question). For course content expected to be unassessed, learners might be "selectively negligent" [38] or "do it in a perfunctory way" [105]. Forbes and Spence [34] finds that learners stopped doing their weekly problem sheets when the teachers were too busy to grade their work. Peer-assessment was evaluated as a potential solution and led to increases in learners' engagement levels and higher final exam grades than teacher-graded assessment.

Extrapolating these findings to MOOCs, we expect this behavior change on assessed and unassessed content to manifest itself similarly with regard to learner engagement before and after reaching a passing grade. Activity that happens after learners have clinched a passing grade is not required and therefore enables us to examine the learners' motivation for the course.

Kovacs [59] studies how in-video quizzes affect learners' behaviors in viewing lecture videos, but this study only focuses on short-term behaviors around in-video quizzes. Whereas Kovacs [59] focuses specifically on behavior within videos containing quiz questions, such as seeking behaviors or quiz-driven video navigation strategies, the present research differs in that we chiefly consider behavior on a course-long scale and how that is affected by the attainment of a passing grade.

In our work, we develop a novel typology of MOOC learner behavior to classify learners into different groups based on their behavior patterns before and after "passing", or clinching a passing grade in a given course.

Once we are able to use proper learner stereotypes to represent MOOC learners’ longitudinal behaviors, we can build adaptive applications that use learner stereotypes in MOOCs, inspired by e-learning systems that did this before. Kizilcec et al. [56] make a first contribution to learner stereotypes in MOOCs, even though they do not consider the impact of assessment on MOOC learners’ longitudinal behaviors. In our work, inspired by previous studies on learners’ behaviors on assessed content and unassessed content, we focus on the change of learners’ behavior before and after “passing” and classify learners based on their pre-passing and post-passing behavior patterns.

2.3 MOOC Datasets

We analyze the log trace data of 4,000 learners who successfully completed one of four MOOCs offered on the edX platform—they are summarized in Table 2.2. Each course is set up as an xMOOC [101] with weekly releases of lecture videos and graded³ quizzes. The quizzes are composed of automatically assessed multiple choice and fill-the-blank questions, and none of the MOOCs have a final exam. The assessment is exclusively based on the scores learners reached in the graded quizzes. In each MOOC learners can continuously check their scores by accessing their course “Progress” page.

For three of the MOOCs (FP, DA and SEW) the passing threshold is $\tau_{pass} = 60\%$, for SE it is $\tau_{pass} = 58\%$. Previous work [25] has shown that learners who pass a MOOC do follow the designed learning path of the course much closer than learners who do not pass. Thus, we can assume that the temporal sequence of course activities passers follow is in line with the design of the course.

As the distribution of possible scores shows in Figure 2.1 (and with $\tau_{pass} = 58\%$ and $\tau_{pass} = 60\%$ in mind), all four MOOCs can be passed well before the final unit.

In Figure 2.2 we plot the *total* number of learners who earned a certificate by the end of each (weekly) unit⁴—starting at the first possible certificate-earning unit. We make two key observations from this sample of courses: (1) many learners earn the certificate at the earliest opportunity—for both FP

³Although some ungraded quizzes exist as well, we ignore them in this analysis, as only activities on graded quizzes bring learners closer to the passing threshold.

⁴To be precise shown in Figure 2.2 are the **Unit-n passers** as defined in the upcoming sub-section **Concept Definitions**.

Table 2.2: Overview of the four MOOC datasets analyzed. **#Questions** refers to graded questions in MOOCs. **#Attempts** refers to how many attempts a learner has for each graded question. **#Engaged** refers to learners who watched at least one lecture video or answered one question. **Compl Rate** refers to the completion rate of a MOOC.

ID	Name	Start	End	Units	Videos	Questions	Attempts	#Learners			Compl Rate
								Registered	Engaged	Passed	
FP	<i>Introduction to Functional Programming</i>	10/2015	01/2016	8	39	285	1	25,188	9,900	1,143	4.54%
DA	<i>Data Analysis: Take It to the MAX()</i>	09/2015	11/2015	8	60	137	2	23,813	9,780	1,156	4.85%
SEW	<i>Treatment of Urban Sewage</i>	04/2016	06/2016	7	79	36	1	11,006	2,589	361	3.28%
SE	<i>Solar Energy</i>	09/2015	12/2015	8	61	128	1 – 3	26,718	12,723	1,346	5.04%




Figure 2.1: Overview of the fraction of scores that learners can earn in each unit. The passing threshold for SE is $\tau_{pass} = 58\%$, while it is $\tau_{pass} = 60\%$ for the other three MOOCs. Best viewed in color.

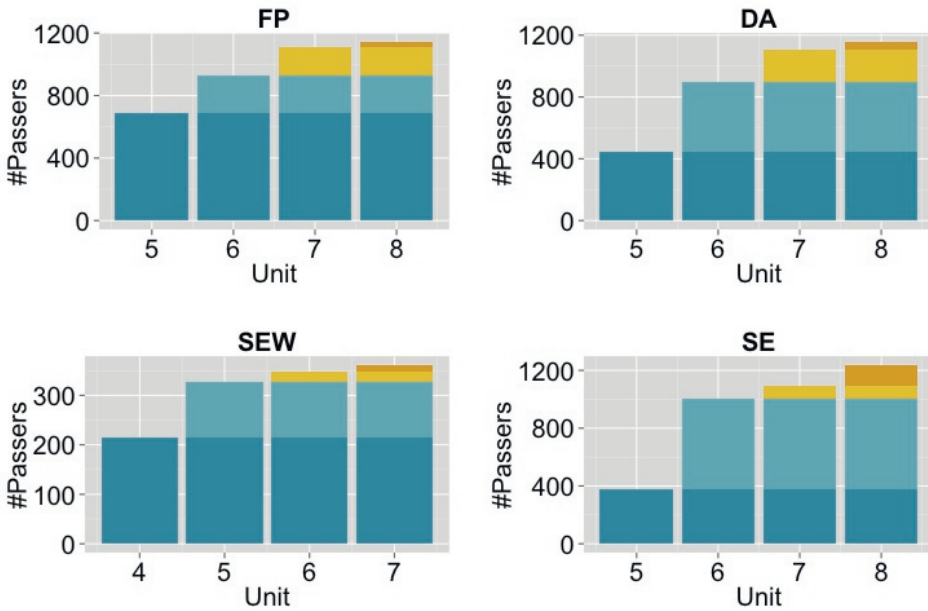


Figure 2.2: Total number of certificate earners (i.e. “passers”) at the end of each unit.

and SEW this is true for approximately 60% of the learners, for DA and SE it holds for 40% and 30% of the learners respectively; (2) only a very small minority of learners pass in the final two units.

2.4 Methodology

In this section, we first formally define the core concepts which we will use throughout our work and then describe how we conduct the analyses to answer our research questions.

2.4.1 Concept Definitions

MOOCs & MOOC units: A MOOC M consists of a sequence of m units, i.e. $M = (U_1, U_2, \dots, U_m)$. Each unit contains videos and/or quizzes and is typically designed to be completed over the course of one calendar week.

Unit- n quizzes & videos: According to [2, 56], there are two core components of xMOOCs⁵: (1) lecture videos, and (2) quizzes. Quizzes and lecture videos included in a weekly unit U_i are represented as $U_i = \{V_i, Q_i\}$.

Learner’s Activities: We consider quiz scores and time spent on videos as the main measurements for learner activity on a MOOC platform, i.e. for each learner l and MOOC unit $U_i \in M$, the normalized quiz score is denoted as Q_i^l . A learner’s l normalized time spent on watching the video of a given unit U_i is debited by V_i^l , where $V_i^l = 1.0$ represents watching the full length of all videos of a unit at normal speed. Thus, watching all videos twice results in $V_i^l = 2.0$, and skipping half of the videos and watching the remainder at double speed results in $V_i^l = 0.25$. We compute these normalized video watching times from analyzing all learner event log files from edX, extracting and aggregating all interactions with the video player component.

Passers: Passers P are learners who are eligible to receive a MOOC certificate at the end of the MOOC as their assessment scores reach the defined threshold τ_{pass} (independent of the unit they reach the threshold). In the present research, only these learners are considered.

Unit- n passers: Given τ_{pass} , unit- n passers P_n are passers whose achieved assessment scores reach at least τ_{pass} only considering units up to U_n and whose scores up to unit U_{n-1} are not sufficient, i.e.

$$P_n = \left\{ p \in P \mid \sum_{i=1}^{n-1} Q_i^p < \tau_{pass} \wedge \sum_{i=1}^n Q_i^p \geq \tau_{pass} \right\}$$

⁵xMOOCs are heavily relying on video lectures and quizzes to convey knowledge, in contrast to cMOOCs which rely on learners’ self-formed communities and peer teaching.

Note once more that the actual time the quizzes are completed by the passers can vary (a quiz released in unit n may be completed a week or two after its initial release). This, however, has little impact on our work as passers usually follow the predefined sequences of MOOC units [25].

Pre-passing activities: The pre-passing activities A_{pre}^p of a passer $p \in P_n$ include all quiz & video activities up to & including unit n .

Post-passing activities: The post-passing activities A_{post}^p of a passer $p \in P_n$ include all quiz and video activities starting in unit $n + 1$. A passer who passes in the final unit has no post-passing activity.

We denote the previously introduced concepts with the respective MOOC label when appropriate, e.g. P_{DA} or P_{FP} for referring to a specific passer group, or $Q_{5,DA}^p$ for referring to the quiz score of learner p for unit U_5 of the course DA.

2.4.2 From Concepts to Analyses

Recall that (in traditional classroom contexts) learners engage differently with assessed course content than they do with unassessed content [38, 105, 34, 51]. Applying this same concept to the MOOC context, we expect to observe a difference in the way learners behave before and after clinching a passing grade.

To address **RQ 1.1**, we operationalize *behavior* in this case as a learner’s engagement with course quizzes and video—the two most prominent activities in most MOOC settings [12, 112]. We then identify the unit in which they clinched a passing grade and group them accordingly. Finally, we plot the distribution of their quiz scores and video watching activity over time.

In the next step, we zoom in and explore the *individual* learner behavior. In order to determine whether behavioral changes can be observed on individual learners, we represent each passer p by a vector of her normalized quiz scores. Then, we resort to k-means clustering (also employed in [4, 56] for analyzing learners’ activities) of each unit- n passer group to cluster learners with similar feature vectors. We measure the distance between learner feature vectors by their Euclidean distance⁶. As we do not know in advance how many different prototypical learner behaviors exist (i.e., the best number of clusters is unknown), we generate multiple k-means clusterings with

⁶We also explored Dynamic Time Warping [135], a specialized distance function for time-series data—this did not yield a higher silhouette score.

$k = [1, 2, \dots, 7]$. For each of these seven clusterings, we assess the clustering quality using silhouette coefficients [104], an effective technique for assessing the quality of a clustering result. Our final clustering is the one with the highest silhouette score.

2.4.3 Definition of Behavior Patterns

Whereas the prior work in modeling learner behavior considers learner activity from the entire course duration [71, 131, 50, 56, 2, 137], we break the sequence into two parts: pre- and post-passing. This segmentation allows us to examine any effects or changes in behavior stemming from the attainment of a passing grade.

In addressing **RQ 1.2** we now conceptually define a number of behavior patterns (based on the literature and our own findings) and then classify our passers into their closest matching pattern. While the clustering just describe provide us with meaningful insights (as will become evident in §2.5.1), these clusters do not allow us to explicitly *model* learner behaviors.

Once more, we restrict ourselves to quiz-score based behavior patterns. Concretely, we define five pre-passing behavior patterns and six post-passing behavior patterns which we summarize in Table 2.3 and 2.4. We deliberately make this split of behavior prior to and after passing as the clustering results show (cf. §2.5.1) a divergent behavior almost exactly at the point of passing.

For pre-passing, we define: *keeping high scores* for learners who exhibit high quiz scores for all units before passing (which might indicate highly motivated or effective learners); *keeping mid scores* for learners analogously keeping medium scores (which might indicate reduced but constant motivation, or problems with the complexity of the topic); *raising scores* for slow starters who begin the MOOC with low scores, but then increase their scores until they pass; *reducing scores* for learners who start with high quiz scores which then steadily decline until the course is finally passed (this may be indicate of slowly waning motivation); and *unstable scores* which represents no clear behavior pattern (like achieving high scores in one unit, and then skipping the next unit all-together).

Analogously, we also define six post-passing behaviors. The behavior pattern we add over the pre-passing patterns is *keeping low scores* which is consistent low to zero scoring behavior after the passing stage (an impossibility in the pre-passing behavior pattern set, as we only consider learners that eventually pass).

Table 2.3: Overview of the five pre-passing behavior patterns for normalized quiz scores.

Pre-passing Behaviors	Explanations	Definitions
Keeping high scores	Passers start with high scores and keep high scores.	$Std_{pre}^p \leq t_{std} \wedge Avg_{pre}^p \geq t_{high}$
Keeping mid scores	Passers start with middle scores and keep middle scores.	$Std_{pre}^p \leq t_{std} \wedge Avg_{pre}^p < t_{high}$
Raising scores	Passers start with mid scores or low scores but increase scores to high scores.	$Std_{pre}^p > t_{std} \wedge (Slp_{pre}^p \geq t_{slope} \vee (Slp_{pre}^p > 0 \wedge LRE_{pre}^p < t_{tre}))$
Reducing scores	Passers start with high scores but reduce their scores to middle or low scores.	$Std_{pre}^p > t_{std} \wedge (Slp_{pre}^p \leq -t_{slope} \vee (Slp_{pre}^p < 0 \wedge LRE_{pre}^p < t_{tre}))$
Unstable scores	Passers' scores are not stable and do not show clear trends.	<i>learners who are not successfully assigned to a previous pattern</i>

Table 2.4: Overview of the six post-passing behavior patterns for normalized quiz scores.

Post-passing Behaviors	Explanations	Definitions
Keeping high scores	Passers keep high scores to the end.	$Std_{post}^p \leq t_{std} \wedge Avg_{post}^p \geq t_{high}$
Keeping mid scores	Passers keep middle scores to the end.	$Std_{post}^p \leq t_{std} \wedge t_{mid} \leq Avg_{post}^p < t_{high}$
Keeping low scores	Passers keep low scores or 0 scores to the end.	$Std_{post}^p \leq t_{std} \wedge Avg_{post}^p < t_{mid}$
Raising scores	Passer' scores show increasing trends to the end.	$Std_{post}^p > t_{std} \wedge (Slp_{post}^p \geq t_{slope} \vee (Slp_{post}^p > 0 \wedge LRE_{post}^p < t_{tre}))$
Reducing scores	Passer' scores show decreasing trends to the end.	$Std_{post}^p > t_{std} \wedge (Slp_{post}^p \leq -t_{slope} \vee (Slp_{post}^p < 0 \wedge LRE_{post}^p < t_{tre}))$
Unstable scores	Passers' scores are not stable and do not show clear trends.	<i>learners who are not successfully assigned to a previous pattern</i>

Having defined these patterns, we now manually classify all learners' exhibited pre- and post-passing behavior sequences into these patterns; this classification is crisp: each learner exhibits exactly one of the defined pre-passing patterns and one of the defined post-passing patterns. We manually determined the rules and the best setting of the threshold values by sampling a small number of to-be-classified quiz score series, hand-labeling them and creating rules and thresholds accordingly. The rules and thresholds are the same for all four MOOCs. The resulting rules are listed for each behavior pattern in Table 2.3 and 2.4. Our rules are based on analyzing average val-

ues, deviations, and linear regression of quiz scores. The rules contain the following additional concepts:

Average/std. normalized quiz score: A passer's p average and standard deviation normalized quiz score for all pre-passing and post-passing scores Avg_X^p , Std_X^p with $X = \{pre, post\}$.

Stability: Threshold t_{std} determines whether a given series of pre- or post-activity scores are stable based on the respective standard deviation. When $Std_X^p \leq t_{std}$, we classify p as keeping scores.

Level: Thresholds t_{high} and t_{mid} to indicate high or medium normalized quiz scores. If a learner's behavior was classified as "keeping" based on Std_X^p , these thresholds are used on the average Avg_X^p to determine the correct "keeping" class.

Score slope: A passer's p score slope Slp_X^p with $X = \{pre, post\}$ is the slope of the linear regression of the pre-passing or post-passing normalized quiz score series. LRE_X^p with $X = \{pre, post\}$ is the least squared error of that linear regression. We consider linear regression if a behavior is not "keeping scores" because of high standard deviation Std_X^p .

Instability: Threshold t_{lre} determines, based on the standard error LRE_X^p of the linear regression, if a behavior can be fit close enough for being considered raising/reducing, or if it should be unstable instead. For example, the score sequence (100%, 66%, 33%, 0%) has a very low regression error and is "reducing", while (100%, 0%, 66%, 33%) has a high regression error and is thus considered "unstable".

To classify a series of data points, we evaluate the patterns from top to bottom and the first pattern whose corresponding rule evaluates to *true* is considered as the pattern to classify the data points into. As visible in the rule definitions, the final pattern (*unstable*) ensures that all data points are classified into one of the patterns.

2.5 Results

2.5.1 Pre/Post-Passing Behaviors

Recall, that in **RQ 1.1** we are concerned with the question whether or not passers behave differently before and after having reached the passing threshold.

Observation Analysis


The distribution of quiz scores and video consumption for our learners grouped by passing unit are shown in Figure 2.3, Figure 2.4, Figure 2.5 and Figure 2.6. Here, each row shows the behavior of one group of passers (e.g. the top row in Figure 2.3(a) shows the quiz scoring activities of all unit-5 passers of FP) while each column shows the behavior of all passers in a particular unit (e.g. the last column of Figure 2.3(a) shows the behavior of all passers in unit 8).

Across all courses we find learners who pass in early units (top two rows in each sub-figure of Figure 2.3 and Figure 2.4) to score in a narrow range of high scores before passing—this is of course a prerequisite for passing early. After the minimum passing threshold is reached, however, the variance of scores increases drastically, with a number of learners starting to score very low. For example, 6% of $P_{5,FP}$ learners (i.e. learners who passed in week 5) score less than 20% of the available points in Q_6 and 22% of $P_{6,FP}$ learners (who passed in week 6) score less than 20% of the available points in Q_7 . In contrast to DA and SEW, in FP and SE we observe a larger number of learners who maintain high scores after passing than learners who score low after passing. Concretely for FP, in the final unit, more than two thirds of the $P_{5,FP}$ passers score 80% or higher on the remaining quizzes.


The video consumption behavior of passers across MOOCs is also noteworthy: in every MOOC a small to medium fraction of passers does not watch any⁷ of the unit’s videos —3.4% in FP, 3.0% in DA, 10.8% in SEW and 20.0% in SE. In Figure 2.5 and Figure 2.6, we report on the video watching behavior of all those passers with at least one video activity in our logs. Across the four courses the trend over time is similar: the number of passers who do not watch lecture videos increases in the final units. With respect to the completeness of lecture video consumption we find a clear divide between DA & SE and SEW & FP: in DA & SE learners’ normalized video consumption peaks around 1.0 (indicating that many learners watch the whole video lecture at normal speed), while in SEW & FP for most passers the normalized duration is below 1.0 indicating that they skip at least parts of the videos.

We can conclude that learner behaviors on quizzes are distinctive before and after passing. We also find (not unexpectedly) marked differences between the quizzing behavior of passers and not-yet-passers in the same unit. At the same time, we fail to observe the same clear differences in the

⁷We note that an alternative explanation for the zero peak may be that learners download videos for offline learning as suggested by [2], which is not captured in the edX logs. While this may be true for some learners, this cannot explain the change in behavior after the passing threshold is reached.

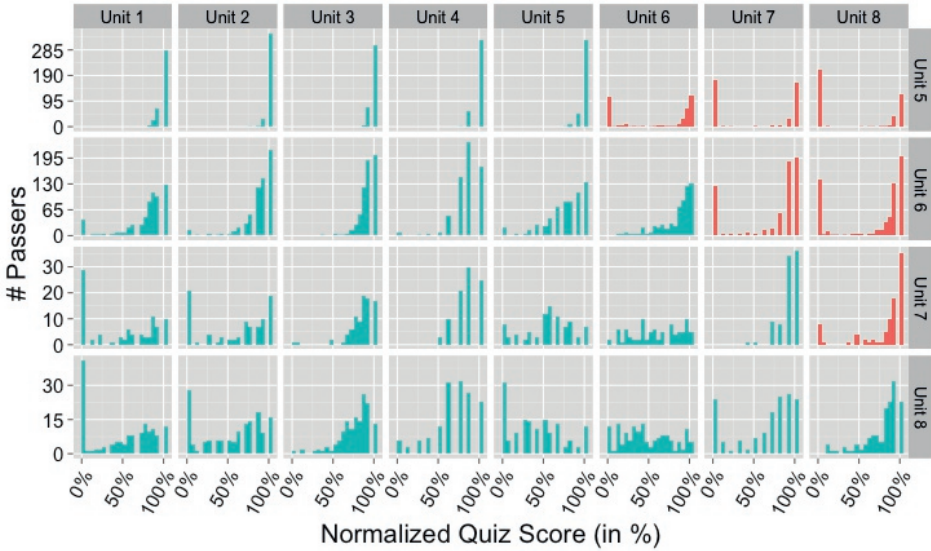


(a) FP




(b) DA

Figure 2.3: Quiz score distribution of FP and DA: passers are binned according to their passing unit. Rows represent groups of passers, columns represent one particular unit. Red plots show groups of passers that reached the passing threshold in a previous unit.




(a) SE




(b) SEW

Figure 2.4: Quiz score distribution of SE and SEW: passers are binned according to their passing unit. Rows represent groups of passers, columns represent one particular unit. Red plots show groups of passers that reached the passing threshold in a previous unit.

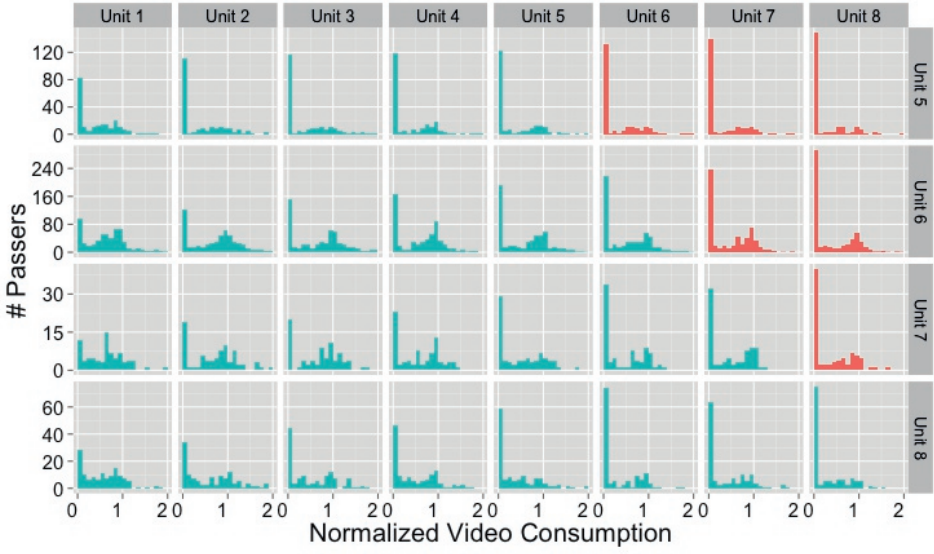


(a) FP




(b) DA

Figure 2.5: Video consumption distribution of FP and DA: passers are binned according to their passing unit. Rows represent groups of passers, columns represent a particular unit. Red plots show groups of passers that reached the passing threshold in a previous unit.



(a) SE



(b) SEW

Figure 2.6: Video consumption distribution of SE and SEW: passers are binned according to their passing unit. Rows represent groups of passers, columns represent a particular unit. Red plots show groups of passers that reached the passing threshold in a previous unit.

video consumption. Based on this result, in our further analyses we focus exclusively on passers quiz behaviors.

Clustering Analysis

Based on the clustering described in §2.4.2 we visualize the resulting normalized quiz score clusters in Figure 2.7 for the four courses: each unit in each cluster is represented by the average score learners in that cluster achieve in that unit with their respective confidence bands. The key insights of Figure 2.7 with respect to **RQ 1.1** are:

- For passers who pass MOOCs early (i.e. the first two unit-n passers groups), the clusters share very similar activity levels before passing, but begin to differ immediately at the passing unit.
- For nearly all unit-n passer groups and MOOCs, choosing $k = 2$ clusters yields the best clustering fit. This strongly indicates that for early passers, there are two dominant behavior patterns: “reducing scores” (rapidly declining quiz scores for the units following the passing event) and “keeping scores” (the averaged scores of passers in one cluster stay more or less stable at a high level) after passing.
- There are exceptions to the two-cluster rule: $P_{5,SE}$ and $P_{7,SE}$ split into many small clusters the latter can be attributed to the overall low number of learners to be clustered. The five clusters observed in $P_{5,SE}$ are explained by the special setup of SE with “exams” appearing in units 3, 6 and 8 which not only cover the material of the current unit but also of previous units. Passers of unit 5 fall into different clusters depending on whether or not they “take the exams” in units 6 and 8.
- The MOOCs differ in the dominant post-passing behavior, while for $P_{5,FP}$ and $P_{6,SE}$ the dominant cluster is “keeping scores”, in DA across all groups the “reducing scores” passers dominate over those that keep participating in the assessment (not shown here). This may hint at different motivations for taking the course (gaining knowledge vs. gaining a certificate).
- In $P_{7,DA}$ we also observe a behavior unique to DA: a group of learners starting off slowly (low scores in units 1 and 2) and finishing strong (high scores in starting in unit 3).

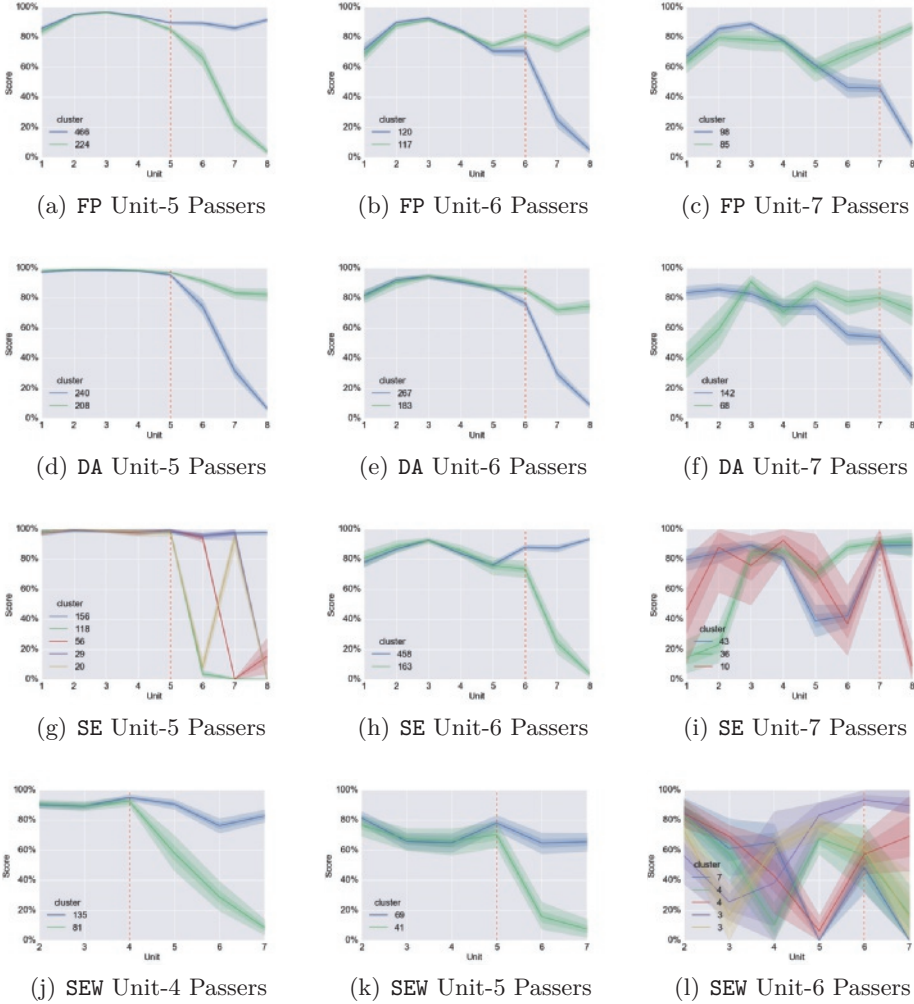


Figure 2.7: K-means clustering of learners normalized quiz score feature vectors for the first three unit- n passers groups (in SEW, learners' scores can reach τ_{pass} already in Unit 4). The cluster label in each graph shows the number of passers in each cluster. The vertical red line indicates the unit in which passers reached the passing threshold. The shaded areas around the lines show the upper (99%) and lower (70%) confidence bounds. Best viewed in color.

These results show that indeed, we find significant changes in learner behavior after the passing event. We conducted a similar analysis for video consumption, but as expected based on the observation analysis, we did not observe meaningful clusters or behavioral changes after passing.

2.5.2 Learners' Core Behavior Patterns

To answer **RQ 1.2**, based on rules we designed in §2.4.3, we now classify all *early passers*, i.e. those that pass within the first two possible units, into their pre- and post-behavior patterns. We restrict ourselves to early passers as (i) the vast majority of passers fall into this category (cf. Figure 2.2), and (ii) we have sufficient post-passing behavior data points for them. Based on our five pre-passing and six-post passing behavior patterns, we have potentially $5 \times 6 = 30$ different “course patterns” the passers fall into. In Table 2.5 we list for each of our four MOOCs what percentage of early passers fall into each of those patterns. We make the following observations:

- Two course patterns (*pre + post*) dominate across all four MOOCs which have already been hinted at in our previous analyses: (1) *keeping-high + reducing*, and (2) *keeping-high + keeping-high* with more than 50% of passers in $P_{4,SEW}$, $P_{5,DA}$, $P_{5,SE}$ and $P_{5,FP}$ respectively falling into one of those patterns.
- Similarly, as already indicated in the cluster analysis, in FP and SE more learners keep high scores after passing (46% of $P_{5,FP}$ passers and 38% of $P_{5,SE}$ passers) than reducing them (24% of $P_{5,FP}$ passers vs. 27% of $P_{5,SE}$ passers), while the opposite is true for DA (among $P_{5,DA}$ 60% reduce and 27% keep scoring high) and SEW (among $P_{4,SEW}$ 30% reduce vs. 20% that keep scoring high).
- Interestingly, among those early passers very few stop submitting assessments immediately after passing (indicated by the low percentages in the post-passing behavior *keeping low scores* rows). This phenomenon is most pronounced for the course pattern *unstable + keeping-low* with more than 10% of passers in three MOOCs falling into this category in $P_{6,DA}$, $P_{6,FP}$ and $P_{5,SEW}$.
- We find that learners whose pre-passing behavior was categorized as *keeping-high* are less likely to drop to zero quizzing after the passing threshold is reached than those learners in other pre-passing categories.

- Among the thirty course patterns we find only one – *raising-scores + raising-scores* – to have never occurred in any of our MOOCs. This is a sensible outcome, as this pattern implies that learners put more effort into the assessments *after* passing than before.

Summarizing this analysis, we find that although many different patterns appear (for 29 out of 30 patterns we observe at least one passer exhibiting it), the majority of passers fall into only *two* course patterns, a result that holds across all four MOOCs.

Table 2.5: Overview of the fraction of early passers falling into each course pattern (combination of pre- and post-passing pattern). P_n represents unit- n passers, shown in brackets (table header) are the number of passers in each group. Marked in bold are all values $\geq 10\%$. Each column sums up to 100%; a – indicates 0%.

		#Passers							
		FP		DA		SEW		SE	
Pre-passing Behaviours	Post-passing Behaviours	P ₅ (690)	P ₆ (237)	P ₅ (448)	P ₆ (450)	P ₄ (216)	P ₅ (110)	P ₅ (379)	P ₆ (621)
Unstable	Unstable	0.6%	6.3%	-	4.7%	1.4%	10.9%	-	5.5%
	Raising	0.1%	0.4%	-	-	-	-	-	0.2%
	Reducing	3.9%	7.2%	-	13.8%	4.6%	10.9%	-	2.6%
	Keeping high	3.9%	9.7%	-	4.4%	1.9%	7.3%	-	21.7%
	Keeping mid	0.6%	2.5%	-	4.4%	0.9%	21.8%	-	2.1%
	Keeping low	0.7%	11.8%	-	11.3%	0.9%	10.0%	-	5.8%
Raising	Unstable	-	0.4%	-	0.7%	1.9%	0.9%	-	1.0%
	Raising	-	-	-	-	-	-	-	-
	Reducing	1.3%	0.4%	-	3.6%	4.6%	1.8%	-	1.3%
	Keeping high	2.9%	2.1%	0.2%	2.7%	4.6%	2.7%	-	6.4%
	Keeping mid	-	-	-	1.3%	0.5%	-	-	0.2%
	Keeping low	0.4%	0.8%	-	2.0%	0.5%	2.7%	-	1.1%
Reducing	Unstable	0.1%	0.8%	-	1.3%	0.9%	1.8%	-	1.1%
	Raising	-	0.4%	-	0.4%	-	0.9%	-	0.3%
	Reducing	0.9%	2.5%	-	3.6%	1.4%	0.9%	-	-
	Keeping high	-	-	-	0.9%	0.5%	-	-	5.3%
	Keeping mid	-	1.7%	-	3.1%	-	4.5%	-	0.6%
	Keeping low	0.7%	7.6%	-	4.2%	-	5.5%	-	6.6%
Keeping high	Unstable	8.4%	4.2%	6.0%	4.4%	14.8%	-	5.8%	2.4%
	Raising	0.4%	-	-	-	-	-	-	-
	Reducing	23.6%	10.1%	59.8%	10.9%	30.1%	-	26.9%	2.3%
	Keeping high	46.4%	15.6%	27.2%	5.6%	20.4%	-	38.3%	26.1%
	Keeping mid	1.6%	3.4%	2.5%	9.8%	4.2%	0.9%	-	0.8%
	Keeping low	3.3%	4.2%	4.2%	6.9%	6.0%	-	29.0%	3.9%
Keeping mid	Unstable	-	2.1%	-	-	-	2.7%	-	0.2%
	Raising	-	-	-	-	-	-	-	0.2%
	Reducing	-	0.8%	-	-	-	1.8%	-	-
	Keeping high	-	1.3%	-	-	-	2.7%	-	1.8%
	Keeping mid	-	1.7%	-	-	-	4.5%	-	0.3%
	Keeping low	-	1.7%	-	-	-	4.5%	-	0.3%

2.6 Conclusions

In this chapter, we investigate long-term behavioral engagement of learners based on their participation in video lectures and quizzes throughout courses. Specifically, learning analytics technology is applied on trace data to evaluate behavioral engagement of passers based on their behaviors on video watching and quiz answering before and after clinching a passing grade.

We find that the vast majority of passers pass a course at the earliest possible point, and after passing they exhibit a certain typology of post-passing behavior patterns which indicate their motivation for the course. Across the courses we explore, we find that the act of clinching a passing grade heavily influence ensuing learner behavior. We also define a number of pre-passing and post-passing behavior patterns and find the majority of learners to fall into a narrow band of behaviors independent of the specific MOOC under investigation.

We find that a certain subset of learners heavily reduce their engagement with quiz questions after clinching a passing grade. Now consider this observation in the context of the value or significance of a course certificate; there exist learners who attained a certificate (and can therefore claim mastery of the course subject) who have been exposed to only 60% of the course materials. Now that universities are beginning to offer official college credits for completing a MOOC [74], this highlights the need for course practitioners to design assessment systems which require learners to display mastery of an entire course subject before earning a certificate.

This is a first step towards gaining more detailed and fine-grained insights into learners' behaviors and motivation. Future work will expand this exploratory research to a larger number of MOOCs (from different fields, requiring different types of background knowledge) and take learners' demographic information, prior knowledge and motivations into account (do learners of a certain demographic slow down after passing more than others?). Using such insights, can we then exploit this knowledge to create MOOCs that provide a more sustained learning experience?

Chapter 3

Mobile vs. Stationary Learning

In this chapter, we focus on the impact of mobile learning on behavioral engagement of MOOC learners in a learning session. Mobile technology has become an ubiquitous part of our daily lives and enables us to learn on-the-go. The use of mobile devices for learning on-the-go requires learners to multitask and divide attention between several activities, at least one of which (the learning activity) with high cognitive load. While most MOOC platforms today offer responsive web pages and specific apps to learn via mobile devices, the learning situation and its effect on learners while using mobile devices on-the-go has not been studied in full. In contrast to most existing mobile learning studies which were conducted in the lab, our user study focuses on real-world situations commonly experienced by learners while they learn on-the-go.

This chapter is published as “Stationary vs. Non-stationary Mobile Learning in MOOCs” [149], by Yue Zhao, Tarmo Robal, Christoph Lofi and Claudia Hauff, in Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization, pp. 299-303. ACM, 2018, and “Can I have a Mooc2Go, please? On the Viability of Mobile vs. Stationary Learning” [148], by Yue Zhao, Tarmo Robal, Christoph Lofi and Claudia Hauff, in European Conference on Technology Enhanced Learning, pp. 101-115. Springer, 2018.

3.1 Introduction

With the rapid advancement of mobile technology, the use of mobile devices has become ubiquitous around the world—about 98% of the population in developed countries, and 50% of the population in developing countries had mobile-broadband subscriptions in 2017 [109]. This development has affected the way people exploit mobile technology to learn new skills—a significant number of people use mobile devices for learning. A survey on lifelong learning by Tabuenca et al. [129] finds that 56% of learners used their smartphone on a daily basis, whilst a study on mobile language learning by Dinger et al. [29] reports that about 38% of learning sessions took place while in transit. According to O'Malley et al. [81], mobile learning refers to “*any sort of learning that happens when the learner is not at a fixed, predetermined location, or learning that happens when the learner takes advantage of the learning opportunities offered by mobile technologies.*”

The start of the MOOC movement in 2011 vastly widened the learning opportunities for people across the world outside of a formal education setting. While in the early years MOOC platforms lacked support for mobile devices, by 2015, most well-known platforms (such as edX, Coursera and Udacity) offered a mobile learning experience [66], either in the form of responsive web pages or native mobile apps (for Android and iOS), thus further expanding the possibilities to learn anywhere and anytime.

Critical for mobile learning [118, 117, 94, 119] is the *learning situation*—a set of environmental and intentional constraints [7]—in which learning occurs. A learner’s available time, the employed device type(s), and the frequency of interventions or distractions are only a few of those constraints that affect learning. One common learning situation for MOOC learners is *stationary learning*: here, learners use a device with a large screen to access course materials whilst being stationary in a comfortable environment (e.g. at their desk), enabling them to focus on the learning activity. In the mobile learning situation², the conditions are quite different—mobile devices have considerably smaller screens and they are used in various and possibly changing environments which require learners to multitask (e.g. learning whilst walking or transiting). In terms of learning, this situation results in an increase in interruptions and distractions [117], an increase in cognitive load [129, 16, 29], and increased frustration [23].

²In the remainder of this chapter, we refer to learning in a non-stationary situation with a mobile device as *learning on-the-go*.

Existing studies on mobile learning in MOOCs focus on the design and delivery of course content for mobile devices [97, 66] as well as the learning experience on mobile devices [90, 142, 23, 142]; the latter though is typically studied *in the lab*, instead of real-world environments. Thus, little is known about how multitasking and a multitude of overlapping *real-world conditions* affect MOOC learning on-the-go compared to stationary learning. This knowledge gap serves as the core motivation for our work.

More specifically, we focus on the impact of the learning situation on learners' performance and interactions, the effect of different environmental variables on the learning on-the-go process, and the correlation between learners' perceived workload and their performance/interactions. We analyze both trace data and survey data collected from a user study with 36 participants, each of whom completed two mini-MOOCs (one in stationary and one in the on-the-go condition³ at specific times of the day to control for daylight and crowdedness), guided by the following research questions:

RQ 2.1: To what extent does learning on-the-go (compared to stationary learning on a mobile device) affect MOOC learners' learning gain, learning efficiency and interactions with the course content?

RQ 2.2: How do learners perceive their workload (physical as well as mental) in the stationary and learning on-the-go conditions and how does it relate to their learning performance and interactions?

3.2 Background

Our research addresses the following aspects of online learning: multitasking and attention fragmentation, and the use of mobile devices in different learning situations, with a focus towards learning in MOOCs.

3.2.1 Multitasking and Divided Attention

Interacting with a mobile device while on-the-go requires the ability to multitask and divide one's attention between several tasks efficiently at once. Multitasking—the act of attempting to engage simultaneously in two or more tasks that have independent goals [36]—is directly connected to our research on mobile learning from MOOCs.

³In this condition our participants physically explored the university campus.

Multitasking is tightly coupled with the attention level and situational awareness. Studies on walking and mobile use have highlighted the increase of cognitive load and a necessity to divide attention, thus forcing mobile users to correct their gait and walk slower while performing tasks on mobile devices [60, 61].

Multitasking also incurs a cost on performance and accuracy for other tasks as our ability to effectively process two or more attention-demanding tasks simultaneously is limited [36], and performance across two concurrent tasks is optimized based on perceived priorities [31]. Thus, switching between activity contexts (e.g. in the on-the-go setting switching between reading the slides, paying attention to the traffic, or listening to the video lecture) lowers task effectiveness. Harvey and Pointon [45] investigate the effect of fragmented attention on mobile web search tasks in three different contexts (walking on a treadmill, navigating through an obstacle course, and sitting down) and find that the contextual situation affects user (search) task performance—walking affected participants’ objective and perceived search performance negatively. In addition, participants who performed searches while on the move reported a higher difficulty and cognitive workload in performing the tasks than those sitting. In MOOC learning, which requires a high degree of attention and commitment, this indicates a potential for less effective learning in the on-the-go condition compared to the stationary one. Xiao and Wang [142] investigate the impact of divided attention on the learning process and learning outcomes for mobile MOOCs, and propose to detect divided attention via monitoring learners’ heart rate. In their study with 18 participants under lab conditions, they observed divided attention to hurt learners’ performance.

With respect to multitasking and fragmented attention, our study explores the effect and extent learning on-the-go has on learners’ ability to comprehend course content, and on their cognitive learning performance.

3.2.2 Mobile Learning

Mobile learning (i.e. learning with a mobile device) stresses the possibility to learn across time and space, and commonly assumes that learners are on the move [118]. What mainly distinguishes mobile learning from traditional classroom learning is the variety and unpredictability of the situations in which learning can take place [117] which places different demands on learners’ attention level, body posture, environment, and social context whilst learning.

Mobile technology has enabled context-sensitive learning and the use of sensor data of mobile devices to enrich the learning experience [119]. Dinger et al. [29] implement an Android app to collect sensor data (e.g. location, ringer mode, or motion) in order to detect learners' contexts and boredom levels during learning sessions on mobile devices. Based on a user study, the authors conclude that while on mobile and in transit people are more open to engage in quick learning sessions, and context information retrieved from phone sensors can be helpful for mobile learning.

Learning tasks that are cognitively demanding (e.g. reading and writing scientific essays) seem to be incompatible with the use of mobile phones while on-the-go, whereas activities that are less cognitively demanding (e.g. social networking, texting, or taking pictures) are compatible with body movement [16]. Music et al. [76] attempt to detect changes in user attention by exploiting smartphone accelerometers to trace changes in user gait patterns as a response of interaction with a mobile device. In a traditional study setting (e.g. a library, classroom), the use of mobile phones whilst learning has been found to be a distraction for most learners [6]; the same can be said about the mobile MOOC setting as incoming notifications, messages, news, etc. can take learners' focus away from the actual learning task.

The mobile devices themselves also affect learner perceptions. Dalipi et al. study learners' experience by comparing desktop and mobile platforms of three well-known MOOC environments (edX, Coursera, and Udacity) [23]. They find that learners were more satisfied with the respective desktop variants; mobile platforms with their small screens and a lack of external input devices caused negative emotions as a number of tasks, which were easy on the desktop variants, were rather difficult to execute on the mobile variants. In a similar vein, Becking et al. [7] argue that learning situations for learning on-the-go are uncomfortable because of the lack of space for taking notes, and the potential for interruptions.

In our study, we explore learning with a mobile device in two different settings: (i) on-the-go and (ii) in a seated and more convenient condition close to traditional online learning, yet with a mobile device. In the former condition, we do not confine our participants to the lab (e.g. by using a treadmill or an obstacle course), but instead ask them to physically explore the university campus whilst learning.

3.3 Study Design

3.3.1 Learning Situations

Inspired by the mobile search study conducted by Harvey and Pointon [45] (who find walking to impact participants workload perception and search effectiveness), we investigate whether learning on-the-go has any measurable impact on learning gain, effectiveness and perceived workload compared to stationary learning in the MOOC setting. We consider the following two learning situations (or scenarios) in our user study:

Stationary Scenario (*StaSc*): Learners study MOOCs while sitting in the office with a mobile device. This scenario is used as the baseline in order to measure the impact moving around has on learning.

Moving Scenario (*MovSc*): Learners study MOOCs with a mobile device while on-the-go. Participants are asked to learn whilst walking from one building to another on campus at their normal walking speeds, while paying attention to the traffic.

To eliminate the effects of learning behaviors unrelated to the use of mobile devices (e.g. taking notes on a piece of paper) and of different types of mobile devices, we instructed our study participants to perform all learning tasks exclusively on the same mobile device⁴ in both *StaSc* and *MovSc*. We hypothesize—in line with the findings in [142]—that compared to *StaSc*, the necessary multitasking and the possible interruptions and distractions in *MovSc* negatively affect MOOC learners' learning gain. We also hypothesize that participants in *MovSc* require more time to consume the course materials (due to the divided attention) than those in *StaSc*. In line with the previous hypothesis, we anticipate participants in *MovSc* to revisit the video page more often and rewind the video more often than those in *StaSc* to refresh their memory (which is impaired due to the distractions on-the-go).

3.3.2 Learning Materials

We prepared four mini-MOOCs on different topics (Table 3.1) for our user study and deployed them on edX Edge, a low-visibility clone of the edX platform.

⁴A Samsung S5 smart-phone with 1,080 × 1,920 pixels, 5.1" display screen, 2GB RAM, 2.50 GHz CPU, Google Android 6.0.1 and the Chrome browser installed.

All four mini-MOOCs had the same structure: one lecture video and 20 knowledge questions about the video content. To ensure similar difficulty across the four mini-MOOCs, we selected them from a pool of introductory MOOC video lectures produced by the Delft University of Technology for the edX platform. We chose those four based on their similar *amount of unfamiliar terminology* as labelled by three annotators with computer science degrees. Each question was a multiple-choice question (almost all with four answer options in addition to *I don't know*). These questions were not only used in the mini-MOOCs (right after the video lecture) but also in the pre-study questionnaire, which enables us to compute the knowledge gain in a straight-forward manner. This setup also means that the questions covered key knowledge concepts discussed in the respective lecture, instead of specific video details (such as the number of instructors, or the color of the background). Each question could be attempted once in the pre-study questionnaire and MOOC.

The pre-study questionnaire thus contained $4 \times 20 = 80$ questions about the four topics; we used answers to those questions to select for each study participant the two mini-MOOCs with the *lowest* prior knowledge levels. This setup leads to large potential knowledge gains. Table 3.1 lists the pre-study knowledge scores for the four mini-MOOCs across our 36 participants. Note that the maximum obtainable score for the questionnaire was 20 for each topic. The Qubit topic proved to be the most difficult, with more than half of the participants answering 0 or 1 question correctly; in contrast, water quality aspects proved to be the easiest topic with half the participants answering between 7 and 11 questions correctly.

Table 3.1: Overview of our mini-MOOCs, the video length per MOOC and the minimum/median/maximum of participants' prior knowledge test scores on the topics. The highest possible score per topic is 20.

Mini-MOOC	Video Length	Pre-study Scores		
		Min.	Median	Max.
Radioactive decay	6m53s	0.0	3.0	9.0
Qubit	12m24s	0.0	1.5	16.0
Water quality aspects	10m45s	1.0	7.0	11.0
Sedimentary rocks	5m03s	0.0	4.0	10.0

3.3.3 Environmental Conditions

In our study, next to stationary and on-the-go, we focus on the impact of two additional environmental variables—the *light condition* and the *crowdedness*

of the surrounding. It is known that daylight can affect the visibility of the screen on mobile devices [143] and the visibility of the surroundings during learning. The crowded learning situation may lead to intensive interruptions and distractions in *MovSc*. We thus hypothesize daylight and crowdedness to lead to reduced learning gains. Note that these environmental conditions only apply to *MovSc*.

Table 3.2: Number of participants under different experimental conditions.

Mini-MOOC	<i>MovSc</i>				<i>StaSc</i>
	Daylight & Crowded	Daylight & Uncrowded	Dark & Crowded	Dark & Uncrowded	
Radioactive decay	3	1	4	2	15
Qubit	3	5	3	4	13
Water quality aspects	0	2	0	0	2
Sedimentary rocks	2	0	3	4	6
Total	8	8	10	10	36

Study participants were randomly assigned to one of four groups based on the time of the experiments for *MovSc*: (i) 8:45 am (crowded time with daylight), (ii) 11:00 am (uncrowded, daylight), (iii) 5:45 pm (crowded, no daylight⁵), and (iv) 8:00 pm (uncrowded, no daylight). Table 3.2 shows the distribution of study participants across the four groups.

3.3.4 User Study Steps

In our experiments, each participant was guided through the following steps.

1. Pre-study questionnaire: 80 knowledge questions plus questions on demographics, experience with mobile devices, mobile learning and MOOCs;
2. In random order, complete *StaSc* and *MovSc* with the two mini-MOOCs that exhibited the lowest prior knowledge levels. During a mini-MOOC, participants were allowed to switch between the video and questions. Each of the two scenarios was assigned a 30-minute time block.
3. Post-MOOC questionnaires: after each of the two scenarios a NASA TLX workload assessment form⁶ [43] had to be completed. It assessed

⁵We conducted this user study in December 2017 and January 2018 in Delft, the Netherlands.

⁶<http://www.nasatlx.com/>

the workload during learning in each scenario on six aspects: mental demand, physical demand, temporal demand, performance, effort, and frustration.

3.3.5 Metrics

We now describe how we measure participants' learning gain, learning efficiency and interactions. To measure the statistical significance of the difference between groups of learners, we employ the *Mann-Whitney U test*.

In our study we use *absolute learning gain (ALG)* and *realized potential learning (RPL)* to measure participants' **learning gain** [127]. *ALG* refers to the number of questions that were answered *incorrectly* in the pre-study questionnaire and *correctly* in the mini-MOOC, normalized by the total number of questions (20). *RPL* refers to the *absolute learning gain* normalized by the maximum possible learning gain⁷.

We measure **learning efficiency** through the efficiency of (i) course material consumption and (ii) learning gain. For the former, the time participants spent on watching videos (i.e., *video duration* and *normalized video duration*) and answering questions (i.e., *question duration*) are calculated—as we deployed our mini-MOOCs on edX Edge, we have access to all tracking data logged by edX. As shown in Figure 3.1, *video duration (VD)* refers to the minutes a participant spent watching the lecture video. *Normalized video duration (NVD)* refers to *VD* normalized by the video length, which measures the proportion of the video consumed. *Question duration (QD)* refers to the minutes a participant spent on the questions, including any time spent on video rewinding. To compute the **efficiency of the learning gain**, we divide *RPL* by *VD* and *NVD*.

As **interactions** metrics we consider those that lead the participant away from the default mini-MOOC path (i.e. watch the video lecture and answer the 20 quiz questions). Specifically, we use the times participants revisited the video page during question answering (i.e. *#video page revisiting*, *#V_revisit* in short) and the minutes participants spent on video rewinding for questions (i.e. *video rewinding duration*, *VRD* in short) as metrics.

⁷For example, if in the pre-study questionnaire a learner answers 2 out of 20 questions correctly, the maximum possible learning gain is 18. If in the MOOC quiz two more questions are answered correctly, then *ALG* is $\frac{2}{20}$ and *RPL* is $\frac{2}{18}$.




Figure 3.1: An example of a participant’s learning progress. In this example, *video duration (VD)* is $V_1 + V_2 + V_3 + V_4$, *initial video watching duration* is $V_1 + V_2$, *video rewinding duration (VRD)* is $V_3 + V_4$, *question duration (QD)* is $Q_1 + V_3 + Q_2 + V_4 + Q_3$, and *question answering duration* is $Q_1 + Q_2 + Q_3$.

3.3.6 Study Participants

We recruited study participants within the faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology through flyers and mailing lists. 36 learners participated in our study: 9 women and 27 men. Their average age was 24.4 (std.dev. 2.7; min. 19; max. 30). Most participants were Master students, the highest educational degree (so far) was: high school (5 participants), Bachelor’s degree (21) and Master’s degree (10). On average, the participants had been using smart-phones for 7 years; all indicated to use them daily. 27 participants had used their mobile device for a learning activity within the last seven days before the user study. 26 participants had registered to at least one MOOC, 13 had made use of their mobile devices to learn in a MOOC and 11 participants had successfully completed at least one MOOC.

On average, each participant took about two hours to complete the entire experiment (recall, that each mini-MOOC was given a thirty minute time limit, however additional time was required for the pre-study questionnaire, switching scenarios, explanations by the experimenter, post-MOOC questionnaires and so on). Participants received a payment of €15. To motivate participants to learn, we provided a bonus payment of €5 for the participant achieving the highest learning gain overall.

3.4 Results

3.4.1 Learning Gain, Efficiency and Interactions

In Table 3.3 (rows 1 & 2) we report our learning gain metrics across the two learning scenarios and the different environmental conditions, aggre-

gated over all participants and topics. We find that, overall the learning gain achieved in the *MovSc* setting ($ALG = 0.47$) is slightly lower than in *StaSc* ($ALG = 0.5$). The difference is not significant though; similarly, the environmental conditions exhibit no consistent tendency. More concretely, as in our setup (20 questions per mini-MOOC), an ALG value of 0.05 represents one question answered correctly in the mini-MOOC but not the pre-study questionnaire, the recorded difference between *StaSc* and *MovSc* means that on average not quite one more question is answered correctly in the stationary learning scenario—this is in contrast to our hypotheses, where we expected to find considerable differences in learning gain across the two learning scenarios. The findings also hold for *RPL*; here a value of 0.05 means that 5% of those questions not answered correctly in the pre-study questionnaire are answered correctly in the mini-MOOC.

In terms of **learning efficiency**, the results in Table 3.3 (rows 3 to 7) show that in line with our hypotheses, participants in the *MovSc* scenario did take slightly more time to consume the lecture videos than those in the *StaSc* scenario. Importantly, participants spent significantly more time on questions in *StaSc* (on average 16 minutes) than in *MovSc* (13 minutes), a finding that corresponds to the results in [45] where stationary and on-the-go mobile web search tasks were compared. This result can be explained by the fact that a comfortable and stationary environment allows participants to engage with in-depth tasks requiring a lot of focus. Remember though, that this additional time spent on questions did not result in significantly higher learning gains as seen in our previous analyses. Once again, when considering the impact of the environmental variables, we do not observe a consistent trend, one way or another.

To determine the **efficiency of learning gain**, we measure how much participants learn from video watching. We hypothesize that *MovSc* has a negative impact on participants' efficiency of learning gain. *RPL/VD* refers to participants' learning gain *per minute of video watching*. We find that on average participants in *StaSc* reach a 40% higher efficiency (statistically significant) than in *MovSc*. We again do not observe clear trends for the different environmental variables.

When we consider **learners' interactions** in Table 3.3 (rows 8 & 9) it is evident that on average participants in *StaSc* spent nearly twice as much time rewinding the videos than those in *MovSc*. The same trend holds for the number of times participants revisited the video playing page during question answering. Both of these findings indicate that in *StaSc* participants put more effort on finding relevant information for question answering than in

Table 3.3: The average value and standard deviation of metrics on participants' learning gain, learning efficiency and interactions under different experimental variables. † indicates significance at $p < 0.1$ level. ‡ indicates significance at $p < 0.05$ level. ◊ indicates significance at $p < 0.01$ level.

Metrics	Learning Situation			MovSc with Different Environmental Variables					
	StaSc (S)	MovSc (M)	Daylight & Crowded (DIC)	Daylight & Uncrowded (DIU)		Dark & Crowded (DkC)		Dark & Uncrowded (DkU)	
				Uncrowded	Crowded	Uncrowded	Crowded	Uncrowded	Crowded
ALG	0.504(±0.130)	0.474(±0.145)	0.463(±0.155)	0.463(±0.074)	0.480(±0.164)	0.485(±0.178)			
RPL	0.575(±0.140)	0.533(±0.164)	DkU†S†0.484(±0.161)	0.550(±0.125)	0.536(±0.177)	0.554(±0.195)			
VD (minutes)	10.796(±3.929)	11.883(±4.125)	10.881(±4.577)	S†13.179(±1.937)	11.068(±5.131)	12.463(±4.139)			
NVD	1.304(±0.572)	1.407(±0.519)	DkU†1.312(±0.468)	DkU†1.187(±0.189)	1.457(±0.716)	S†1.609(±0.486)			
QD (minutes)	16.284(±6.754)	S†12.581(±6.323)	S†12.142(±6.983)	13.913(±6.551)	S†12.703(±6.833)	S†11.745(±5.916)			
RPL/VD	0.074(±0.080)	S†0.053(±0.029)	0.053(±0.029)	S†0.043(±0.013)	0.063(±0.040)	0.050(±0.026)			
RPL/NVD	0.583(±0.531)	S†0.419(±0.170)	S†DU†0.384(±0.126)	0.475(±0.138)	0.459(±0.236)	S†DU†0.363(±0.141)			
VRD (minutes)	4.515(±4.514)	S◊2.284(±3.416)	S†2.102(±3.523)	S†2.048(±3.485)	S†2.698(±4.406)	S†2.203(±2.568)			
#V_Revisit	5.056(±5.270)	S◊2.250(±2.708)	S†2.500(±3.546)	S◊1.125(±1.356)	S†2.700(±3.093)	S†2.500(±2.506)			

MovSc. In order to understand participants' interactions in more detail, in Figure 3.2 we plot on a per-participant basis their (i) video watching duration before they started question answering (i.e. *initial video watching duration*), (ii) their *video rewinding duration* during question answering and (iii) their time spent on question answering only (i.e. *question answering duration*).

Compared to *StaSc*, it is evident that participants in the *MovSc* scenario tended to spend more time on video watching before they started question answering and less time on question answering. During question answering, most participants in *MovSc* revisited the video playing page fewer times and spent less time on video rewinding than in *StaSc*. This finding shows that participants in *MovSc* tended to switch less between the video playing page and the question page than those in *StaSc*. An explanation for the long question answering duration in *StaSc* can be that question answering is an activity with higher cognitive demand than video watching, which is not as compatible as video watching with walking with a mobile device [16].

3.4.2 Learning and Perceived Workload

We now investigate the relationships between participants' learning and their workload perception. Concretely, we report the Pearson correlation coefficient between our learning & interaction metrics and the six aspects of workload participants self-reported via the NASA TLX form. The results are shown in Figure 3.3; here, *TLX score* is the overall score of workload, and *MentDmd*, *PhysDmd*, *TempDmd*, *Perform*, *Effort*, *Frustr* are participants' workload scores on mental demand, physical demand, temporal demand, performance, effort, and frustration respectively.

When comparing *StaSc* and *MovSc* we observe sensible results with respect to mental demand and physical demands: in both scenarios the mental demand was found to be the most important one, followed by the physical demand in *MovSc* (in contrast to *StaSc*, where the physical demand received the lowest average weighting).

In *StaSc* we find performance (*How successful were you in accomplishing what you were asked to do?* with answer options ranging from *Poor* to *Good*) to be negatively correlated with learning gain, i.e. our participants were not able to estimate their own learning success very well. In contrast, performance is positively correlated with *normalized video duration*, indicating that participants estimated their learning performance to at least some extent based on how much of the video content they watched.




Figure 3.2: The time participants spend on difference activities in *StoSc* and *MonSc*.




Figure 3.3: Linear correlation coefficient between participants’ learning performance, interactions and their perceived workload as measured through the NASA TLX form. The x-axis label also shows the average score of each workload dimension across our participants.

In the *MovSc* scenario, participants were also not able to self-estimate their learning gains (we find a slight negative correlation between *ALG/RPL* and performance); most interesting though is the positive correlation between *frustration* and question duration, i.e. the longer participants in the on-the-go condition spent answering questions, the more frustrated they felt (though overall frustration was not a major workload dimension).

3.5 Conclusions

In this chapter, we investigate the impact of mobile learning on learner behavioral engagement in learning sessions. Due to the requirement for divided attention and multitasking in learning on-the-go, learners cannot be fully engaged in a learning session. With a controlled study, the impact of mobile learning on learner engagement can be evaluated based on learners’ performance and behaviors.

Concretely, our study focuses on to what extent learning on-the-go (compared to stationary learning on a mobile device) affects MOOC learners’ learning gain, learning efficiency and interactions with course content. Our investigation includes a foray into the influence environmental variables (light conditions and crowdedness) have on mobile learning. A second research

question we consider is the relationship between learners' perceived workload and their learning.

In order to explore these questions, we designed a user study with 36 participants; each participant “followed” two mini-MOOCs deployed on the edX Edge platform: one in the on-the-go condition (learning on a mobile device while walking) and one in the stationary condition (learning on a mobile device while being stationary). We measure participants' learning through a set of pre/post-study multiple choice question sets. Our analyses result in the following key findings:

- On average, learning on-the-go (*MovSc*) resulted in a lower (-6% in *ALG*) learning gain than stationary learning (*StaSc*) with a mobile device.
- Compared to *MovSc*, *StaSc* participants spent 29% more time on answering questions and reached a 40% higher learning efficiency.
- When it comes to workload perception, participants in both conditions were not able to estimate their performance (wrt. learning gain) well; *MovSc* participants reported higher physical demands and slightly higher frustration than participants in the *StaSc* condition, though the differences in learning gains were small (first key finding).
- The environmental variables we investigated (daylight and crowdedness) did not have a consistent impact on any of the metrics investigated.

Our study has several limitations, among them the size of the user study (36 participants in total) which provides us with trends but few significant differences. A second limitation is the simplification of the on-the-go scenario to a walk on the campus (which does improve though—in terms of realism—on the lab conditions in prior studies). As pointed out by Becking et al. [7], the learning situation might be more complicated and unstable in many situations. Learners may walk, wait or take a bus or train while learning with a mobile device. Additionally, we only considered two environmental variables—the light condition and the crowdedness; other variables such as the weather and the temperature (recall that we conducted the experiments during December/January, i.e. the winter season in Europe) were not considered, although they are likely to also affect our participants' behaviour. For example, two participants who were assigned the 8 pm timeslots for the study told us that they aimed to finish their learning sessions as quickly as

possible due to the bad weather. In the future to measure learners' interactions in more complex learning situations, a dedicated mobile app may be needed to record fine-grained details of learners' contexts and actions whilst on-the-go.

Chapter 4

Eye-Tracking Based Inattention Detection

In this chapter, we focus on short-term behavioral engagement of MOOC learners. If we were able to track short-term behavioral engagement of MOOC learners in (near) real-time, intervention can be provided to learners once they are disengaged, which might improve learner engagement in MOOC learning.

On some MOOC platforms (e.g. edX in our study), trace data generated by learner interactions (which is used in our studies in Chapter 2 and Chapter 3) is not sufficient to track learner behavioral engagement in near real-time by using learning analytics technology. For example, when a learner is fully engaged in video watching, she may focus on following the lecturer without having any operations on video playing. Consequently, her activities for this video watching are only recorded as two events *video start* and *video end* in the log trace data of edX. With these event records, the extent of learner engagement in video watching cannot be accurately tracked in near real-time by using learning analytics technology on their trace data on edX.

In particular, we investigate the webcam-based eye-tracking method for tracking learners' inattention in video watching. Since most MOOCs on premier MOOC platforms are centered around video lectures, in which learners might become distracted and lose their attention frequently without realizing it.

Inattention (also known as *mind-wandering* in psychology) is a frequently occurring experience for many learners and negatively impacts learning outcomes. While in traditional classroom contexts, a skilled teacher may be

able to observe and react to learners' loss of attention, no such intervention is possible (yet) in MOOCs. Previous studies suggest a strong relationship between learners' inattention and their gaze, making it possible to detect inattention in real-time using eye-tracking devices. Existing research in this area though has made use of *specialized* (and expensive) hardware, and thus cannot be employed in MOOC scenarios due to the inability to scale beyond lab settings. In order to make a step towards *scalable* inattention detection among MOOC learners, we propose the use of webcams. In our user study, we compare the accuracy of inattention detection based on gaze data recorded through a consumer grade webcam and a specialized and high-quality eye tracker.

This chapter is published as “Scalable Mind-Wandering Detection for MOOCs: A webcam-Based Approach” [147], by Yue Zhao, Christoph Lofi, and Claudia Hauff, in European Conference on Technology Enhanced Learning, pp. 330-344. Springer, 2017.

4.1 Introduction

Inattention (mind-wandering) is an essential part of human behavior consuming up to 50% of everyday thoughts [55], and can be described as *thoughts and images that arise when attention drifts away from external tasks and perceptual input toward a more private, internal stream of consciousness* [69]. While inattention can also have positive effects (such as fostering creativity [132]), many educational tasks including following a lecture or solving an assignment require active attention and focus on reaching the desired learning outcomes. For these tasks, excessive inattention has disastrous effects on learning efficiency [123].

In traditional classroom contexts, attention lapses have been studied for a long time, e.g. [14, 140]. Although researchers do not yet agree on the actual attention span of learners, several past works have found attention among students during lecture time to vary in a cyclic manner.

For online courses and MOOCs, this problem is even more severe as they are consumed using digital display devices. This mode of consumption is particularly prone to inattention. Likely due to the ubiquity of smartphones and digital content, a significant subgroup of online users adopt a “heavy media multitasking” behavior [63], making it challenging for them to focus on a single multimedia content unit. This finding is also supported by our work, where learners frequently lose their attention even in short video clips of around seven minutes.

In order to detect inattention among online learners during their consumption of digital materials, we require an approach that is *scalable* (it can be deployed to thousands of learners), *near real-time* (inattention is detected as soon as it occurs), *unobtrusive* (learners are not distracted by the detection procedure) and *autonomous*. In addition to providing insights into learners’ behaviors, such a method would also enable real-time interventions that lower the amount of inattention taking place. As a concrete example we envision an intelligent MOOC video player: the player (via the webcam feed) monitors a learner’s attention state and when a loss of focus is detected, the player pauses the video automatically in order to avoid skipping over relevant content. In order to ensure learners’ privacy, all necessary processing will be client-side (i.e. executed within the browser).

To this end, previous research shows that by analyzing people’s gaze data, inattention can be detected, e.g. whilst reading texts on screen [8], or watching (non-educational) films [9]. These results can be attributed to the eye-mind link effect [96], which states that *there is no appreciable lag*

between what is fixated and what is processed. Existing works usually rely on expensive and specialized eye-tracking hardware (e.g. a Tobii eye tracker) to obtain gaze data, which is not available to the average MOOC learner. It is therefore still an open question whether eye-tracking based inattention detection can be performed in a scalable manner.

Our goal in this chapter is to develop a fully automatic method for detecting loss of attention in near real-time using only low-end webcams ubiquitously found on laptop computers. To this end, we conduct a lab study with 13 participants, collecting a dataset of gaze features (i.e. features extracted from gaze data) and self-reported inattention. To motivate this approach, Figure 4.1 visualizes the gaze of two of our study participants through heatmaps. The MOOC video shown has several relevant visual areas, including the lecture slides, the subtitles, and the speaker’s face. In the depicted scene, a changing set of examples is shown on the slides which are important to grasp the lecture content. The participant who reported inattention in the 30-second interval intently gazed on a spot on the speaker’s face, ignoring the slides and the shown examples, while the second participant who reported no inattention focused on all relevant areas of the video. Our proposed approach employs supervised machine learning to automatically learn such inattention patterns based on gaze features.


Our contributions in this work are as follows:

1. We create an elaborate gold dataset to foster eye-tracking based inattention research, featuring 13 participants watching two MOOC videos each in a controlled lab setting, reporting feedback on inattention in brief intervals. In addition to these inattention reports, we provide video and gaze data as recorded and analyzed by a professional eye tracker as well as gaze data recorded by a webcam and processed by an open-source gaze library. We make this data available on our companion Web page [146].
2. We implement and evaluate an approach to automatically detect inattention based on gaze data (i) collected with a specialized eye-tracking device (*Tobii X2-30*²), relying on the results and best practices published in [9], and (ii) collected with a standard webcam.
3. We extensively discuss and evaluate both approaches, and argue that our webcam-based method is indeed suitable for large-scale deployment outside a controlled lab setting.

²<https://www.tobii.com/product-listing/tobii-pro-x2-30/>



(a) Reported Inattention



(b) No Inattention

Figure 4.1: Gaze heatmaps of two study participants over a 30-second interval

4.2 Background

Different data collection methods have been used to study inattention of students in traditional classrooms since the 1960s, such as the observation of inattention behaviors [53], the retention of course content [68], using direct probes in class [126, 62], or relying on self-reports from students [14]. A common belief was that learners' attention might decrease considerably after 10 – 15 minutes of the lecture, which was supported by [126]. However, Wilson and Korn [140] challenge this claim and argue that more research is needed. In a recent study, Bunce et al. [14] asked learners to report their inattention voluntarily during 9 – 12 minute course segments. In their experiments, three buttons were placed in front of each learner, representing attention lapses of 1 minute or less, of 2 – 3 minutes and of 5 minutes or more. During the lectures, learners were asked to report their inattention by pressing one of three buttons once they *noticed* their inattention. This setup leads Bunce et al. [14] to conclude that learners start losing their attention early on in the lecture and may cycle through several attention states within the 9 – 12 minute course segments.

In online learning environments, inattention may be even more frequent. Risko et al. [98] used three 1-hour video lectures with different topics (i.e. psychology, economics, and classics) in their experiments. While watching the videos, participants were probed four times throughout each video. The inattention frequency among the participants was found to be 43%. Additionally, Risko et al. [98] find a significant negative correlation between test performance and inattention. Szpunar et al. [128] investigate the impact of interpolated tests on learners' inattention within online lectures. In their study, participants were asked to watch a 21-minute video lecture (4 segments with 5.5 minutes per segment) and report their inattention in response to random probes (one probe per segment). The inattention frequency found in their experiments was about 40%. Loh et al. [63] also employ inattention probes to measure learners' inattention and find a positive correlation between media multitasking activity and learners' inattention (average frequency of 32%) whilst watching video lectures. Based on these considerably high inattention frequencies, we conclude that reducing inattention in online learning is an important approach to improve learning outcomes.

Inspired by the eye-mind link effect [96], a number of previous studies [8, 9, 73] focus on the automatic detection of learners' inattention by means of gaze data. In [8, 9], Bixler and D'Mello investigate the detection of learners' inattention during computerized reading. To generate the ground

truth, the study participants were asked to manually report their inattention when an auditory probe (i.e. a beep) was triggered. Based on those reports, the inattention frequency ranged from 24% to 30%. During the experiment, gaze data was collected using a dedicated eye tracker. In contrast to [8, 9], Mills et al. [73] mainly focus on the relationship between a participant’s gaze and areas of interest (AOIs), specific areas in the video a participant should be interested in. Mills et al. asked the study participants to watch a 32-minute, non-educational movie and self-report their inattention throughout. In order to detect inattention automatically, statistical features and the relationship between gaze and video content were considered.

4.3 Methodology

In our study, we focus on the automatic detection of learners’ inattention through webcam-based eye tracking. The scenario we consider is video lecture watching, which is the most common manner of conveying lecture content in MOOCs [98]. We collect data through a lab study with 13 participants who were asked to watch two lecture videos and regularly report their inattention during this time. We recorded their gaze data with a dedicated high-quality eye tracker and a standard webcam. In this chapter, gaze data refers to both gaze points (the points on the screen a participant is actively looking at) and gaze events (i.e. fixations and saccades). Fixation refers to the action that concentrates the gaze points on a single area, and saccade refers to the quick and simultaneous movement of both eyes between two or more phases of fixations.

Compared to previous works [98, 128, 63, 73], the two MOOC lecture videos in our study are considerably shorter—they are between six and eight minutes in length, in line with standard MOOC practices today. To collect the ground-truth (did inattention occur in the last n seconds?), we rely on inattention probes which have proven to be effective in traditional classroom contexts [15, 126, 62] and online learning [8, 9]. Probes (regularly and actively seeking input from the study participants) are more reliable than self-caught reports which require study participants to think about their loss of attention and about reporting it [122]. In response to our probes (in the form of an auditory signal—a bell) during video lecture playback, participants were asked to press a key to indicate that they experienced inattention in the past 30 seconds. Participants who did not experience inattention were asked to ignore the bell and continue watching.

Having collected the ground truth data, we next turn to the extraction of features from gaze data, following [73]. In line with previous works, we extract features from gaze events. These gaze events are generated by gaze points. Note that gaze points are not measured directly - they are estimated from the recorded eye and iris movements; we use the existing software libraries of our dedicated high-quality eye tracker and our open-source webcam-based framework to turn eye and iris movements into gaze points.

Finally we employ the ground truth data and extract features in a supervised machine learning task to explore to what extent the automatic detection of inattention in this setting is possible.

The overview of the processing pipeline is shown in Figure 4.2. In the following sections, we first describe in more detail the experimental design of our study, and then elaborate on the features we extract.

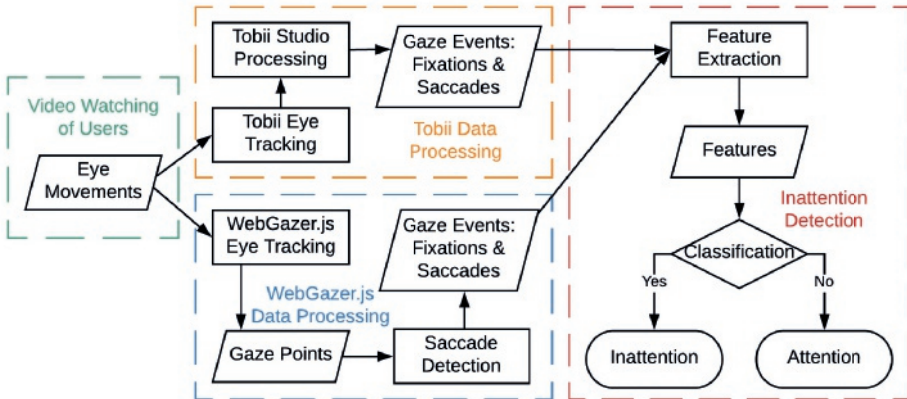


Figure 4.2: Overview of the processing pipeline

4.3.1 Study Setup

Our study is built around two introductory videos taken from two different xMOOCs [101] professionally produced and offered by the Delft University of Technology on the edX platform. One video, (taken from the *Understanding Nuclear Energy* MOOC), covers the basics of the atomic model with a length of 6 : 41 minutes; the second one (part of the *Solar Energy* MOOC and 7 : 49 minutes long) introduces the concept of energy conversion. We select those videos specifically as they contain rich visual lectures slides overlaid with the speaker (see Figure 4.1). They cover topics we consider interesting to

a wider audience and do not require extensive prior knowledge due to their introductory nature. All study participants watched both videos; their order was randomized to avoid order effects.

We used two eye-tracking devices in the experiment, a high-quality one as a reference and a low-quality webcam. Concretely, we made use of the professional *Tobii X2-30* eye tracker and its corresponding software Tobii Studio to estimate participants' gaze points. Our webcam was the built-in camera of our experimental laptop, a Dell Inspiron 5759 with a 17-inch screen and a $1,920 \times 1,080$ resolution. To estimate the gaze points based on a live webcam feed, we relied on *WebGazer.js*³ [83], an open-source eye-tracking library written in JavaScript. We built a Web application closely resembling existing MOOC lecture video players with additional logging capabilities. In order to alert our participants to each inattention probe, we included a medium-volume acoustic bell signal played by the Web application. After the bell, participants reported their inattention in the past 30 seconds by pressing a feedback button. The next bell signal occurred after another 30 – 60 seconds. The actual time was randomized within those boundaries, as previous research [8, 63] suggests that participants perceive interruptions which are not perfectly periodic as less interrupting. In order to further limit the mental annoyance of this process, participants were only asked to actively report in case they had indeed experienced inattention. This process resulted in inattention reports for each participant, including the bell signals and participant responses with respect to inattention as shown in Figure 4.3.




Figure 4.3: An example inattention report

We recruited our study participants (six females, seven males, all with a computer science background) through an internal mailing list and did not pay them. After a pre-study briefing, we asked our participants, 6 of whom wore glasses or contact lenses, to sit stable and comfortably in front of the laptop (with a distance of 52 – 68 cm between eyes and screen). The study consisted of pre- and post-study questionnaires, an instruction phase

³<https://webgazer.cs.brown.edu>

by the experimenter, a calibration phase (to calibrate the eye trackers) and the watching of the two lecture videos; overall, participants spent about 35 minutes in the experiment. We conducted all experiments during daylight hours with both office lights and natural daylight contributing to our lighting.

The data generated by Tobii Studio during the study includes (among others) the estimated 2D coordinates of gaze points for each eye, the duration and coordinates of gaze events (i.e. fixations and saccades), the eye and pupil positions of the participant as well as the distance between the participant and the camera with a sample rate of 30 samples/second. In contrast, the data extracted from our webcam-based eye-tracking solution only includes the estimated 2D coordinates of gaze points of both eyes sampled at a rate of 5 samples/second.

4.3.2 Inattention Detection with Gaze Features

To realize eye-tracking based inattention detection using the professional eye tracker and our webcam-based solution, we turn the task into a standard supervised machine learning task. Our classifiers are trained using the aforementioned inattention reports as reference labels, and extracted gaze features for each time span between two bell signals as collected by either technique as input.

Given Tobii Studio's gaze data and inspired by [8, 9] we extracted 58 features in total. These features can be classified into two groups, global features and local features. The global features refer to features which are independent of the current content of the MOOC video, and are as shown in Table 4.1 based on fixations and saccades. The feature vector of a given bell time span covers statistical aggregates of fixation and saccade data such as maximum, minimum, mean, median, standard deviation, range, kurtosis and skew of fixation durations, saccade durations, saccade distance and saccade angles.

Local features are mainly based on the relationship between fixations/saccades and the areas of interest (AOIs) in the MOOC video, i.e. local features correlate gaze data with the current video content. There are certain areas of a video where a focused learner should focus her attention (e.g. the slides) in order to follow the content, while others are less interesting. While this opens a complex design space for engineering features, we opt for a simple implementation in which we manually define three fixed areas of interest: the instructor's face, subtitles, and the lecture slides. The resulting local features include then the number and length of saccades and fixations which focus on

different areas of interest for a given time span. Recall once more that all saccade and fixation data are computed by Tobii Studio with high precision for each bell time span based on a raw sample rate of 30 Hz.

Table 4.1: Features leveraged in the detection of participants' inattention

Feature Name	Explanation
Global Features	
Fixation Duration	the durations (ms) of fixations
Saccade Duration	the durations (ms) of saccades
Saccade Distance	the distances (pixel) of saccades
Saccade Angle	the angles (degree) between saccades and the horizon
Number of Saccade	total number of saccades
Horizontal Saccade Ratio	the proportion of the number of saccades which have saccade angles less than 30 degree
Fixation Saccade Ration	the ratio of the durations of fixations to the duration of saccades
Local Features	
Saccade Landing	the proportion of the number of saccades landing in different areas
Fixation Duration AOI	the durations (ms) of fixations located in different areas

Due to limitations of the WebGazer framework⁴, we only achieve a sample rate of 5 Hz for our webcam-based experiments. As changes of fixations and saccades usually happen within the range of 200 ms to 400 ms [107], reliable gaze data comparable to the one provided by the high-speed Tobii tracker is impossible to obtain using such a low sample rate and thus needs to be estimated algorithmically. For this purpose we implement micro-saccade detection as discussed in [30]: we first determine whether the movement between two consecutive gaze points is a saccade based on the movements' velocity. Then we treat gaze points between two saccades as a fixation. If there is only a single gaze point between two saccades, we assume this gaze point is a fixation with a duration between this gaze point and the previous gaze point. After the detection of saccades and fixations, we can generate the same 58 features as already shown in Table 4.1. Intuitively, the feature vectors from the webcam-based solution are less precise (as the sampling rate is much lower), however, we will show later that they still show comparable classification performance as we aggregate features over the time spans between consecutive bells, thus this imprecision carries little weight.

To train classifiers, we adopt leave-one-participant-out cross-validation [73]. In each run, the data of one participant is selected as test data and the data of all other participants is used for training. Based on the results reported in previous works [8, 9, 73], the collected data on learners' inattention is usually unbalanced with considerably less than 50% of probes resulting in reported inattention. We counter the effects of this imbalance by apply-

⁴It is based on an iterative algorithm that each detection runs after the previous detection is finished.

ing the oversampling method *Synthetic Minority Over-sampling Technique (SMOTE)* [17].

We have two requirements for our choice of classifiers as follows:

1. The selected models trained with our data can be used effectively to infer inattention in data of unseen participants.
2. The selected models trained with our data can be used in *real-time* inattention detection.

For the first requirement, we consider the bias-variance trade-off of machine learning models and the data size in our experiments. We select Logistic Regression, Linear SVM and Naive Bayes classifiers in our experiments as they have a low variance on small datasets like ours. These classifiers are also suitable for our second requirement.

Since the trained models are small and require few inference steps, they can easily be integrated into Web applications within MOOC platforms.

In order to determine the effect of different feature types, we evaluate different subsets of features in our experiments: (i) global features only (G), (ii) local features only (L) and (iii) the combination of global and local features ($G+L$). Since we also include SMOTE as a pre-processing step to deal with the unbalanced nature of our data, overall we report results on six different setups.

4.3.3 Research Questions

We address two main research questions:

RQ 3.1: How often do MOOC learners experience inattention within video watching?

RQ 3.2: How well do our webcam-based inattention detection methods perform?

For **RQ 3.2**, we first compare the overall effectiveness of our three selected classifiers with different sets of gaze features. Then, we delve deeper into the inattention detection results. Considering that the inattention reports are not evenly distributed among participants nor across the entire length of the lecture videos, we address two sub-questions **RQ 3.2.1** and **RQ 3.2.2**. A final sub-question is dedicated to the generalizability of our trained models.

RQ 3.2.1: Does inattention detection perform equally well across all participants?

RQ 3.2.2: Does inattention detection perform equally well across the entire length of a lecture video?

RQ 3.2.3: Does an inattention detection model trained on one video perform well to detect inattention on a different video?

4.4 Results

4.4.1 Exploratory Analysis of Inattention Reports

In order to answer **RQ 3.1**, we now analyze our participants' inattention behaviour while watching the two MOOC lecture videos.

In Figure 4.4, the distributions of participants' reported inattention events over the course of each of the two videos are shown. As discussed in the last section, participants were shown both videos in a random order, which is also reflected in the diagram. As the number of participants in each of the experimental groups is very small, no statistically significant conclusions can be drawn. However, it is visible that inattention is indeed a rather frequent occurrence even for very short video lectures of roughly 7 minutes: our measured inattention rate is 29%; i.e. in 71% of all bell time spans, our subjects actually stayed focused. In addition, it appears that our participants tire considerably during the second video when the experiment draws to its conclusion. This feedback was pro-actively provided by several of our participants in a post-experiment questionnaire, and seems to be at least anecdotally confirmed by the presented inattention reports.

4.4.2 Inattention Detection

In order to answer **RQ 3.2**, we investigate how accurately we can detect participants' inattention based on gaze data extracted by *WebGazer.js* compared to *Tobii X2-30*. The results are shown in Table 4.2. The results are based on the nested leave-one-participant-out cross-validation, which means that a leave-one-participant-out cross-validation is used as the inner cross-validation for model selection and a leave-one-participant-out cross-validation is used as the outer cross-validation for measuring performance of the selected model.

As a *baseline method*, we use a random classifier which includes the knowledge that the inattention rate is 0.29 and thus each feature vector is labeled as inattention with a probability of 0.29. Since accuracy is not a suitable metric for unbalanced data, the average precision, the average recall, and the average F1-measure of the nested leave-one-participant-out cross-validation are reported.

Based on our results in Table 4.2, all our methods are significantly better than the random baseline according to all three metrics. We do not observe a large impact of SMOTE: applying the SMOTE pre-processing method on Tobii data slightly increases *Precision*, however it has no effect on the detection results on WebGazer data. The combination of local and global features does not benefit the detection on Tobii data nor the detection on WebGazer data.

The highest F1 scores of each group of features are slightly lower than F1 scores reported by previous research [9] which relies on similar features and classifiers. We believe the difference (0.1 in F1 score) to be due to the slightly different data collection setup: Bixler et al. [9] utilize a short movie instead




Figure 4.4: Overview of the reported inattention reports across the MOOC videos. Due to the randomized video order in the experiment, we partition the results according to whether the video was shown first (“order 1”) or last (“order 2”). The video time displays the number of seconds since the start of the video.

Table 4.2: Inattention detection results based on gaze data (*G* means global features and *L* means local features)

Data	Feature	SMOTE	Avg.Precision	Avg.Recall	Avg.F1	Classifier
Baseline	-	-	0.290	0.291	0.290	-
Tobii	G	-	0.316	0.515	0.350	Logistic Regression
	G	-	0.224	0.197	0.132	Linear SVM
	G	-	0.343	0.462	0.325	Naive Bayes
	G	✓	0.358	0.487	0.336	Logistic Regression
	G	✓	0.207	0.345	0.250	Linear SVM
	G	✓	0.245	0.510	0.304	Naive Bayes
Data	L	-	0.263	0.625	0.309	Logistic Regression
	L	-	0.160	0.444	0.222	Linear SVM
	L	-	0.397	0.298	0.245	Naive Bayes
	L	✓	0.284	0.367	0.258	Logistic Regression
	L	✓	0.212	0.526	0.252	Linear SVM
	L	✓	0.294	0.682	0.364	Naive Bayes
Data	G+L	-	0.302	0.505	0.330	Logistic Regression
	G+L	-	0.115	0.369	0.167	Linear SVM
	G+L	-	0.342	0.486	0.335	Naive Bayes
	G+L	✓	0.301	0.487	0.310	Logistic Regression
	G+L	✓	0.346	0.502	0.330	Linear SVM
	G+L	✓	0.286	0.526	0.325	Naive Bayes
WebGazer	G	-	0.319	0.595	0.357	Logistic Regression
	G	-	0.059	0.308	0.093	Linear SVM
	G	-	0.309	0.671	0.395	Naive Bayes
	G	✓	0.239	0.384	0.246	Logistic Regression
	G	✓	0.232	0.607	0.319	Linear SVM
	G	✓	0.306	0.744	0.405	Naive Bayes
Data	L	-	0.294	0.534	0.366	Logistic Regression
	L	-	0.203	0.332	0.189	Linear SVM
	L	-	0.313	0.650	0.394	Naive Bayes
	L	✓	0.256	0.425	0.310	Logistic Regression
	L	✓	0.145	0.500	0.210	Linear SVM
	L	✓	0.320	0.691	0.403	Naive Bayes
Data	G+L	-	0.270	0.553	0.310	Logistic Regression
	G+L	-	0.154	0.429	0.217	Naive Bayes
	G+L	-	0.289	0.696	0.378	Naive Bayes
	G+L	✓	0.261	0.535	0.299	Logistic Regression
	G+L	✓	0.228	0.642	0.303	Naive Bayes
	G+L	✓	0.286	0.674	0.378	Naive Bayes

of MOOC lectures and free self-reporting instead of periodic self-reporting to obtain inattention reports. With respect to the evaluated classification methods, we find that the Gaussian Naive Bayes models outperform the other approaches on WebGazer data in every feature set combination.

The most surprising finding in this experiment is that compared to the Tobii data we achieve higher *Recall* and *F1* scores based on the gaze features extracted from WebGazer data. Based on our intuition, features extracted from the data which is generated from the high-quality eye tracker *Tobii X2-30* should lead to a more accurate detection of inattention, than features extracted from the data which is generated by a standard webcam. A possible reason for this experimental artifact is the small number of participants in our study; in future work we plan increase our participant pool to at least 100 participants.

Based on Table 4.2, we now delve deeper into our inattention detection results. In order to answer **RQ 3.2.1**, we investigate the detection results on each participant separately. For this step, we select the best-performing models for each data source (based on F1 scores reported in Table 4.2). For the detection on Tobii data, we use Gaussian Naive Bayes with local features and the SMOTE method. For the detection on WebGazer data, we use Gaussian Naive Bayes with global features and the SMOTE method. The results are shown in Table 4.3. We observe that across all metrics, the minimum observed accuracy is zero (for both Tobii and WebGazer data), which implies that there are participants for whom our prediction is not working at all. At the same time, we observe that at best a participant’s inattention can be detected with high accuracy with an F1 of 0.7 (Tobii data) and 0.8 (WebGazer data) respectively. The large standard deviations across the three metrics - 0.2 to 0.35 - further show that the accuracy of our detector varies widely between participants. Therefore, we conclude that *the detection does not work equally well for all participants in our experiments*.

Table 4.3: Statistics of detection results on individual participants ($P_{highest}$ shows the detection results of the participant with highest F1-measure, P_{lowest} with lowest)

Data	Metrics	Max	Min	Mean	Std	$P_{highest}$	P_{lowest}
Tobii Data	Precision	0.714	0	0.294	0.198	0.600	0
	Recall	1.000	0	0.682	0.357	0.857	0
	F1	0.706	0	0.364	0.200	0.706	0
WebGazer Data	Precision	0.700	0	0.306	0.209	0.700	0
	Recall	1.000	0	0.744	0.354	1.000	0
	F1	0.824	0	0.405	0.244	0.824	0

Based on the analysis in § 4.4.1, we find that inattention is not evenly distributed throughout a video. This leads to our **RQ 3.2.2**. We split each

Table 4.4: Detection across the entire length of the video (*Part 1* means the first half part of the video, and *Part 2* means the second half part of the video)

Data	Metrics	Solar Energy		Nuclear Energy	
		Part 1	Part 2	Part 1	Part 2
Tobii Data	Precision	0.147	0.410	0.276	0.321
	Recall	0.308	0.763	0.397	0.462
	F1	0.195	0.474	0.285	0.369
WebGazer Data	Precision	0.365	0.240	0.295	0.327
	Recall	0.615	0.500	0.462	0.615
	F1	0.438	0.285	0.344	0.416

video into two parts with the same length. Then, for each part of the video, we use the data of the other part and the data of the other video to train the model and to detect the inattention in this specific left-out part of the video. The models, feature sets and the SMOTE method used in this experiments are same as in **RQ 3.2.1**. The results are shown in Table 4.4.

We conclude that *the detection of inattention cannot be made equally well across the entire length of the lecture videos in our experiments*. For Tobii data, we find the results of the inattention detection in the second part of the same video to be much better than the first part. For WebGazer data, we observe no trend, the results vary depending on the lecture video. We hypothesize this result to be connected to the fact that different participants were shown the videos in different orders.

Our last experiment answers **RQ 3.2.3**. So far we have shown that our method can detect a participant’s inattention based on a model trained on the gaze data and inattention reports of other participants. To scale out, we need to determine to what extent we can detect learners’ inattention in video lectures of one course with a model trained in lecture videos of other courses. If we were to obtain good detection results for such scenarios, there may be a general model which can be used in different lecture videos at scale (i.e., “train once, deploy everywhere”). In this experiment, the experimental settings for classifiers, feature sets and the SMOTE method on different kinds of data are same as in our previous experiments (**RQ 3.2.1** and **RQ 3.2.2**). We evaluate the cross-video performance by training our model on one video, and test the performance of the model using the other video. The results of all video combinations are shown in Table 4.5. For reference, this table also includes training and testing using the same video, using leave-one-participant-out cross-validation.

Based on results in Table 4.5, we find the model trained on WebGazer data to be more robust to a change of video context than the model trained

Table 4.5: Detection with model translation (i.e. using a model on a different video than it was trained on)

Data	Metrics	Trained on Solar		Trained on Nuclear	
		Used in Solar	Used in Nuclear	Used in Nuclear	Used in Solar
Tobii Data	Precision	0.267	0.171	0.294	0.149
	Recall	0.705	0.372	0.410	0.205
	F1	0.355	0.229	0.296	0.150
WebGazer Data	Precision	0.240	0.298	0.346	0.344
	Recall	0.679	0.692	0.596	0.667
	F1	0.317	0.401	0.392	0.423

on Tobii data. We also observe that it does matter whether we train on video A and test on B or vice versa as results are comparable. Overall, we believe that *a model trained on the WebGazer data collected on one video can lead to good predictions in other videos*, at least if the videos share similarities with respect to style and type as in our scenario.

4.5 Conclusions

In our work, to evaluate short-term behavioral engagement of MOOC learners, we focus on tracking learners' inattention in video watching. We compare the effectiveness of a webcam plus the open-source library *WebGazer.js* to the effectiveness of the specialized (and expensive) *Tobii X2-30* for the task of inattention detection in a lab study. In our experiments, we could show that the accuracy of our webcam-based approach is on par with the specialized eye-tracking device. This opens the way for large-scale experiments in real-world MOOCs, allowing for both investigating learners' inattention behavior and investigating the effectiveness of interventions based on inattention detection in future research under realistic conditions.

Our work has a number of limitations, including the small pool of participants all sharing similar educational backgrounds. Similarly, the number of evaluated MOOC videos is very limited and both videos have a comparable (but very common) style. Thus, it is unclear how well our approach can be applied to completely different types of videos or user groups. In addition, we rely on a number of established and straightforward-to-implement features; we expect a further boost in detection accuracy when more sophisticated features are introduced.

A core contribution provided by our work is the published repository of data collected during our controlled lab study. In addition to including

the inattention reports of our experiment's participants, we also provide the full set of gaze data obtained by the *Tobii X2-30* and our webcam on our companion Web page [146].

Chapter 5

Face-Tracking Based Inattention Detection

In this chapter, we follow Chapter 4 and investigate short-term behavioral engagement of MOOC learners based on learner inattention in video watching. We propose a face-tracking approach to improve the accuracy and speed of inattention detection in near real-time. In recent years, researchers have begun to exploit eye-tracking and gaze data generated from webcams as part of complex machine learning solutions to detect inattention or loss of focus. Those approaches however tend to have a high detection lag (e.g. 30 seconds usually), can be inaccurate (e.g. detection accuracies of 14% – 35% reported in Chapter 4), or are complex to design and maintain (e.g. the processing pipeline in Chapter 4). In contrast, we explore the possibility of a simple alternative—the presence or absence of a face—to detect a loss of attention in MOOC learning in this chapter.

This chapter is published as “Webcam-based Attention Tracking in Online Learning: A Feasibility Study” [100], by Tarmo Robal, Yue Zhao, Christoph Lofi, and Claudia Hauff, 23rd International Conference on Intelligent User Interfaces, pp. 189-197. ACM, 2018.

5.1 Introduction

In recent years, a number of works have investigated inattention prediction based on various signals, including heart-rate data [142], EEG data [144], skin conductance and temperature [10] as well as computer mouse pressure data [141]. While insightful, none of these approaches can be applied at scale in an online learning environment in the near future and thus, most of the existing research on inattention detection relies on eye-tracking data, including [5, 8, 9, 73, 22, 115, 147]. In our study in Chapter 4, inattention detection with webcam-based eye tracking is investigated. Here, the eye-mind link [96] is exploited as the eye gaze usually correlates well with a person’s focus.

A major issue of existing eye-tracking based inattention detection approaches is the lack of *real-time* detection capabilities (30 – 60 second delays are common) (e.g. our study in Chapter 4). An additional point of concern in our setting are the privacy requirements of MOOC environments—to ensure a learner’s privacy all necessary computations should be conducted within the learner’s *browser environment* (the alternative approach of streaming a learner’s webcam data to a high-performance server has severe privacy implications, while requiring the installation of dedicated software packages hampers usability).

In this chapter, we explore a significantly simpler alternative approach towards detecting inattention whilst learning in a MOOC environment: we use the departure of a user’s *face* from the webcam’s viewport as a proxy for learner inattention—a user whose face is not aimed at the screen is unlikely to pay attention to a video playing on it. It turns out, that even this deceptively simple detection task is challenging in a MOOC environment where we have to consider widely varying consumer-grade hardware and browser software. In this chapter we conduct an extensive study involving two open-source browser-based software frameworks for gaze and face detection, *WebGazer.js* and *tracking.js*, as well as a third hardware-based solution (a *Tobii X2-30* eye tracker) to determine an upper performance bound². Both software-based frameworks can be integrated into current MOOC environments, and perform all their processing on the user’s computer without the need for a server infrastructure or additional browser plugins. We benchmark the ability of the three frameworks to reliably detect a user’s focus towards the screen content (using the presence/absence of a face as proxy) across a variety

²Both *WebGazer.js* and *Tobii X2-30* are used for eye-tracking based inattention detection in Chapter 4.

of common MOOC user activities such as watching a MOOC video whilst leaning on one’s hand, checking the weather report on a smart-phone or drinking coffee. Specifically, we address the following research question in our work:

RQ 3.2: How well do our webcam-based inattention detection methods perform?

To test inattention detection methods with different software frameworks, we compile a benchmark suite of 50 typical MOOC learner activities, partitioned into activities that are indicative of (i) focus, (ii) certain loss of focus and (iii) likely loss of focus. We conduct an extensive lab study involving *tracking.js* and *WebGazer.js* as well as a professional eye tracker (our upper bound in terms of performance). A total of 20 study participants execute the benchmark suite of activities in a controlled environment.

We find that in our setup, *tracking.js* performs significantly better than *WebGazer.js*, achieving a median detection accuracy of 62% across all fifty tasks (for the most difficult task detection accuracy was 17%), with the professional hardware-based eye tracker achieving a median accuracy of 72.5% (the most difficult task resulted in 27% accuracy). The observed detection delay is below 2 seconds for *tracking.js*, making it a viable choice for webcam-based attention detection (using face detection as a proxy). At the same time, the reported accuracy numbers suggest that current software and hardware solutions still struggle to provide a consistently high detection quality across all tasks.

5.2 Background

Different data collection methods have been used to study the loss of attention of learners in traditional classrooms since the 1960s, such as the observation of inattention behaviors [53], the retention of course content [68], using direct probes in class [126, 62], or relying on self-reports from learners [14]. A common belief was that learners’ attention might decrease considerably after 10–15 minutes of the lecture, which was supported by [126]. However, Wilson and Korn [140] challenge this claim and argue that more research is needed. In a recent study, Bunce et al. [14] asked learners to report their attention loss voluntarily during 9 – 12 minute course segments. Three buttons were placed in front of each learner, representing attention lapses of 1 minute or less, of 2 – 3 minutes and of 5 minutes or more. During the lectures, the

learners were asked to report their loss of attention by pressing one of three buttons once they *noticed* their attention loss. This leads Bunce et al. [14] to conclude that learners start losing their attention early on in the lecture and may cycle through several attention states within the 9 – 12 minute course segments.

In online learning environments, attention loss may be even more frequent. Risko et al. [98] used three 1-hour video lectures with different topics (i.e. psychology, economics, and classics) in their experiments. While watching the videos, participants were probed four times throughout each video. The attention-loss frequency among the participants was found to be 43%. Additionally, Risko et al. [98] find a significant negative correlation between test performance and loss of attention. Szpunar et al. [128] investigate the impact of interpolated tests on learners' loss of attention within online lectures. The study participants were asked to watch a 21-minute video lecture (4 segments with 5.5 minutes per segment) and report their loss of attention in response to random probes (one probe per segment). The inattention frequency reported in their experiments was about 40%. Loh et al. [63] also employ probes to measure learners' loss of attention and find a positive correlation between media multitasking activity and learners' loss of attention (average frequency of 32%) whilst watching video lectures. Based on these considerably high inattention frequencies, we conclude that reducing learner inattention in online learning is an important approach to improve learning outcomes.

In traditional classroom contexts, a teacher has the ability to detect and regain learner attention through various pedagogical approaches. This is not applicable in MOOC environments due to the nature of online learning. Various technological approaches have been explored to detect and record signals of user (in)attention in the past besides eye tracking, including heart-rate tracking through mobile cameras [142], brain activities through EEG analysis [144], skin conductance and temperature [10], posture and body pressure sensing and pressure applied on a computer mouse [141]. As already implied, most of the existing research though focus on either face or eye-gaze detection [5, 8, 9, 73, 22, 115, 147].

Inspired by the eye-mind link effect [96], a number of previous studies [8, 9, 73] focus on the automatic detection of learners' loss of attention by means of gaze data. In [8, 9], Bixler and D'Mello investigate the detection of learners' loss of attention during computerized reading. To generate the ground truth, the study participants were asked to manually report their loss of attention when an auditory probe (i.e. a beep) was triggered. Based on

those reports, the loss of attention frequency ranged from 24% to 30%. During the experiment, gaze data was collected using a dedicated eye tracker. In [73], Mills et al. asked the study participants to watch a 32-minute, non-educational movie and self-report their loss of attention. In order to detect loss of attention automatically, statistical features and the relationship between gaze and video content were considered. In contrast to [8, 9], Mills et al. mainly focus on the relationship between a participant’s gaze and areas of interest, i.e. specific areas in the video a participant should be interested in. In Chapter 4, we present our method for detecting learner inattention similar to the studies in [73], but our method is adapted and optimized for the MOOC setting.

All mentioned approaches relying on the eye-mind link share two common issues: they are usually unable to provide real-time feedback as they are trained on eye-gaze recordings with sparse manually provided labels (e.g. most approaches have a label frequency of 30–60 seconds, which directly translates into a detection delay of similar length), and the reported accuracy is too low for practical application (e.g. Chapter 4 reports that the detection accuracy ranges from 14% to 35% depending on training and video). As a result, we choose a different approach as discussed in the following sections.

5.3 Eye/Face-Tracking Frameworks

Recall that we employ face presence and absence as proxies of learner attention and inattention respectively. Next to explicit face-tracking software frameworks, eye-tracking frameworks are suitable for our work as well, as in the absence of a face, no eye tracking is possible.

In order to determine an upper performance bound, we use the professional high-end hardware eye tracker *Tobii X2-30*. *Tobii X2-30* uses its own proprietary analytic software Tobii Studio to analyze the gathered eye-tracking data.

Although there exist a number of different eye/face-tracking software solutions, our choice is limited by the typical MOOC environment (which runs within the browser, and thus we require browser-based software frameworks), privacy aspects (all computations have to be performed on the user’s device) and the variety of hardware capabilities we can expect MOOC learners’ devices to have (the computations should not require too many resources). Evidently, JavaScript-based solutions fit the task description. Libraries such as *CCV.js*, *headtrackr*, *ObjectDetect*, *tracking.js*, and *WebGazer.js* (with *clm-*

trackr for face detection) are thus potential candidates. After an initial testing phase of all mentioned frameworks we settle on two suitable ones: *WebGazer.js* [83] and *tracking.js*³.

5.3.1 WebGazer.js

WebGazer.js is an open-source eye-tracking library written in JavaScript that is able to infer eye-gaze locations in real-time. Use-case specific extensions, e.g. to track users' web search behaviour [84] exist as well. *WebGazer.js* can be configured with different components to track gaze, pupils, or faces. We use the *clmtrackr* component⁴, a face fitting library (referred to as *CLM* in the following), which has previously been used among others in works on camera-based emotion detection [110], and intelligent public displays in city environments [80]. *CLM* tracks a face and the coordinate positions of a face model, as shown in Figure 5.1. Using this face model, *WebGazer.js* can extrapolate the user's gaze (i.e. the point of the screen on which a user's gaze focuses) by estimating the face's distance and orientation from the screen. A weakness of *CLM* is its "aggressive" face-fitting algorithm that often attempts to fit a model even when no face is present. This leads to many potential problems where random background elements (like posters, plants, furniture) are mistaken for faces, and sometimes even preferred over a real user's face clearly visible in the camera's viewport.

5.3.2 Tracking.js

tracking.js is a JavaScript-based face-tracking library (*TJS* in the following), which has been employed, among others, in security systems for identity verification [46] and object recognition tasks [130]. With respect to eye/face-tracking, this library offers a significantly less powerful feature set than both *Tobii X2-30* and *CLM*, as it can only detect the presence and location of the boundary box of an object—in our case the face—in a video stream (see Figure 5.2). While it can also be employed to track the eyes' locations (but not the gaze), we do not use that feature in this study. We hypothesize that the simplicity of *TJS* leads to more reliable face presence and face absence detection.

³<https://trackingjs.com>

⁴<https://github.com/auduno/clmtrackr>



Figure 5.1: Face fitting model generated by *CLM*. This example shows a common face fitting error due to hand positioning.

5.3.3 Detecting Face-Miss Events

We define a *face-miss* event to be an event of a user’s face turning or moving away from the computer screen. The differences in the three evaluated frameworks (*Tobii*, *CLM*, and *TJS*) leads to different heuristics for detecting a *face-miss*:

Tobii: A *face-miss* event is detected if the proprietary Tobii Studio software cannot determine gaze point coordinates. This usually represents a problem with detecting the users’ eyes by the tracker hardware (e.g. they are not within the camera viewport, they are closed, or obstructed by an object). At times, while the eyes can be found by *Tobii*, no gaze coordinates can be determined as the gaze direction is unclear. We cannot distinguish this case from a case where there is no face at all. In our experience, the presence of gaze coordinates is a very reliable proxy for the presence of a face (low false positive rate), while the lack of coordinates does not necessarily imply the absence of a face.




Figure 5.2: Face Boundary Box detected by tracking.js

CLM: Similar to *Tobii*, we define a *face-miss* event as the software’s inability to fix exact gaze coordinates, which also means a failure in reliable face detection in case of *CLM*. In contrast to *Tobii*, due to the aggressiveness of the face fitting algorithm, *CLM* is quite prone to detect faces where in reality there are none present (high false positive rate).

TJS: We define a *face-miss* event as the library’s inability to fix a face boundary box in the webcam’s video stream. We do not try to track eyes or gaze.

The video or eye tracker stream is continuously processed while it is recorded. The *Tobii* system relies on dedicated hardware support for this task (which partially contributes to its high retail price), and is thus able to guarantee a sampling rate of 30 samples per second mostly independent of the computer hardware. For the webcam-based solutions, image processing of the video stream needs to be handled by the system’s CPU and within the browser’s environment. As a result, only low sampling rates are possible without overwhelming low-end computer systems. For this reason, we fix the

sampling rate to 4 samples per second. However, due to the unreliability of the JavaScript timer events under high system loads, the standard deviation of the targeted sampling time of 250 ms is 48 ms in our experiments (described further in §5.4). Furthermore, we have extreme cases where the sampling times increased up to 1,157 ms, i.e. less than one sample per second. Therefore, *Tobii* should be able to react with significantly lower delays than the webcam-based frameworks.

5.4 Methodology

In order to evaluate the suitability of the chosen webcam toolkits for face and gaze tracking, we develop a **benchmark set of tasks**, which we argue represent common behaviours of online learners in front of their laptops. For each of the tasks we define the desired behaviour: the eye-tracking devices should either report the loss of the face/gaze (in the case of *face-miss* tasks) or keep detecting the face/gaze (in the case of *face-hit* tasks). We exclude mobile learners from these tasks as desktop learners are still the vast majority of learners in today’s MOOC environments⁵.

We design a total of 50 tasks together with a small sample of regular MOOC learners (graduate students in our research lab). These tasks are—to some extent—abstract versions of the behaviour MOOC learners exhibit when watching lecture videos, one of the most common activities in xMOOCs. The task descriptions we developed are shown in Table 5.3. They fall under three broad categories:

Face-miss tasks: describe those user behaviours that *should* result in the loss of a detected face/gaze. 21 tasks belong to this category; examples include *Take a sip from the cup [next to you] while turning away from the camera* or *Look straight up to the ceiling for 8 seconds*.

Likely-face-miss tasks: should result in our frameworks reporting a mix of face hit and face miss samples. Two examples among the 14 tasks in this category are *Lean back and put your hands behind your neck for 5 seconds* and *Draw a square on the paper*.

Face-hit tasks: describe user behaviours that should not influence our frameworks’ ability to detect the face, though they may influence gaze detection.

⁵Concretely, based on a sample of twenty edX MOOCs offered at Delft University of Technology, fewer than 20% of learners accessed the course content via mobile devices.

15 tasks belong to this category, for example *Reposition yourself in the chair* and *Stare at the camera for 3 seconds*.

We develop a dedicated Web application as testing ground. The 50 tasks are presented as virtual “cue cards” to study participants and both *TJS* and *CLM* are included as webcam-based eye/face-tracking solutions. The design of the application is modular, additional frameworks can easily be evaluated as well. We have open-sourced our application at https://github.com/trx350/xMOOC_benchmark.

Welcome to the Intellieye Pilot Study #1

In the following you will be requested to perform different activities you could be engaged with while otherwise normally watching a video. You will only be shown short instructions what to do and not the associated activity itself. Please read the instructions in yellow and wait for the sound alert.

Act only and immediately after the bell ring sounds. Test the sound here:

Resume focusing for the next task after a 'ding' sound. Test the sound here:

Thank you for your co-operation!

Prior to study, please select the values that best describe the study environment:

I am wearing:

The background behind me is:

The light in the room is majority:

Calibrate the system beforehand: * as instructed

Experiment ID: Prediction rate:

Figure 5.3: Opening screen of the user study

The opening screen of the application is shown in Figure 5.3; an example task cue card is shown in Figure 5.4. The task order is randomized. The procedure for each task Q_i is the same: the task description is shown and five seconds later a bell sound indicates the start of the task at time $t_{start}^{Q_i}$: at the sound of the bell the participant is expected to perform the task. Another bell sound (different to the one indicating the start) indicates to the participant when the task has been finished at time $t_{end}^{Q_i}$, and this is followed by the next task description. Task durations differ, depending on the specific task, e.g. Q_{31} requires a participant to look at a certain angle for 5 seconds while Q_{39} asks a participant to check his or her phone for 10 seconds.

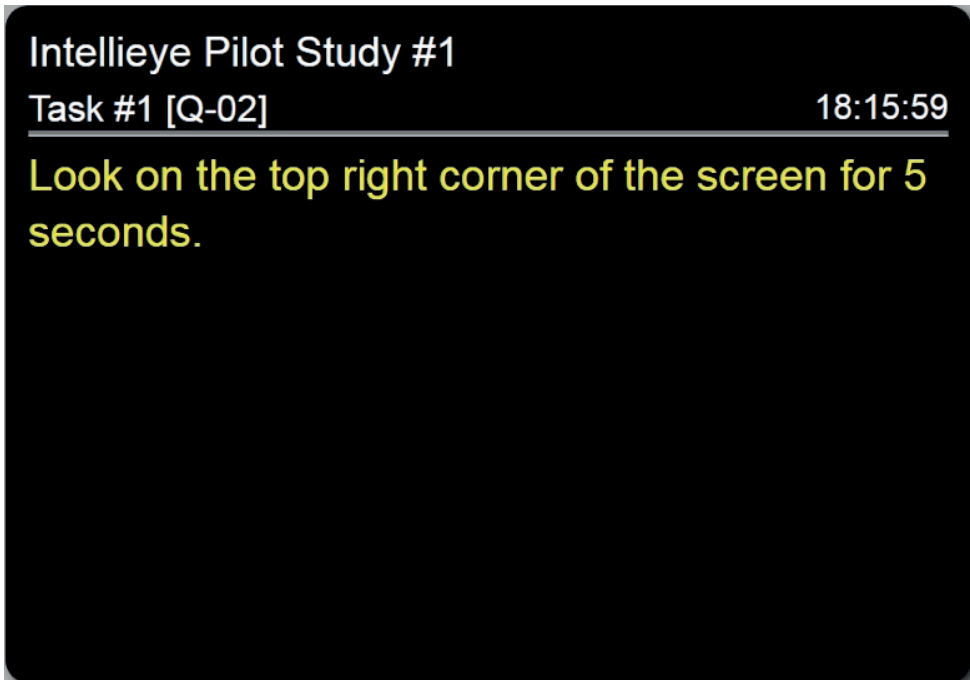


Figure 5.4: Example task “cue card” of the user study.

5.4.1 Study Setup

We conducted all our experiments on a Dell Inspiron 5759 laptop (with built-in webcam situated in the center of the top screen bezel) with a 17-inch screen and a $1,920 \times 1,080$ resolution running Windows 10. The Tobii eye tracker was placed on the lower screen bezel.

The study was conducted across a one week period: 20 participants were recruited among graduate students and staff members of Delft University of Technology via email lists. The participants did not receive compensation and spent less than an hour on this study. Among 20 participants, 9 wore glasses and 2 had contact lenses. In 10 of the sessions the background behind the test subject had a uniform (light) color, in another 10 cases a poster or photographic background was observed. We recorded these settings in our study as we had conducted preliminary experiments which indicated that eye trackers (especially the software-based ones) can be misled by noisy backgrounds.

As this is a controlled study, in order to facilitate the proper execution of the tasks, the participants were provided with the necessary tools to perform

all necessary behaviours, including a sheet of paper and a pen (required for Q_{22} , Q_{24} & Q_{25}), a cup (Q_{41} & Q_{42}) and a phone (Q_{39}).

The Tobii requires a calibration step which participants concluded at the start of the study. The *CLM* framework can also be calibrated in a light-weight manner: five red dots were shown on the screen that have to be clicked one after the other. To test the effect of the calibration we randomly switched on the calibration step for 8 of the 20 learners.

To prepare the participants for the tasks, each participant was trained on two tasks before the start of the actual study. The participants were reminded repeatedly to only start executing a task’s required behaviour after the sound of the bell and to keep executing the behaviour until the ending sound occurred.

5.4.2 Detection Accuracy

For every task and participant, we determine the eye trackers’ *face-hit/face-miss* predictions from the collected logs between the $t_{start}^{Q_i}$ and $t_{end}^{Q_i}$ timestamps. As the eye trackers vary in their sampling rate they all produce a differing amount of labels (face-hit, face-miss) for each sample interval. We evaluate the accuracy of the produced labels by computing the percentage of correct predictions (as defined by the type of task) in the task interval. For example, in a 5-second task slot the webcam-based approach takes a sample once every 250 ms (on average), and thus we collect approximately 20 predictions. For a *face-miss* task, if 14 of the 20 predictions are a miss, the detection accuracy will be 70%. Lastly, we average the accuracy for each task across all participants.

Table 5.1: Tobii’s delay between the start of a face-miss/likely-face-miss task and the first face-miss event. The data is averaged across all participants of a single task.

Delay (Seconds)	1	2	3	4	5+
% of Tasks	53%	28%	6%	3%	9%

Table 5.2: Overview of the impact of the participants’ background on *TJS*’s and *Tobii*’s accuracy.

Background	#Participants	Accuracy in %	
		<i>TJS</i>	<i>Tobii</i>
Solid light	10	61.5	68.6
Poster/photo	10	55.7	67.8

Table 5.3: Overview of all fifty benchmark tasks, and the accuracy (in %) of *CLM*, *TJS* and *Tobii* averaged across the 20 participants in our user study. For “(Likely) Face-Miss” tasks we report the percentage of detected face misses, i.e. the eye tracker flags frames as not containing a face. For “Face-Hit” tasks we report the percentage of detected face hits. A higher percentage indicate a better performance. The best performance per task is shown in bold. † Note that for tasks Q_1 and Q_2 Tobii’s camera was not covered and the detection reflects participants’ hand moving through Tobii’s camera viewport to cover the webcam on the top bezel of the experimental laptop; for task Q_{43} there was no gaze to detect for Tobii.

QID	Task	Accuracy in %		
		<i>CLM</i>	<i>TJS</i>	<i>Tobii</i>
Face-Miss Tasks				
Q_1	Cover the camera for 2 seconds	12	45	†7
Q_2	Cover the camera for 5 seconds	28	73	†17
Q_3	Cover your face with both hands for 5 seconds	17	67	75
Q_4	Look what is under your table (3 sec)	3	64	81
Q_5	Stand up for 5 seconds	10	68	71
Q_{20}	Tilt your head to the right for 3 seconds	15	59	38
Q_{21}	Check if there is a HDMI port on the laptop	12	56	77
Q_{26}	Look straight up to the ceiling for 8 seconds	12	72	92
Q_{27}	Tilt your head back for 5 seconds (face ceiling)	10	68	84
Q_{28}	Tilt your head back for 2 seconds (face ceiling)	5	51	66
Q_{29}	Look down for 3 seconds	4	35	78
Q_{32}	Look left for 2 seconds	7	50	72
Q_{33}	Look left for 8 seconds	14	69	88
Q_{35}	Look over your right shoulder	13	50	72
Q_{36}	Look right for 10 seconds	13	77	90
Q_{37}	Look right for 3 seconds	14	64	79
Q_{38}	Look right for 5 seconds	7	63	83
Q_{39}	Check your phone for 10 seconds	7	42	89
Q_{40}	Check your phone, return after the ding	13	37	87
Q_{42}	Take a sip from the cup while turning away from the camera, return after the ding	5	40	51
Q_{47}	Look up and return immediately	8	49	68
Likely Face-Miss Tasks				
Q_6	Lean back and put your hands behind your neck for 5 seconds	2	67	63
Q_7	Lean closer to the screen and immediately back	3	17	27
Q_{13}	Rapidly lean back and forth until the ding sounds	6	37	57
Q_{18}	Tilt your body to the left and stay for 3 seconds	13	50	57
Q_{19}	Tilt your body to the right and return immediately	6	41	55
Q_{22}	Draw a square on the paper	9	45	67
Q_{23}	Write down 5 keys left from letter A, focus back to the screen only after the ding	4	19	61
Q_{24}	Write down a sentence about weather	15	47	73
Q_{25}	Write down <i>I love Intel!ieye!</i>	10	45	78
Q_{30}	Look half-left and return	7	36	64
Q_{31}	Look half-right and stay for about 5 seconds	7	42	77
Q_{41}	Face the camera and take a sip from the cup until you hear the ding	8	30	35
Q_{46}	Cover the left side of your face with left hand over cheek and eye	8	38	43
Q_{48}	Look around in the room to every direction	10	63	82
Face-Hit Tasks				
Q_8	Open browser and navigate to www.weather.com. Return after the ding. (15 sec)	94	97	80
Q_9	Open new browser tab and return to this after the ding	95	89	87
Q_{10}	Open some program window on top of study window and return after the ding	99	87	94
Q_{11}	Feeling sleepy? Yawn and cover your mouth with a hand. (3 sec)	94	66	64
Q_{12}	Grab the tip of your nose for 3 seconds	100	64	71
Q_{14}	Reposition yourself in the chair	98	77	61
Q_{15}	Scratch the top of your head (or nape) for 3 seconds	94	69	85
Q_{16}	Scratch the lower part of your left leg for 2 seconds	93	79	64
Q_{17}	Slowly lean back and stay for about 2 seconds	96	32	38
Q_{34}	Look on the top right corner of your screen for 5 seconds	95	86	96
Q_{43}	Rest your eyes for 5 seconds (close them)	95	84	14†
Q_{44}	Scratch your left cheek for 3 seconds	95	74	89
Q_{45}	Sit still and face the camera for 5 seconds	94	87	90
Q_{49}	Grab your ears with both of your hands for 3 seconds	95	76	85
Q_{50}	Stare at the camera for 3 seconds	95	89	88

5.5 Results

To answer **RQ 3.2**, in this section we report the outcomes of our user study along three dimensions: (i) accuracy across tasks, (ii) reaction times and (iii) the influence of the participants’ background on the accuracy levels.

5.5.1 Accuracy

The first question we consider is the accuracy of the three eye trackers under investigation across the 50 tasks of our benchmark suite. Table 5.3 lists the detection accuracy for each task, aggregated across the 20 study participants. As expected, *Tobii* achieves the highest accuracy, with an average of 68.2% across all tasks. Among the two software solutions, *TJS* clearly outperforms *CLM*, achieving an average accuracy of 58.6% compared to *CLM*’s 35.4%. If we were only to focus on the tasks where face misses and likely face misses form the ground truth, *CLM*’s accuracy would drop to 9.6%. The reason for this poor performance is *CLM*’s approach to face and gaze detection: it will try to match anything in the video frame to a potential face area, a separate face detection phase is not performed. This also explains its high accuracies in the face hits tasks. Note that the calibration step performed by some of our participants for *CLM* does not result in a different outcome.

The comparison between *Tobii* and *TJS* shows a relatively small performance gap between the webcam-based face tracker and the high-end device. While *Tobii* outperforms *TJS* in 39 of the 50 tasks, in many instances the difference in accuracies is rather small. Using *Tobii* as a reference point, *TJS* is able to conform with 77.8% of *Tobii*’s detected labels.

Due to the clear performance differences between *TJS* and *CLM*, in further analyses we focus exclusively on *TJS* and its performance compared to *Tobii*.

5.5.2 Reaction Times

As one of the potential reasons for *TJS*’s lag in performance compared to *Tobii* we investigate the reaction times of both users and frameworks. More specifically, we measure the delay between the *instructed* start time of the task (i.e. the timestamp $t_{start}^{Q_i}$) and the first time a framework detects a face-miss. This time delta of course consists of both the user delay (i.e. the time it took for the study participant to finally start performing the task, which for some tasks—e.g. Q_{23} & Q_{46} —show a considerable delay) and the

actual detection delay imposed by the framework. We average the delays of all participants for a task and report the percentage of tasks whose average delay is up to 1 second, up to 2 seconds, etc. in Table 5.1. For the majority of tasks, *Tobii* is able to detect the first face-miss within 1 second of the start of the task.

The *Tobii* eye tracker runs with a very high fixed sampling rate of 30 samples per second, and is mostly unaffected by the current CPU load of the host machine. Therefore, we make the assumption that the delays in Table 5.1 represent the user delay. In contrast, *TJS* and *CLM* can have very low sampling rates depending on the current system load (we aim at 4 samples per second, but we also experienced significantly lower rates). By comparing the times of detecting the first face-miss of both *TJS* and *CLM* with *Tobii*, we can obtain an intuition of the delays imposed by those frameworks. For *TJS*, this resulted in a delay of 0.6 ± 1.1 seconds, and for *CLM* in 1.3 ± 1.0 seconds. While these detection delays are not instantaneous, the delays are short enough for practical applications.

5.5.3 Background as an Influencing Factor

As we conduct the user study in different rooms on different times of the day, we also record our participants with various backgrounds. In Table 5.2 we partition our participants according to the background they sat in front of during the study. All participants reported their background to be either of a solid light color (as present in many offices) or contain a poster and/or photo. This factor had an impact on the eye trackers' accuracy: while *Tobii*'s accuracy remained unaffected by the background, the *TJS* eye tracker considerably degraded when the background was noisy.

5.6 Conclusions

In this chapter, to evaluate short-term behavioral engagement of MOOC learners, we present a face-tracking method for the inattention detection. To enable real-time attention tracking in a standard MOOC environment, we introduce the presence or absence of a face in a learner's webcam viewport as a simple proxy of learner attention or inattention to improve the accuracy and speed of the inattention detection method proposed in Chapter 4.

We compare three potential technical solutions for this task: using the high-end professional eye tracker *Tobii X2-30*, and using two software-based

solutions that analyze the video stream of a consumer-grade webcam. We conduct a lab study with 20 participants, who had to perform a controlled benchmark suite of 50 realistic tasks, which introduced several challenging factors such as body movement, partially covering the face, noisy backgrounds, and crooked body postures. This benchmark suite and the accompanying Web application allows for a standardized and fair comparison of different approaches for *face-hit* and *face-miss* detection, and we provide it under an open-source license to foster future research.

Our experiments show that the professional dedicated hardware solution outperforms the open-source software-based solutions both in respect to detection performance and processing speed, but is of course unsuitable for a large-scale deployment outside of a controlled lab setting. For the software-based solutions which can indeed run on typical hardware used by MOOC learners, the complicated *CLM* gaze tracking as employed by *WebGazer.js* introduces many complications, resulting in poor detection performance both for the presence and absence of a user's face. In contrast, the face-tracking library *TJS* shows significantly higher performance for nearly all benchmark tasks. Additionally, both software libraries incur an additional time delay of around 1 – 2 seconds over the nearly instantaneous detection response of the hardware solution. With careful design, this delay should be easily manageable in a future MOOC learner attention detection component.

In our future work, we plan an implementation of an attention tracker suitable for a large-scale MOOC deployment on the basis of the *TJS* framework. Beyond purely technical or methodical challenges, this allows us to tackle additional interesting research questions: Would MOOC learners be willing to accept and use such an attention detection tool? What are the reasons why they would like/or refuse to use such technology? And of course finally, if learners accept the use of such tools, does this indeed positively impact their learning outcomes?

Chapter 6

Near Real-Time Inattention Detection Widget

In this chapter, to evaluate short-term behavioral engagement of MOOC learners, we follow Chapter 4 and Chapter 5 and focus on learner inattention in video watching. We design and deploy a webcam-based inattention detection widget *IntelliEye* in a real-world MOOCs based on the face-tracking method proposed in Chapter 5.

IntelliEye is a privacy-aware system that makes use of learners’ webcam feeds to determine—in near real-time—when they no longer pay attention to the lecture videos. *IntelliEye* makes learners aware of their attention loss via visual and auditory cues. We deploy *IntelliEye* in a real-world MOOC and explore to what extent MOOC learners accept it as part of their learning and to what extent it influences learner behaviour.

This chapter is published as “IntelliEye: Enhancing MOOC Learners’ Video Watching Experience through Real-Time Attention Tracking” [99], by Tarmo Robal, Yue Zhao, Christoph Lofi, and Claudia Hauff, in Proceedings of the 29th on Hypertext and Social Media, pp. 106-114. ACM, 2018.

6.1 Introduction

In this chapter, we present *IntelliEye*², a system we design to directly tackle the “loss of attention” issue during MOOC lecture video watching by detecting it in real-time and alerting the learner to it.

How exactly can we detect learners’ loss of attention *in real-time* and *at scale*? How can we alert the learner to her loss of focus? One answer to these questions lies in the ubiquitous availability of webcams in today’s laptops: *IntelliEye* employs the webcam feed to observe learners’ activities during their time on the MOOC platform and intervenes (e.g. by delivering an auditory signal) if it detects a loss of focus. All of these actions are performed by *IntelliEye* in a *privacy-aware* manner: none of the data or computations leaves a user’s machine. Prior studies [8, 9, 73, 147, 100] exploit eye/face tracking to determine a user’s attention state, though these studies are either conducted with commercial high-quality hardware eye-tracking devices and/or well-settled experimental lab conditions (e.g. our studies in Chapter 4 and Chapter 5). In contrast, in this chapter we make use of commonly available webcams and deploy *IntelliEye* “in the wild”, to 2,612 MOOC learners in an actual MOOC, instead of a controlled lab study.

We conduct our analyses of *IntelliEye*’s use along three dimensions: (1) the **technological capabilities** of MOOC learners’ hardware, (2) the **acceptance** of *IntelliEye* by MOOC learners, and, (3) the **effect** of *IntelliEye* on MOOC learners’ behaviour. Specifically, we investigate the following research questions:

RQ 3.3: To what extent is MOOC learners’ hardware capable to enable the usage of technologically advanced widgets such as *IntelliEye*?

RQ 3.4: To what extent do MOOC learners accept technology that is designed to aid their learning but at the same time is likely to be perceived as privacy-invading (even though it is not)? Are certain types of MOOC learners (e.g. young learners, or highly educated ones) more likely to accept this technology than others?

RQ 3.5: What impact does *IntelliEye* have on learners’ behaviours and actions? To what extent does *IntelliEye* affect learners’ video watching behaviour?

Our main findings can be summarized as follows:

²*IntelliEye* is open-sourced at <https://github.com/Yue-ZHAO/IntelliEye>.

- We find that most learners (78%) use hardware and software setups which are capable to support such widgets, making the wide-spread adoption of our approach realistic from a technological point of view.
- The majority of learners (67%) with capable setups is reluctant to allow the use of webcam-based attention tracking techniques, citing as main reasons privacy concerns and the lack of perceived usefulness of such a tool.
- Among the learners using *IntelliEye* we observe (i) high levels of inattention (on average one inattention episode occurs every 36 seconds—a significantly higher rate than reported in previous lab studies) and (ii) an adaptation of learners' behaviour towards the technology (learners in conditions that disturb the learner when inattention occurs exhibit fewer inattention episodes than learners in a condition that provides less disturbance).

6.2 Related Work

6.2.1 Attention Loss in Learning

Identifying and tracking learners' loss of attention in the classroom has been explored in a myriad of ways since the 1960s, including the analysis of students' notes [44, 65], the observation of inattention behaviors (by observers, stationed at the back of the classroom) [53], the retention of course content [68], probes (requiring participants to record their attention at particular given points in time) [126, 62] and self-reports (requiring participants to report when they become aware of their loss of attention) [14]. A common belief was that learners' attention might decrease considerably after 10 – 15 minutes into the lecture [126]. Wilson and Korn [140] challenge this claim and argue that more research is needed, a call picked up by Bunce et al. [14] who find that learners start losing their attention early on in higher-education lectures and may cycle through several attention states within 9 – 12 minute course segments.

With the advent of online learning, the issue of attention loss, how to measure it and how it compares to classroom attention lapses receive renewed attention. Different studies show that in online learning environments (often simulated in lab settings where participants watched lecture videos), attention lapses may be even more frequent than in traditional classroom contexts.

Risko et al. [98] used three 1-hour video lectures with various topics (i.e. psychology, economics, and classics) in their experiments, probing participants four times throughout each video. The attention-loss frequency was found to be 43%. In addition, Risko et al. report a significant negative correlation between test performance and loss of attention. Szpunar et al. [128] study the impact of interpolated tests on learners' loss of attention within online lectures, asking participants to watch a 21-minute video lecture (4 segments with 5.5 minutes per segment) and report their loss of attention in response to random probes (one per segment). In their experiments, the loss of attention frequency was about 40%. Loh et al. [63] also apply probes to measure learners' loss of attention, finding a positive correlation between media multitasking activity and learners' loss of attention (average frequency of 32%) whilst watching video lectures. Based on these considerably high loss of attention frequencies, we conclude that reducing loss of attention in online learning is an important approach to improve learning outcomes.

6.2.2 Automatic Detection of Attention Loss

Inspired by the eye-mind link effect [96], a number of previous studies [8, 9, 73] focus on the automatic detection of learners' loss of attention by means of gaze data. In [8, 9], Bixler and D'Mello investigate the detection of learners' loss of attention during computerized reading. To generate the ground truth, the study participants were asked to manually report their loss of attention when an auditory probe (i.e. a beep) was triggered. Based on those reports, the loss of attention frequency ranged from 24% to 30%. During the experiment, gaze data was collected using a dedicated eye-tracker. In contrast to [8, 9], Mills et al. [73] mainly focus on the relationship between a participant's gaze and areas of interest (AOIs), specific areas in the video a participant should be interested in. Mills et al. asked study participants to watch a 32-minute, non-educational movie and self-report their loss of attention throughout. In order to detect loss of attention automatically, statistical features and the relationship between gaze and video content were considered.

In Chapter 4, we present a method to detect inattention similar to the studies in [73], but optimized for a MOOC setting (including the use of a webcam alongside a high-quality eye-tracker). All mentioned approaches relying on the eye-mind link share two common issues: (i) they are usually unable to provide real-time feedback as they are trained on eye-gaze recordings with sparse manually provided labels (e.g. most approaches have a label frequency of 30 – 60 seconds, which directly translates into a detection delay of similar length), and (ii) the reported accuracy is too low for

practical application (e.g. Chapter 4 reports that the detection accuracy is 14% – 35%). We note that besides the eye-mind link, another recent direction is the use of heart rate data (measured for instance by tracking fingertip transparency changes [89]) to infer learners’ attention. Lastly, our study in Chapter 5 presents a face-tracking approach to detection learner inattention and designs a benchmark with a series of learner activities related to the attention/inattention of MOOC learners.

6.2.3 MOOC Interventions

We now discuss MOOC interventions, especially those geared towards video watching and towards improving self-regulated learning. Existing research on MOOC videos is largely concerned with the question of what makes a MOOC video engaging and attractive to learners; examples include the overlay of an instructor’s face over the lecture slides [57], shorter video segments instead of one long lecture video [41], and the overlay of an instructor’s gaze to enable learners to more easily follow the video content [116].

Few studies consider the issue of self-regulated learning in MOOCs, largely because this requires approaches that are personalized and reactive towards each individual learner. Simply informing learners about the best strategies for self-regulated learning at the beginning of a MOOC is not sufficient [58]. Davis et al. [27] design a visual “personalized feedback system” that enable learners to learn how well they were doing compared to successful passers from a previous MOOC edition (in terms of time spent on the platform, their summative assessment scores and so on). This comparison, even though this feedback moment was rare (once a week), enabled learners to self-regulate their learning better, leading to significantly higher completion rates for learners exposed to the feedback system. A prior study by Davis et al. [26] indicate that non-compliance among learners is a difficult obstacle in very simple interventions: the authors had included an extra question in each week of a MOOC, asking learners to write about their study plans (and thus make learners think about those plans). Few learners saw the benefit of this question (it was ungraded) and thus very few complied.

Overall, we have shown that attention lapses are a regular occurrence in the classroom and occur with even greater frequency in online learning, where learners are prone to digital multitasking. We have also presented some drawbacks of sophisticated eye-tracking based inattention detectors (accuracy and timeliness of detection) and finally we have pointed out the difficulty of bringing self-regulated learning into the MOOC scenario due to

learners’ non-compliance. In response to these findings, we design *IntelliEye*, a robust inattention (by using face detection based on our study in Chapter 5) detector that requires no additional actions by the learners beyond what they usually do on a MOOC platform, provides personalized feedback, is privacy-aware and detects a loss of attention in near real-time (with at most 2 seconds delay).

6.3 *IntelliEye*

6.3.1 Architecture

The goal of *IntelliEye* is to provide real-time feedback on learner’s attention, and is based on a set of heuristics *reliably implementable* on a wide variety of hardware setups: (1) if the browser tab/window containing the lecture video is not visible to the learner, *IntelliEye* triggers an inattention event; (2) we assume a learner is inattentive if her face cannot be detected for a period of time, i.e. we employ face tracking as a robust proxy of attention tracking³; (3) if the face-tracking module detects a loss of the face we consider the mouse movements as a safety check: if no face is detected but the mouse is being moved, no event is triggered.

The resulting high-level architecture is shown in Figure 6.1. *IntelliEye* is implemented in JavaScript, as the edX platform allows custom JavaScript to be embedded in course modules—thus providing us with an easy way to “ship” *IntelliEye* to all learners in our MOOC. As visible in Figure 6.1, *IntelliEye* resides exclusively on the client to ensure learners’ privacy; usage logs are sent to our dedicated *IntelliEye* log server for the purpose of evaluating *IntelliEye*, though this communication is not necessary for *IntelliEye* to function. This setup requires *IntelliEye* to be light-weight and resource-saving as all computations are carried out on the learner’s device and within the resource limits of a common Web browser. We now describe the seven architecture modules that *IntelliEye* consists off.

Profiling Module

In order to provide a smooth user experience for MOOC learners, we limit the full usage of *IntelliEye* to devices that fulfill certain device setup re-

³We note that this is a lower-bound for inattention, as learners watching the video may still not pay attention.

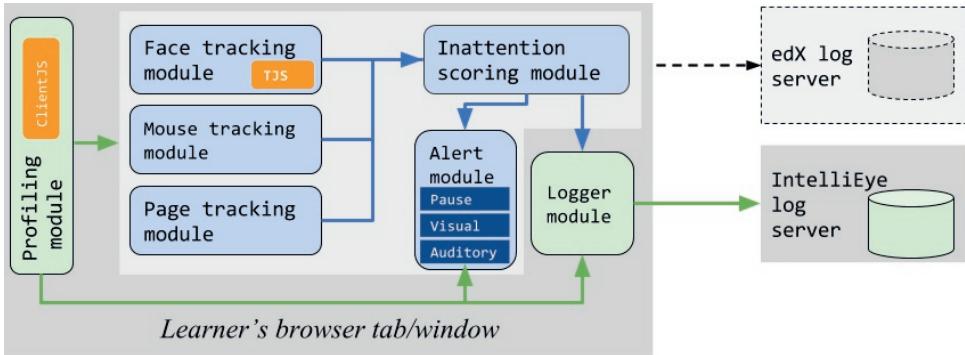


Figure 6.1: *IntelliEye*'s high-level architecture. The profiling and logger modules are always active; the attention tracking and alerting modules are only enabled if supported setup is detected and learner has granted access to webcam feed.

quirements, a situation we call *supported setup*. We rely on the `ClientJS`⁴ library to determine the device type, operating system and browser version of the learner's device and activate the inattention tracking modules only if a supported setup is detected. The requirements are as follows:

1. The device is not a mobile device and is not running iOS or Android, due to their incompatibility with *IntelliEye*.
2. The browser used is either: Chrome 54+ (i.e. version 54 or higher), Firefox 45+ or Opera 41+ to ensure the availability of JavaScript dependencies necessary for *IntelliEye*.
3. The device has at least one usable webcam as detected via the Media Capture and Streams API.

If the profiling yields an *unsupported setup*, a log entry is sent to our *IntelliEye* log server and no further modules are activated.

The profiling module is also responsible for extracting the learner's edX user ID, which in turn determines which alert type the learner receives in our experiments.

Face-Tracking Module

In *IntelliEye*, we use face tracking to proxy inattention detection, thus aiming at overcoming the reported shortcomings of gaze tracking with respect to

⁴<https://github.com/jackspirou/clientjs>

response time and reliability: if a learner's face is not visible in front of the screen when a lecture video is playing, we argue that she is likely not paying attention.

We choose the open-source library *tracking.js* [64] (or *TJS* for short) for this purpose. The detection accuracy and the delay of *tracking.js* are evaluated in Chapter 5 based on a benchmark with 50 behaviours that learners typically execute in front of their computer (e.g. *Check your phone*; *Look right for 10 seconds*; or *Reposition yourself in the chair*). *TJS* has a competitive accuracy: it is able to detect 77.8% of the face hit/face miss behaviours that the *Tobii X2-30* was identifying correctly. The delay of *TJS* in detecting inattention is $0.6 \pm 1.1s$.

The module performs face presence detection (via *TJS*) from the webcam feed every 250 ms and reports a boolean (face present or absent) to the *Inattention scoring module*. We choose this time interval not to overburden the computational resources of the learner's device.

Mouse Tracking Module

This module acts as a sanity check for the face-tracking module: if the face-tracking module reports loss of a face and the learner is still moving the mouse in the active MOOC window, we assume that the face-tracking module misclassified the situation and do not raise an inattention alert. This module tracks the absence or presence of mouse movements every 250 ms and reports it to the *Inattention scoring module*.

Page Tracking Module

This module tracks the visibility of the browser window or tab that contains the edX page (and thus the lecture video) using the *document.hidden()* Web API call. A value is produced every 250 ms and forwarded to the *Inattention scoring module*.

Inattention Scoring Module

This module estimates inattention of a learner by aggregating the data obtained from the tracking modules based on the heuristics already introduced at the start of § 6.3.1: a learner is inattentive if her face is not trackable unless there is mouse movement and the video player browser window is visible. The input from the three scoring modules is aggregated over a sliding time

Algorithm 1 Inattention detection mechanism in *IntelliEye*

Require: \mathcal{F} , \mathcal{M} , \mathcal{V} , \mathcal{L} —threshold value, \mathcal{S} —scores for \mathcal{F} , \mathcal{M} , \mathcal{V}

$\mathcal{T} = (t_1, t_2, \dots, t_k)$ score queue of the trending functionality;

- 1: $inAttention \leftarrow False$
- 2: $n \leftarrow 20$
- 3: $S_F \leftarrow \sum_i f_{n-i}(n-i)/n$
- 4: $S_M \leftarrow \sum_i m_{n-i}(n-i)/n$
- 5: $S_V \leftarrow v_n$
- 6: $trend_F \leftarrow 0$
- 7: $T.dequeue(t_1); T.enqueue(t_k \leftarrow S_F)$
- 8: $(t_k > t_{k-1}) \Rightarrow trend_F \leftarrow 1$
- 9: $(t_k < t_{k-1}) \wedge (t_{k-1} < t_{k-2}) \Rightarrow trend_F \leftarrow -1$
- 10: $Q \leftarrow (S_F < \mathcal{L} \wedge trend_F < 1)$
- 11: $(Q \wedge S_M < \mathcal{L} \wedge S_V) \vee (Q \wedge \neg S_V) \vee (S_F > \mathcal{L} \wedge \neg S_V) \Rightarrow inAttention \leftarrow True$

window of 5 seconds—we chose this time window based on our user study with 50 typical activities during MOOC video watching, where we found the longest activity to take approximately 5 seconds. Recall that each module has a fixed sampling rate of 250 ms, and thus our sliding window takes into account 20 measurement points from each tracking module.

More formally, the input to this module are the boolean values (i) for face presence $\mathcal{F} = (\dots, f_{n-20}, f_{n-19}, \dots, f_n)$, (ii) mouse movement $\mathcal{M} = (\dots, m_{n-20}, m_{n-19}, \dots, m_n)$, and (iii) page visibility $\mathcal{V} = (\dots, v_{n-20}, v_{n-19}, \dots, v_n)$. To conserve computational resources, the module computes the attention state once a second. Algorithm 1 outlines the inattention decision process employed by the *Inattention Scoring module*. In essence, a weighted score is computed for the face presence and mouse movement values (lines 3 & 4), giving higher weights to more recent values. The visibility score of the video window is simply the last recorded value (line 5). Lines 6-9 compute face-tracking trends over time. The role of the face-tracking trend computation is to minimize the volume of false positives driven by learner behaviour, in particular sudden movements, bad position in front of the webcam, or a temporary short time failure of *TJS* in detecting the face in webcam video feed. Lines 10-11 show the rules the module employs to determine inattention based on the predefined threshold (which represents the minimum accepted score that is considered as attention, in our case $\mathcal{L} = 2.92$), computed scores and the trend. The threshold and rules are another outcome of our user study—they led to the highest accuracy in distinguishing between attention and inattention behaviours [100].

Note that the level of thresholding (\mathcal{L}) determines the sensitivity of *IntelliEye*—lowering the value will make the system less rigorous, increasing this value will on the other hand increase system responsiveness to learner behaviour.

Alert Module

We explore three different mechanisms—with varying levels of disruption—to raise learners’ awareness about their detected loss of attention; none of these requiring an action from the user beyond returning their attention to the video at hand. In our experiment each learner was assigned to a single alert type, depending on their edX user ID detected by the *Profiling module*.

Pausing the video: When attention loss is detected *IntelliEye* will pause the currently playing lecture video. Once *IntelliEye* detects re-gained attention on the video, playing is resumed. At what position playing is resumed depends on *how long* the learner was not paying attention since pausing. The video is rewound to between 0 and 10 seconds before the attention loss was detected; we define three different configurations: (i) if the inattention period is less than 1.5 seconds, the video continues from where it was paused as it would be annoying for a learner to review content just seen and available in her short-time memory, but also to avoid repetitive ‘rewind-and-play’ situations; (ii) if the inattention lasts more than 10 seconds, the video is rewound 10 seconds which is the approximate lower level of human short-time memory (reported in between 10 – 30 seconds [88, 72]); and (iii) in all other cases it is rewound 3 seconds—rewind a little for rapid recall in case of distraction. This scheme ensures that the video will restart at a familiar point for the learner. The drawback of this mechanism is the severity of false alerts as the video will pause and thus the learner is disturbed if inattention was falsely determined.

Auditory alert: In this setup, the video keeps playing but an additional sound effect (a bell ring) is played repeatedly as long as inattention is detected. This setup is not as “annoying” as falsely pausing the video, but can still substantially disturb the learner.

Visual alert: In this version, *IntelliEye* visually alerts the learner by repeatedly flashing a red border around the video as long as inattention is detected. Figure 6.3 shows an example of this alert. This scenario is the least intrusive in case *IntelliEye* falsely detects inattention. It may also be

the least effective, as learners who look away from the screen or minimize the browser tab/window will not be able to view the alert.

Logger Module

This module is responsible for logging *IntelliEye*'s usage. These logs are sent to our dedicated log server. Specifically, the following actions lead to logging (for log entries with categorical values we list all possible values within {...}):

Loading: When *IntelliEye* is loaded due to a learner accessing a course subsection⁵ containing one or more video units we log (timestamp, alertType {pause, visual, auditory}, userID, deviceSetup).

Video status change: Every change in the video's status (e.g. from paused to play) for a learner with supported setup leads to a log of the form (videoID, timestamp, videoStatus {play, pause, seek, end}, videoTime, videoLength, videoSpeed, subtitles {on, off}, fullScreen {on, off}). The videoTime entry refers to the point in time within the video the status changed.

***IntelliEye* status change:** When a learner with a supported setup changes the status of *IntelliEye* (e.g. from disabled to enabled), we log (videoID, timestamp, videoTime, videoLength, IntelliEyeStatus {allow, disallow, start, pause, resume, end}). Information on the video is logged as most interactions with *IntelliEye* occur within the edX video player (cf. § 6.3.2).

Inattention status change: This log event occurs when the attention status of a learner with a supported setup changes: (videoID, timestamp, videoTime, videoLength, inattention {start, stop}). Here, start indicates that inattention has been detected. The next event is generated when the status changes back to attention again (stop). As long as the inattention state is maintained, no further log events are generated.

Finally, we note that beyond the *IntelliEye* logs (cf. Figure 6.1), we also have access to the official edX logs, which contain information on all common actions learners perform within a MOOC on the edX platform such as quiz submissions, forum entries, clicks, views, and so on—data we use in some of our analyses.

⁵A set of course elements semantically belonging together, cf. § 6.4.

6.3.2 User Interface

Having described *IntelliEye*'s architecture, we now turn to its user interface. Figure 6.2 shows *IntelliEye*'s welcome screen (potentially shown every time a MOOC learner opens a course subsection with one or more video units), describing its capabilities, and the positive impact it can have on learning. The learner has four choices: (i) to enable *IntelliEye* for this particular video only, (ii) to disable *IntelliEye* for this video only, (iii) to enable *IntelliEye* for all videos, and, (iv) to disable *IntelliEye* for all videos. If a learner opts for (iv), we ask her for the feedback on the decision (“*You have disabled IntelliEye. Please tell us why.*”).

Once a learner enables *IntelliEye*, the face-tracking module attempts to access the webcam feed, which in all supported browsers triggers a dialogue controlled by the browser (*Will you allow edx.org to use your camera?*); once the learner chooses *Allow*, *IntelliEye* is fully functioning.

Figure 6.3 shows how *IntelliEye* embeds itself in the edX video player. Here the learner can return to the welcome screen and change her enable/disable decisions (via the “eye” icon) and switch *IntelliEye* on or off on the fly. *IntelliEye*'s status is visible at all times: either ‘Active’ (*IntelliEye* is enabled, the video is not playing at the moment), ‘Playing’ (*IntelliEye* is enabled), or ‘Not Active’ (*IntelliEye* is disabled). Note that this change in the video player interface is only visible to learners with a supported setup. Learners on non-supported setups will receive the original edX video player without alterations.

6.4 MOOC Setting


We deployed *IntelliEye* in the MOOC *Introduction to Aeronautical Engineering (AE1110x)* offered by Delft University of Technology on the edX platform. The MOOC's target population are learners who are looking for a first introduction to this particular field of engineering. The MOOC requires around 80 – 90 hours of work and consists of 104 videos and 332 automatically graded summative assessment questions. The MOOC is *self-paced*, that is, the MOOC is available for learners to enroll for up to 11 months. In contrast to the more common six to ten week MOOCs, learners can set their own schedule and their own pace. The MOOC was opened for enrollment on May 1, 2017 and remained so until March 31, 2018. *IntelliEye* was deployed for ten weeks (October 5, 2017 to December 17, 2017); it was available for all videos within the MOOC. A total of 2,612 different learners visited the



IntelliEye: an experimental add-on to improve your learning

Imagine someone looking over your shoulder while you learn in this MOOC, reminding you to pay attention and alerting you when you become distracted. This would probably make you learn more efficiently!

IntelliEye is a first step towards this vision: an intelligent video player add-on we have developed at the Delft University of Technology. It will become active when you watch a lecture video: whenever the add-on detects a loss of focus on your part it will visually alert you by repeatedly flashing a red border around the video until it detects your focus again. IntelliEye makes use of your Webcam to track your focus and attention. IntelliEye is privacy-aware: none of the Webcam data leaves your computer, all computations are made on your device.



How does IntelliEye do it? [Learn more](#)

What do I need to do to use IntelliEye? [Learn more](#)

Who do I contact in case of concerns/issues/ideas? [Learn more](#)

Anything else? [Learn more](#)

Enable IntelliEye

Disable IntelliEye

Always ask me

Remember my choice

Figure 6.2: *IntelliEye* welcome screen.

MOOC during the deployment period and were exposed to *IntelliEye*. We deployed *IntelliEye* in three different variants according to the manner of alerting learners to their lack of attention: video pause, auditory alert and visual alert (§ 6.3.1). We conducted an inter-subject study: each learner was randomly assigned (based on their learner ID) to one of the three conditions. Once assigned, a learner remained in that condition throughout the experiment. Table 6.1 shows the distribution of the 2,612 learners across the three conditions.

Before turning to the analyses section, we introduce the relevant concepts and definitions:

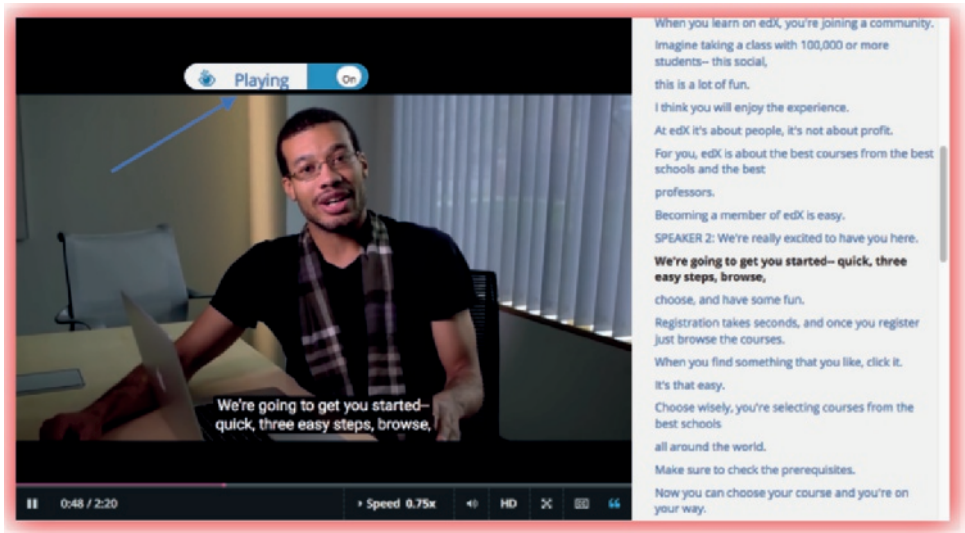


Figure 6.3: *IntelliEye*'s video player interface (arrow) embedded in the edX video player widget. The red hue around the video player is the visual alert we experiment with.

Course subsection: on the edX platform, a course subsection refers to a sequence of course units (such as video units, quiz units and text units) that are grouped together, most likely because they all relate to the same topic. As an example, one of the subsections in our MOOC consists of the following sequence: video →video →text →quiz →video →quiz →text.

Session: refers to a sequence of logs from a single learner (active on a single device), with no more than 30 minutes time difference between consecutive log entries. This means that after 30 minutes of inactivity in the MOOC, we assume a new “learning” session starts (if the learner becomes active again). We combine the logs we retrieved from our *IntelliEye* log server with those collected by edX.

Supported session: refers to a session with a supported setup.

Unsupported session: refers to a session without a supported setup.

Video session: refers to a session in which at least one video was being played by the learner, regardless of the length of video playing.

***IntelliEye* session:** refers to a supported session which is also a video session, and in which *IntelliEye* was running (which means that the learner did accept the terms of use and played a video while *IntelliEye* was active).

Non-*IntelliEye* session: refers to a supported session which is also a video session, and in which *IntelliEye* was not active while the video was playing (this either means that the learner did not accept the terms of use, or manually disabled *IntelliEye*).

6.5 Empirical Evaluation

In this section, we answer three research questions **RQ 3.3**, **RQ 3.4**, and **RQ 3.5** respectively.

6.5.1 Technological Capabilities

The first question we consider is to what extent our MOOC learners (who, according to their edX profiles, hail from 138 different countries) have a supported device setup (**RQ 3.3**): according to Table 6.1, 78% of learners (across all three alert types) logged in at least once with a device supported in *IntelliEye*. Among those 563 learners (22%) who did not have a supported session, 223 of them only accessed the course with a mobile device (that is 9% of the overall learner population). If we drill down on the 340 learners with unsupported sessions on non-mobile devices, the most common reason is an outdated browser we do not support (e.g. Chrome 52, IE 11, Safari 10 and Safari 11), followed by the lack of a webcam (in 118 cases). We do not observe a particular skew towards certain countries or regions; learners from India (104 learners) and learners from the US (93 learners) have the largest number of unsupported setups, which are also the two countries where most learners hailed from (484 learners from India and 334 from the US).

Table 6.1: Learners exposed to *IntelliEye*. Shown is the number of learners: (i) in each alert type condition, (ii) with at least one session with supported setup, (iii) who used *IntelliEye* at least once, and (iv) not accepting *IntelliEye*.

Alert Types	#Exposed Learners	#Learners with 1+ Supported Sessions	#Learners with 1+ <i>IntelliEye</i> Session	#Learners without <i>IntelliEye</i> Session
Video pause	861	681	214	467
Auditory alert	902	703	208	495
Visual alert	849	665	236	429
Total	2,612	2,049	658	1,391
% of Total	—	78%	25%	53%




Figure 6.4: Distribution of video sessions and unique learners. A learner may be listed in more than one session type.

6.5.2 Acceptance of *IntelliEye*

Having established that our hardware requirements are reasonable, we now turn to *IntelliEye*'s acceptance, i.e., are learners willing to enable a widget which observes them via a webcam (**RQ 3.4**). As Table 6.1 shows, 32% of learners (658 out of 2,049) with at least one supported session activated *IntelliEye* at least once.

We had two hypotheses on who engages with our intervention: (1) younger learners are more likely to engage than older ones, and (2) more active learners are more likely to engage than less active ones. To explore these hypotheses we compute various metrics for three different user groups (learners that do not engage with *IntelliEye*, learners that have one or two *IntelliEye* sessions and learners that have three or more *IntelliEye* sessions) as shown in Table 6.2⁶. We observe significant differences across almost all metrics (the exception being age) between those learners not (or hardly) using *IntelliEye* and those using *IntelliEye* three or more times. The number of learners in each group though—highly skewed with more than 1,600 learners in the not/hardly using *IntelliEye* groups and 35 learners in the remaining group—has to serve here as a point of caution. Based on these results, *IntelliEye* appears to be used most often by learners who are already engaged—a finding which is inline with prior MOOC interventions, e.g. [26, 27].

Next, we consider the use of *IntelliEye* across time (Figure 6.4): for each day of our experiment we plot the number of learners exposed to *IntelliEye* and whether they had *IntelliEye* or non-*IntelliEye* sessions. The usage of *IntelliEye* neither increases nor decreases significantly over time.

In Table 6.3, we take a look at learners' decisions of enabling or disabling *IntelliEye* in subsequent video sessions. Learners that enabled *IntelliEye* in a video session, did so again with a probability of 0.35 (6% of learners chose to enable *IntelliEye* for all sessions, 29% chose to enable *IntelliEye* for just the next video session). After enabling *IntelliEye* in a video session, 21% decided to permanently disable *IntelliEye* in the next session. We discuss the main reasons for this decision at the end of this section. Learners that disabled *IntelliEye* in their video session were very unlikely to change their decision in the next video session with 97% of learners sticking to their disable decision.

Next, we consider for *how long* learners were using *IntelliEye* during their video sessions: did they use *IntelliEye* continuously or did they disable it after some time? For all the *IntelliEye* sessions in which *IntelliEye* was

⁶Note that all our analyses consider the 74 days of *IntelliEye*'s deployment only, i.e. the number of sessions, the quiz scores, etc. are only computed for that time period.

Table 6.2: Learner attributes partitioned according to the use of *IntelliEye* (choices made on welcome page are not considered in grouping). Only learners with at least one supported video session are considered. * indicates Student's t-test significance at $p < 0.05$ level. † and ‡ indicate Mann-Whitney U test significance at $p < 0.05$ and $p < 0.01$ levels respectively.

Statistics	# <i>IntelliEye</i> Sessions		
	None	1-2	3+
#learners	1,030	623	35
Median age	23	*None 21	22
Median prior education	*None Associate degree	High school	*None High school
Median av. session length (min)	27.77	27.44	†None,1-2 35.17
Median #sessions	3	3	‡None,1-2 12
Median quiz score	3.0	†None 3.0	‡None,1-2 7.0
Median minutes video watching	21.78	21.87	‡None,1-2 102.82
Median minutes on platform	94.56	90.83	‡None,1-2 542.04

enabled initially (725 sessions from 557 distinct learners), we condense the video session time (which includes video watching as well as other activities on the platform) to video watching time only, based on the edX log data. We then proceed to determine whether *IntelliEye* was consistently enabled throughout, or whether it was disabled in the first, second or the last third of the video. We find (as shown in Table 6.4) that mostly *IntelliEye* was either switched off very early or employed throughout a session. Few learners disabled it well into the video watching experience (beyond the first third of the video). Learners that received the pause alert were more likely to disable *IntelliEye* than learners in the other alert groups; learners in the visual alert condition were most likely to keep *IntelliEye* enabled, reflecting the various levels of disturbance the alerts cause.

Table 6.3: *IntelliEye* usage transition probabilities between subsequent video sessions; E=Enabled, D=Disabled, EF=Enabled Forever, DF=Disabled Forever.

Decision v(i)	Decision v(i + 1)			
	E	D	EF	DF
<i>IntelliEye</i> enabled	0.29	0.43	0.06	0.21
<i>IntelliEye</i> disabled	0.03	0.68	0.00	0.28

As a last analysis of this research question, we focus on the reasons learners provided when disabling *IntelliEye*. Of the 938 learners (248 of them

have at least one *IntelliEye* session) who chose to disable *IntelliEye* forever, 379 provided us with reasons for their decision. With an open card sort, we sort the provided reasons into eight categories shown in Table 6.5. As the vast majority of learners reported a single reason, for the few (< 10) learners who provided a number of reasons we select the one they were most vocal about. Most commonly (35%) learners cited themselves as not needing help to self-regulate their learning (*I never lose my attention because the lecture and the whole course are very interesting.*).

22% of the learners mentioned a non-functioning webcam (e.g. *Because my camera doesn't work well; webcam and audio are easily accessible with WebRTC so I cover and disable it.*), followed by 17% with privacy concerns (e.g. *I feel awkward being observed and controlled.; I don't like the idea of having the webcam on.*) and 9% with *IntelliEye* not performing as expected⁷. Interestingly, conscious multitasking was mentioned several times (*I'm multitasking while doing this.*), showing that at least some learners were very much aware of their learning behaviour and what *IntelliEye* was supposed to do for them. Among the 27 learners who reported being disturbed by the alerts, 12 learners received the pause and 12 learners the auditory alert. Overall, this feedback shows that *IntelliEye* works reasonably well (only 34 out of 248 learners using *IntelliEye* at least once reported issues) and that the largest issue facing future use of *IntelliEye* is learners' *perception* of not requiring an attention tracker during their learning, followed by privacy concerns.

Table 6.4: Number of sessions with *IntelliEye* initially enabled grouped by the time it is switched off in the session.

Disabled during	Pause	Auditory Alert	Visual Alert
1st third of a session	48%	44%	35%
2nd third of a session	6%	10%	7%
Last third of a session	7%	6%	6%
Enabled throughout	39%	39%	52%
Total #sessions	242	207	276

6.5.3 Impact of *IntelliEye*

We now investigate the impact of *IntelliEye* on learners over time and explore whether learners change their video watching behaviour over time (**RQ**

⁷We note that one possible reason is our lack of a calibration step: to make *IntelliEye* easy to use and accessible we did not impose one; *IntelliEye* assumes the learner to be facing the screen and the webcam.

Table 6.5: Reasons provided for disabling *IntelliEye* forever.

Reason	#Learners	[%]
Attention tracking not perceived as useful/needed	131	35%
webcam not functioning	83	22%
Privacy concerns	64	17%
<i>IntelliEye</i> not working well	34	9%
Disturbed by alerts	27	7%
Conscious facing away from the screen	14	4%
Hardware/Internet connection too slow	14	4%
Conscious multitasking	6	2%
Uncomfortable feeling	6	2%
Σ	379	40%
<i>No reason provided</i>	559	60%

3.5). Specifically, we consider all learners with at least two *IntelliEye* sessions (the most active learner in our dataset has six *IntelliEye* sessions); for each learner we bin her sessions into two bins (the first half and the second half). We then proceed to compute for each bin (i) the average number of minutes lecture videos were played, (ii) the average attention duration and inattention duration detected by *IntelliEye*, and, (iii) the average number of inattention alerts occurring per minute of video watching. The results are shown in Table 6.6. Recall that according to the literature, inattention occurs frequently in video watching, though the manner of investigating this (through probes issued at certain times to study participants) [98, 128, 63] does not allow us to draw minute-by-minute conclusions. In contrast, in our work we can now make a statement to this effect: the average number of inattention alerts varies between 0.84 and 2.86 per minute (the latter means that on average a learner gets distracted every 21 seconds in the visual alert condition!). Across all conditions, on average 1.65 inattention alerts are triggered per minute (i.e. one every 36 seconds on average). Interestingly, learners are quickly able to adapt their behaviour towards the offered technology: while the learners in the visual alert type are often alerted (in a manner that is easy to ignore), the learners in the auditory alert conditions receive significantly fewer alerts (cf. row *Mean #inattention per min*); similarly, learners in the pause and auditory alert conditions have significantly shorter inattention spans (cf. row *Mean avg. inattention duration*) than those in the visual condition. As learners were assigned to the conditions randomly we are confident that this behavioural adaptation is due to the different types of alerts.

When comparing the statistics for the two session bins (to detect trends over time), we do not observe a significant decrease over time in the number

Table 6.6: Overview of the impact of *IntelliEye* on learners’ behaviors. There are 37 (pause), 27 (auditory) and 41 (visual) learners in each group. † indicates significance at $p < 0.05$ level between the first half and the second half of the *IntelliEye* sessions (Mann-Whitney U test). * indicates significance at $p < 0.05$ level between the marked group and the visual alert group (Mann-Whitney U test).

Metrics	Alert Type	First 50% <i>IntelliEye</i> Sessions	Last 50% <i>IntelliEye</i> Sessions
Mean avg. video playing length (min)	Pausing	11.93(9.46)	15.96(13.38)*
	Auditory alert	13.38(10.43)	16.16(13.17)
	Visual alert	17.15(16.21)	24.68(20.38)†
Mean avg. attention duration (min)	Pausing	6.71(7.09)	6.70(8.94)
	Auditory alert	9.38(8.76)	9.04(12.13)
	Visual alert	9.33(9.40)	12.53(17.35)
Mean avg. inattention duration (min)	Pausing	0.62(1.45)*	0.50(1.25)
	Auditory alert	0.45(1.94)*	1.07(4.93)*
	Visual alert	3.69(9.03)	3.46(5.29)
Mean avg. #inattention per min	Pausing	1.30(1.96)	1.50(2.13)
	Auditory alert	0.84(2.05)*	0.93(2.14)*
	Visual alert	2.86(4.31)	2.13(3.24)

of inattention triggers per minute and the duration of inattention. There are a number of reasons that can explain this outcome (e.g. as the material becomes more difficult over time, maintaining the same attention levels may already be a success), we will leave this investigation to future work.

6.6 Conclusions

In this chapter, we finalize our study of short-term behavioral engagement of MOOC learners. In these three chapters, short-term behavioral engagement is evaluated based on learner attention/inattention to video lectures in a short time. The main challenge of these chapters is to track learners’ inattention within video watching on a large scale and in real-time. Based on the face-tracking method proposed in Chapter 5, we design *IntelliEye* to increase learner attention while watching MOOC lecture videos by alerting learners to their loss of attention (approximated through face tracking via webcam feeds) in real-time. To re-gain learner attention, we trial three types of interventions—pausing the video with automatic resume once the learner is focusing on the video again, an auditory alert to call learners to attention, and a visual alert around the video widget.

To explore the viability and acceptance of learners towards such an assistive system, *IntelliEye* was deployed in an engineering MOOC across a 74-day period to 2,612 learners.

Our analyses explore three issues: (1) the technological capabilities of our MOOC learners' hardware, (2) the acceptance of *IntelliEye* by MOOC learners, and, (3) the effect of *IntelliEye* on MOOC learners' behaviour. We find the vast majority of learners (78%) to possess hardware capable of running *IntelliEye*; we find fewer—though still a considerable number—learners willing to try such an assistive tool (32% of all learners with supported setups) and among those that did use *IntelliEye* we determine extremely high levels of inattention, on average 1.65 inattention events per minute (i.e. on avg. inattention arises every 36 seconds).

Learners learned to adapt their behaviour as needed: learners in the pausing/auditory conditions had significantly fewer inattention events than learners in the non-disruptive visual alert condition. This though, did not yet translate into learning gains. Learners that opted not to use *IntelliEye* often did not see a need for it and were concerned about their privacy.

Considering the facts that we observe high levels of inattention and that learners once they make a decision on the tool's usage do not change that decision, we need to put more effort into the initial “sign-up” phase of such a tool in future work.

With *IntelliEye* being the first of its kind to address the learner (in)attention problem in MOOCs in real-time and by relying on non-calibrated common webcams and open-source face tracking, we have shown that there is a potential for such a system. In our future work, we will extend the deployment of *IntelliEye* to a larger audience and a wider variety of MOOCs. We will investigate learner incentives and compliance issues to increase the awareness and acceptance of our approach.

Chapter 7

Conclusion

Learner engagement is commonly believed to be essential to the success of learning in both traditional classroom contexts and MOOCs [20, 1, 33, 52, 35, 41]. In MOOCs, educators cannot observe how learners engage in their courses as they usually do in traditional classroom contexts. To better understand learner engagement in MOOC learning, learning analytics technology can be applied on large-scale trace data generated by learner interactions with course materials on MOOC platforms [102]. As mentioned by Fredrick et al. [35], learner engagement can vary in intensity and duration. Therefore, learning analytics technology applied in this field should provide a comprehensive understanding of learner engagement with diverse intensity and duration. In this thesis, we focus on using learning analytics technology to explore learner behavioral engagement—the participation of learners in MOOC learning—at different time scales. We evaluate behavioral engagement of MOOC learners on three time scales: long-term behavioral engagement (i.e. behavioral engagement throughout a course), mid-term behavioral engagement (e.g. behavioral engagement in learning sessions), and short-term behavioral engagement (e.g. behavioral engagement in a short period of time¹). In this chapter, we first summarize our main contributions in this thesis and then sketch some possible future research directions based on our studies.

¹In our studies in Chapter 4, Chapter 5, and Chapter 6, the time period we focus is no longer than 30 seconds.

7.1 Summary of Contributions

7.1.1 Long-Term Behavioral Engagement

We explored behavioral engagement of passers—learners whose scores reached the course requirements before the end of the course—throughout several MOOCs. By investigating learner interactions with video lectures and quiz questions before and after passing, our work contributes new understanding to the impact of the "passing" event on learner behavioral engagement in MOOC learning. Specifically, our study answers the following research questions:

RQ 1.1: Do MOOC learners behave differently after clinching a passing grade?

RQ 1.2: What are the core behavior patterns of MOOC learners before and after passing, and how can learners be classified?

To answer **RQ 1.1** and **RQ 1.2**, we used data-driven approaches to explore passer behaviors based on their trace data throughout several MOOCs. For **RQ 1.1**, our findings reveal that the "passing" event heavily influences ensuing learner behavior, and there are a certain amount of learners whose scores reduced heavily after passing. It means that educators may pass learners whose grasp on course contents is not complete. Moreover, if a large number of learners are only exposed to parts of course contents, educators and course designers may waste a lot of time creating content that actually few learners care about. For **RQ 1.2**, we defined pre-/post-passing behavior patterns and employed a rule-based method to classify learners based on their behavior patterns. The results of our analyses show that most passers fall into a narrow band of behavior patterns. However, the behavior patterns with the majority of passers vary in different courses. For example, in the course *Introduction to Functional Programming* the majority of learners still kept high scores after passing while in *Data Analysis: Take It to the MAX()* the majority of learners reduced their scores sharply. It indicates that the motivation of the majority of passers varies in different courses.

Considering that MOOC learners may pass a course without accessing entire course contents, our study suggests that educators and course designers need to carefully design assessment systems and organize course contents based on the requirements of the course certificate. For example, if all compulsory course contents were arranged in the first few units of the MOOC,

learners have to study all of them before passing. If there were thresholds of assessments in each unit of the MOOC, learners cannot get certificates before they show their mastery of knowledge in each unit.

Our data-driven approaches used in this study can serve as a foundation for learning analytics systems for the observation of learner engagement in MOOCs. With well-designed user interfaces for data visualization and algorithms for handling constantly updated data, educators can observe learners in a scalable way during course runnings. Moreover, they can group learners based on their behavior patterns and explore specific group of learners with different purposes.

7.1.2 Mid-Term Behavioral Engagement

As mobile learning is on the rise in MOOC learning, we explored behavioral engagement of MOOC learners in learning sessions with a mobile device. To evaluate learner behavioral engagement, we selected different metrics of MOOC learners on their learning gain, learning efficiency, and learner interactions in revising course content. Our study contributes to new findings of how learning on-the-go impacts learner behavioral engagement in MOOC learning. Particularly, our study answers the following two research questions:

RQ 2.1: To what extent does learning on-the-go (compared to stationary learning on a mobile device) affect MOOC learners' learning gain, learning efficiency and interactions with the course content?

RQ 2.2: How do learners perceive their workload (physical as well as mental) in the stationary and learning on-the-go conditions and how does it relate to their learning performance and interactions?

To answer **RQ 2.1** and **RQ 2.2**, we designed a user study which asked learners to watch MOOC video lectures and answer follow-up questions on a mobile device in both stationary learning and learning on-the-go.

For **RQ 2.1**, metrics reported in our study show that: 1) learning gain and learning efficiency were both lowered in learning on-the-go, and 2) learners in stationary learning tended to spend more time on question answering. For **RQ 2.2**, we find that learners in learning on-the-go perceived higher physical demands and more frustration than in stationary learning. The positive correlation between the question answering duration and self-rated

frustration in learning on-the-go shows that the more time learners spent on questions while learning on-the-go, the more frustrated they felt.

The experimental design in this study can be extended with different experimental factors and integrated with other types of data (e.g. sensor data) in future studies on mobile MOOC learning.

7.1.3 Short-Term Behavioral Engagement

We investigated behavioral engagement of MOOC learners in a short period of time (≤ 30 seconds in our studies), by using webcams to detect learner inattention during video watching. In our studies, we presented new approaches which could track short-term behavioral engagement (e.g. learner attention in our studies) of MOOC learners in near real-time and on a large scale. More specifically, our studies mainly focus on the following research questions:

RQ 3.1: How often do MOOC learners experience inattention within video watching?

RQ 3.2: How well do our webcam-based inattention detection methods perform?

RQ 3.3: To what extent is MOOC learners' hardware capable to enable the webcam-based inattention detection?

RQ 3.4: To what extent do MOOC learners accept our inattention detection technology that is designed to aid their learning but at the same time is likely to be perceived as privacy-invading (even though it is not)?

RQ 3.5: What impact does the webcam-based inattention detection have on learners' behaviors and to what extent does it affect learners' video watching behaviors?

To answer the above five questions, we conducted a series of studies on webcam-based attention tracking in video watching.

To answer **RQ 3.1** and **RQ 3.2**, we first conducted a user study in which learners were asked to watch a MOOC video lecture and report their inattention after hearing auditory probes. During video watching, both a webcam-based eye-tracking approach and a professional eye-tracker were running to collect learner gaze data. The user study is designed based on the eye-mind link effect [96] that *there is no appreciable lag between what is fixated and*

what is processed. Based on our analyses on data collected from the user study, we find that inattention occurred frequently even for short video lectures. For eye-tracking based inattention detection, our results show that the accuracy of our webcam-based approach was on par with the approach with a professional eye-tracker. Our study indicates that it is indeed possible to use the webcam-based inattention detection in MOOCs on a large scale.

However, the above eye-tracking approach suffers a series of problems: the relatively low accuracy, the long detection lag, and the complexity of the detection process. These problems make the eye-tracking approach difficult to use in the near real-time inattention detection. To better answer **RQ 3.2**, we proposed an alternative approach with face tracking which could increase the detection accuracy and reduce the detection lag. To test the approach with face tracking, we designed a benchmark which contains a set of learner activities related to the attention/inattention of MOOC learners. Our results of experiments on the benchmark show that the face-tracking approach makes an improvement on both detection accuracy and speed, which makes it feasible to track learner inattention in near real-time with webcams.

To answer **RQ 3.3**, **RQ 3.4**, and **RQ 3.5**, we designed *IntelliEye*, a near real-time, scalable, privacy-aware widget which could track learner inattention within video watching only based on a webcam. We deployed *IntelliEye* in a real-world MOOC. The feedback of MOOC learners shows that even though the hardware and software of most learners were capable to run *IntelliEye*, the majority of learners were reluctant to allow the use of *IntelliEye* because of the lack of perceived usefulness and privacy concerns. Our analysis of learner interactions with *IntelliEye* reveals that learners had high levels of inattention and they did adapt their behaviors to our widget.

7.2 Future Work

In this thesis, we contribute new knowledge and novel technical approaches to understanding learner behavioral engagement in MOOCs. In this section, we outline potential future research directions—in which technical approaches presented in our studies can be applied—in the broad area of MOOC learning.

7.2.1 Personal Analytics in MOOCs

Personal analytics, which can provide personalized feedback or guidance to learners based on learning analytics, is an important research direction

in MOOCs which can improve learner awareness in self-regulated learning. Learner awareness is critical for the use of learning strategies in self-regulated learning [150]. For example, learners can select effective learning strategies if they are able to monitor the impact of these strategies on their performance. However, we notice that learners on premier MOOC platforms can only access a small amount of information (e.g. scores received in each part of course content, and time required on different materials) about their learning process.

In the 8th International Learning Analytics and Knowledge Conference (LAK2018), personal analytics supporting self-directed learning is emphasized by organizers of the Hackathon workshop (Hack@LAK18) for user-centred learning analytics. The main purpose of this topic in Hack@LAK18 is to do *learning analytics for the learner* based on data collected from learner's browser and social media account. LAK2018 also held the first international workshop on personalizing feedback which aims to explore future directions to improve the process and richness of personal feedback based on the application of learning analytics.

We envision that an ideal personal analytics system in MOOCs should help learners to not only clearly monitor their learning progress but also fully understand their mastery of knowledge. Moreover, personalized guidance can be generated by this ideal personal analytics system to motivate learners and help them make better learning strategies. Some technologies in previous studies can be applied to building this personal analytics system. For example, Davis et al. [27] build a visual "personalized feedback system" which provides weekly feedback to learners based on the comparison of the learning progress of learners with that of successful passers from a previous run of the same MOOC. Our studies in this thesis can be applied to monitor learner engagement on different time scales in personal analytics. For example, our study in Chapter 2 can be used to analyze the engagement of each learner throughout a course while our work in Chapter 6 can be applied to monitor the engagement of each learner in near real-time. Research on open learner models [11] can be used in MOOCs for the measurement of learner mastery of knowledge. However, to implement the personal analytics system, it still lacks studies on the integration of those approaches and models (e.g. personal analytics dashboard, comprehensive measurements on learner performance and learning progress, or open learner models on learner mastery of knowledge) and the generation of personal guidance in MOOCs.

Further research on personal analytics in MOOCs can be expanded for the following questions. Among different measurements of the learning progress

and the mastery of knowledge, which of them can be perceived by MOOC learners efficiently and effectively? How to generate reasonable personal guidance automatically and in near real-time based on learning analytics on a large scale of MOOC learners? How does such a personal analytics system change learner behaviors and learning outcomes in MOOCs?

7.2.2 Adaptive Learning in MOOCs

To enhance the learning gain and the learning efficiency of learners, adaptive learning is a promising future research direction. Adaptive learning systems refers to systems that *attempt to be different for different students and groups of students by taking into account information accumulated in the individual or group student models* [13]. In LAK2018, conference organizers called for papers with topics about personalized and adaptive learning. Specifically, they were interested in studies on the evaluation of the effectiveness and impact of adaptive technologies. In the UNESCO report about smart learning environments for the 21th century [120], they envision a smart learning environment as an adaptive learning system in which learning can occur anywhere, anytime and at any pace. In the 26th ACM Conference on User Modeling, Adaptation, and Personalization (UMAP2018), technology-enhanced adaptive learning was one of the main track topics, which focused on technological solutions for modeling learner and providing personalized adapted support.

MOOCs are open to diverse learners on a massive scale and these learners have different learning goals/motivations, prior knowledge, or the other backgrounds. The engagement of learners is also diverse in MOOC learning (as shown in our studies in this thesis). With adaptive learning systems, course materials in MOOCs can be customized automatically based on the needs of individual learners. However, adaptive learning has not been fully studied in MOOCs. The possible reason is that there is no integrated adaptive learning system on premier MOOC platforms. Educators and course designers who want to design adaptive MOOCs have to have both an in-depth understanding of the course content and sufficient programming skills to implement a stand-alone adaptive MOOC platform or a plug-in that can be integrated into the current MOOC platforms [86, 103]. Current adaptive MOOC platforms in previous studies [125, 82, 113] require a significant workload from educators and course designers to design course contents which can be integrated into different learning paths for learners with different profiles, while current adaptive plug-ins are in the preliminary stages which only focus on a small

part of the adaptation learning (e.g. the navigation recommendation [86] and the adaptive assessment [103]) in MOOCs.

An adaptive learning system usually consists of two modules: one module for assessing the current profile of each learner and another module which decides what see next [103]. In future studies on assessing learner profile, our studies in Chapter 6 can be applied to evaluate learner engagement as a part of learner profile in near real-time. Further studies also can be expanded on learning paths which decide what learners to see next. For example, how to generate learning paths without much extra workload and adjust them automatically for each individual MOOC learner based on her motivation, learning pace, profile, or preference? How to accurately evaluate the quality of learning paths generated by adaptive learning systems in MOOCs?

7.2.3 Multimodal Learning Analytics in MOOCs

Multimodal learning analytics is an interesting future direction for better understanding the learning progress of MOOC learners. Since learner trace data on MOOC platforms cannot fully describe the learning progress of MOOC learners.

Based on the definition in [79], *multimodal learning analytics works to leverage advances in multimodal data capture and signal processing to address the challenges of studying a variety of complex learning-relevant constructs as observed in complex learning environments*. In the 11th Annual International Educational Data Mining Conference (EDM2018), multimodal learning analytics was one of the main session topics. In this session, studies on motion sensors and eye tracking were discussed. In LAK2018, researchers also held a workshop named as Multimodal Learning Analytics Across (Physical and Digital) Spaces (CrossMMLA). CrossMMLA mainly focused on learner multimodal interactions in real-world learning contexts.

Multimodal learning analytics in MOOCs has not been fully researched. Previous research on multimodal learning analytics in MOOC learning mainly focus on mobile MOOC learning [91] in which learner heart rate data and facial expression data are collected and analyzed. A possible reason is that mobile devices have been ubiquitous in MOOC learning and have different types of sensors to collect multimodal data from both learners and learning environments.

Our study in Chapter 6 provides a direction to use webcams to obtain multimodal data of MOOC learners on a large scale. Once cheap devices with

different interaction technologies (e.g. intelligent personal assistants, activity tracking, eye tracking, augmented reality, virtual reality, EEG or fNIRS) are on the market, these technologies can be applied on a large scale in future MOOCs. Consequently, learners can be modeled by using multimodal learning analytics on not only trace data and survey data but also physiological data, gaze data and environmental data. For example, if a large number of learners use eye-tracking devices and fNIRS devices during watching MOOC videos, educators and researchers can track learners' gaze movements and brain activity on a large scale when different content delivered in videos. Then, they may be able to distinguish between what most learners understand and what they ignore during video watching by multimodal learning analysis. After that, questions following videos in MOOCs can be designed to specifically evaluate content that is important but easily ignored by most learners when watching videos.

Both personal analytics and adaptive learning in MOOCs could also be improved by multimodal learning analytics. For example, with personal analytics based on learner trace data and physiological data, how to detect the obstacles faced by learners and generate personal feedback and guidance? In adaptive learning, how to model environments where learning occurs based on multimodal learning analytics and to what extent course content can be tailored automatically to different learning environments?

Bibliography

- [1] Karl L Alexander, Doris R Entwisle, and Carrie S Horsey. From first grade forward: Early foundations of high school dropout. *Sociology of education*, pages 87–107, 1997.
- [2] Ashton Anderson, Daniel Huttenlocher, Jon Kleinberg, and Jure Leskovec. Engaging with massive online courses. In *Proceedings of the 23rd international conference on World wide web*, pages 687–698. ACM, 2014.
- [3] Monica Anderson. *Technology Device Ownership, 2015*. Pew Research Center, 2015.
- [4] Pavlo D Antonenko, Serkan Toy, and Dale S Niederhauser. Using cluster analysis for data mining in educational technology research. *Educational Technology Research and Development*, 60(3):383–398, 2012.
- [5] Stylianos Asteriadis, Kostas Karpouzis, and Stefanos Kollias. Visual focus of attention in non-calibrated environments using gaze estimation. *International Journal of Computer Vision*, 107(3):293–316, 2014.
- [6] Robert E Beasley, Jacob T McMain, Mathew D Millard, Dylan A Pasley, and Matthew J Western. The effects of college student smartphone use on academic distraction and dishonesty. *Journal of Computing Sciences in Colleges*, 32(1):17–26, 2016.
- [7] Dominic Becking, Stefan Betermieux, Birgit Bomsdorf, Feldmann Birgit, Eberhard Heuel, Patrick Langer, and Gunter Schlageter. Didactic profiling: supporting the mobile learner. In *E-Learn: World Conference on E-Learning in Corporate, Government, Healthcare, and Higher Education*, pages 1760–1767. Association for the Advancement of Computing in Education (AACE), 2004.

- [8] Robert Bixler and Sidney D' Mello. Toward fully automated person-independent detection of mind wandering. In *International Conference on User Modeling, Adaptation, and Personalization*, pages 37–48. Springer, 2014.
- [9] Robert Bixler and Sidney D' Mello. Automatic gaze-based user-independent detection of mind wandering during computerized reading. *User Modeling and User-Adapted Interaction*, 26(1):33–68, 2016.
- [10] Nathaniel Blanchard, Robert Bixler, Tera Joyce, and Sidney D' Mello. Automated physiological-based detection of mind wandering during learning. In *International Conference on Intelligent Tutoring Systems*, pages 55–60. Springer, 2014.
- [11] Robert Bodily, Judy Kay, Vincent Aleven, Ioana Jivet, Dan Davis, Franceska Xhakaj, and Katrien Verbert. Open learner models and learning analytics dashboards: a systematic review. In *Proceedings of the 8th International Conference on Learning Analytics and Knowledge*, pages 41–50. ACM, 2018.
- [12] Lori Breslow, David E Pritchard, Jennifer DeBoer, Glenda S Stump, Andrew D Ho, and Daniel T Seaton. Studying learning in the worldwide classroom research into edx's first mooc. *Research & Practice in Assessment*, 8:13–25, 2013.
- [13] Peter Brusilovsky and Christoph Peylo. Adaptive and intelligent web-based educational systems. *International Journal of Artificial Intelligence in Education (IJAIED)*, 13:159–172, 2003.
- [14] Diane M Bunce, Elizabeth A Flens, and Kelly Y Neiles. How long can students pay attention in class? a study of student attention decline using clickers. *Journal of Chemical Education*, 87(12):1438–1443, 2010.
- [15] Paul Cameron and Dorothy Giuntoli. Consciousness sampling in the college classroom or is anybody listening?. *Intellect*, 101(2343):63–4, 1972.
- [16] Soledad Castellano and Inmaculada Arnedillo-Sánchez. Sensorimotor distractions when learning with mobile phones on-the-move. *International Association for Development of the Information Society*, 2016.
- [17] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.

- [18] Isaac Chuang and Andrew Ho. Harvardx and mitx: Four years of open online courses—fall 2012–summer 2016. 2016.
- [19] Carleton Coffrin, Linda Corrin, Paula de Barba, and Gregor Kennedy. Visualizing patterns of student engagement and performance in moocs. In *Proceedings of the fourth international conference on learning analytics and knowledge*, pages 83–92. ACM, 2014.
- [20] James Patrick Connell, Margaret Beale Spencer, and J Lawrence Aber. Educational risk and resilience in african-american youth: Context, self, action, and outcomes in school. *Child development*, 65(2):493–506, 1994.
- [21] Averil Cook. Assessing the use of flexible assessment. *Assessment & Evaluation in Higher Education*, 26(6):539–549, 2001.
- [22] Matthieu Perreira Da Silva, Vincent Courboulay, Armelle Prigent, and Pascal Estraillier. Real-time face tracking for attention aware adaptive games. In *International Conference on Computer Vision Systems*, pages 99–108. Springer, 2008.
- [23] Fisnik Dalipi, Ali Shariq Imran, Florim Idrizi, and Hesat Aliu. An analysis of learner experience with moocs in mobile and desktop learning environment. In *Advances in Human Factors, Business Management, Training and Education*, pages 393–402. Springer, 2017.
- [24] Lia M Daniels, Catherine Adams, and Adam McCaffrey. Emotional and social engagement in a massive open online course: An examination of dino 101. In *Emotions, technology, and learning*, pages 25–41. Elsevier, 2016.
- [25] Dan Davis, Guanliang Chen, Claudia Hauff, and Geert-Jan Houben. Gauging mooc learners’ adherence to the designed learning path. In *Proceedings of the 9th International Conference on Educational Data Mining*, pages 54–61, 2016.
- [26] Dan Davis, Guanliang Chen, Tim Van der Zee, Claudia Hauff, and Geert-Jan Houben. Retrieval practice and study planning in moocs: Exploring classroom-based self-regulated learning strategies at scale. In *European Conference on Technology Enhanced Learning*, pages 57–71. Springer, 2016.
- [27] Dan Davis, Ioana Jivet, René F Kizilcec, Guanliang Chen, Claudia Hauff, and Geert-Jan Houben. Follow the successful crowd: raising

- moooc completion rates through social comparison at scale. In *Proceedings of the Seventh International Learning Analytics & Knowledge Conference*, pages 454–463. ACM, 2017.
- [28] Sara Isabella de Freitas, John Morgan, and David Gibson. Will moocs transform learning and teaching in higher education? engagement and course retention in online learning provision. *British Journal of Educational Technology*, 46(3):455–471, 2015.
- [29] Tilman Dingler, Dominik Weber, Martin Pielot, Jennifer Cooper, Chung-Cheng Chang, and Niels Henze. Language learning on-the-go: opportune moments and design of mobile microlearning sessions. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services*, page 28. ACM, 2017.
- [30] Ralf Engbert and Reinhold Kliegl. Microsaccades uncover the orientation of covert attention. *Vision research*, 43(9):1035–1045, 2003.
- [31] George D Farmer, Christian P Janssen, Anh T Nguyen, and Duncan P Brumby. Dividing attention between tasks: Testing whether explicit payoff functions elicit optimal dual-task performance. *Cognitive science*, 2017.
- [32] Ángel Fidalgo-Blanco, María Luisa Sein-Echaluce, and Francisco José García-Peñalvo. From massive access to cooperation: lessons learned and proven results of a hybrid xmooc/cmoooc pedagogical approach to moocs. *International Journal of Educational Technology in Higher Education*, 13(1):24, 2016.
- [33] Jeremy D Finn and Donald A Rock. Academic success among students at risk for school failure. *Journal of applied psychology*, 82(2):221, 1997.
- [34] D Forbes and J Spence. An experiment in assessment for a large class. *Innovations in engineering education*. London: Ellis Horwood, pages 97–101, 1991.
- [35] Jennifer A Fredricks, Phyllis C Blumenfeld, and Alison H Paris. School engagement: Potential of the concept, state of the evidence. *Review of educational research*, 74(1):59–109, 2004.
- [36] Adam Gazzaley and Larry D Rosen. *The distracted mind: ancient brains in a high-tech world*. MIT Press, 2016.
- [37] Graham Gibbs. How assessment frames student learning. *Innovative assessment in higher education*, pages 23–36, 2006.

- [38] Graham Gibbs and Claire Simpson. Does your assessment support your students' learning. *Journal of Teaching and Learning in Higher Education*, 1(1):1–30, 2004.
- [39] Graham Gibbs and Claire Simpson. Conditions under which assessment supports students' learning. *Learning and teaching in higher education*, 1:3–31, 2005.
- [40] Philip J Guo and Katharina Reinecke. Demographic differences in how students navigate through moocs. In *Proceedings of the first ACM conference on Learning@ scale conference*, pages 21–30. ACM, 2014.
- [41] Philip J Guo, Juho Kim, and Rob Rubin. How video production affects student engagement: an empirical study of mooc videos. In *Proceedings of the first ACM conference on Learning@ scale conference*, pages 41–50. ACM, 2014.
- [42] Christian Gütl, Rocael Hernández Rizzardini, Vanessa Chang, and Miguel Morales. Attrition in mooc: Lessons learned from drop-out students. In *International Workshop on Learning Technology for Education in Cloud*, pages 37–48. Springer, 2014.
- [43] Sandra G Hart and Lowell E Staveland. Development of nasa-tlx (task load index): Results of empirical and theoretical research. In *Advances in psychology*, volume 52, pages 139–183. Elsevier, 1988.
- [44] James Hartley and Alan Cameron. Some observations on the efficiency of lecturing. *Educational Review*, 20(1):30–37, 1967.
- [45] Morgan Harvey and Matthew Pointon. Searching on the go: the effects of fragmented attention on mobile web search tasks. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 155–164. ACM, 2017.
- [46] Barbara Hauer. Continuous supervision: A novel concept for enhancing data leakage prevention. In *European Conference on Cyber Warfare and Security*, page 342. Academic Conferences International Limited, 2016.
- [47] Sue Helme and David Clarke. Identifying cognitive engagement in the mathematics classroom. *Mathematics Education Research Journal*, 13(2):133–153, 2001.
- [48] Khe Foon Hew. Promoting engagement in online courses: What strategies can we learn from three highly rated moocs. *British Journal of Educational Technology*, 47(2):320–341, 2016.

- [49] Andrew Ho, Justin Reich, Sergiy Nesterko, Daniel Seaton, Tommy Mullaney, Jim Waldo, and Isaac Chuang. Harvardx and mitx: The first year of open online courses, fall 2012-summer 2013. 2014.
- [50] Jui-Long Hung and Ke Zhang. Revealing online learning behaviors and activity patterns and making predictions with data mining techniques in online teaching. *MERLOT Journal of Online Learning and Teaching*, 2008.
- [51] K Innis. Diary survey: how undergraduate full-time students spend their time. *Leeds Metropolitan Univ., Leeds, UK*, 1996.
- [52] Shane R Jimerson, Emily Campos, and Jennifer L Greif. Toward an understanding of definitions and measures of school engagement and related terms. *The California School Psychologist*, 8(1):7–27, 2003.
- [53] Alex H Johnstone and Frederick Percival. Attention breaks in lectures. *Education in chemistry*, 13(2):49–50, 1976.
- [54] Katy Jordan. Initial trends in enrolment and completion of massive open online courses. *The International Review of Research in Open and Distributed Learning*, 15(1), 2014.
- [55] Matthew A Killingsworth and Daniel T Gilbert. A wandering mind is an unhappy mind. *Science*, 330(6006):932–932, 2010.
- [56] René F Kizilcec, Chris Piech, and Emily Schneider. Deconstructing disengagement: analyzing learner subpopulations in massive open online courses. In *Proceedings of the third international conference on learning analytics and knowledge*, pages 170–179. ACM, 2013.
- [57] René F Kizilcec, Kathryn Papadopoulos, and Lalida Sritanyaratana. Showing face in video instruction: effects on information retention, visual attention, and affect. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 2095–2102. ACM, 2014.
- [58] René F Kizilcec, Mar Pérez-Sanagustín, and Jorge J Maldonado. Recommending self-regulated learning strategies does not improve performance in a mooc. In *Proceedings of the Third (2016) ACM Conference on Learning@ Scale*, pages 101–104. ACM, 2016.
- [59] Geza Kovacs. Effects of in-video quizzes on mooc lecture viewing. In *Proceedings of the Third (2016) ACM Conference on Learning@ Scale*, pages 31–40. ACM, 2016.

- [60] T Krasovsky, PL Weiss, and R Kizony. Effect of aging, mixed reality and dual task prioritization on texting while walking. In *2017 International Conference on Virtual Rehabilitation (ICVR)*, pages 1–6. IEEE, 2017.
- [61] Sammy Licence, Robynne Smith, Miranda P. McGuigan, and Conrad P. Earnest. Gait pattern alterations during walking, texting and walking and texting during cognitively distractive tasks while negotiating common pedestrian obstacles. *PLOS ONE*, 10(7):1–11, 07 2015.
- [62] Sophie I Lindquist and John P McLean. Daydreaming and its correlates in an educational environment. *Learning and Individual Differences*, 21(2):158–167, 2011.
- [63] Kep Kee Loh, Benjamin Zhi Hui Tan, and Stephen Wee Hun Lim. Media multitasking predicts video-recorded lecture learning performance through mind wandering tendencies. *Computers in Human Behavior*, 63:943–947, 2016.
- [64] Eduardo Lundgren, Thiago Rocha, Zeno Rocha, Pablo Carvalho, and Maira Bello. tracking.js: A modern approach for computer vision on the web. *Online*. *Dosegljivo: <https://trackingjs.com/>*[Dostopano 30. 5. 2016], 2015.
- [65] H Maddox and Elizabeth Hoole. Performance decrement in the lecture. *Educational Review*, 28(1):17–30, 1975.
- [66] Félix Albertos Marco, Víctor MR Penichet, and José A Gallud. What happens when students go offline in mobile devices? In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct*, pages 1199–1206. ACM, 2015.
- [67] Ference Marton and R Säljö. On qualitative differences in learning – ii outcome as a function of the learner’s conception of the task. *British journal of educational psychology*, 46(2):115–127, 1976.
- [68] John McLeish. *The lecture method*. Cambridge Institute of Education, 1968.
- [69] Rebecca McMillan, Scott Barry Kaufman, and Jerome L Singer. Ode to positive constructive daydreaming. *Frontiers in psychology*, 4:626, 2013.

- [70] Judith L Meece, Phyllis C Blumenfeld, and Rick H Hoyle. Students' goal orientations and cognitive engagement in classroom activities. *Journal of educational psychology*, 80(4):514, 1988.
- [71] Carolyn ML Miller and Malcolm Parlett. Up to the mark: A study of the examination game. 1974.
- [72] George A Miller. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological review*, 6a3(2):81, 1956.
- [73] Caitlin Mills, Robert Bixler, Xinyi Wang, and Sidney K D' Mello. Automatic gaze-based detection of mind wandering during narrative film comprehension. *EDM*, 16:30–37, 2016.
- [74] Anka Mulder. Moocs for credit .. goes live. *TU Delft Anka Mulder Weblog*, 2016.
- [75] Almedina Music and Stéphan Vincent-Lancrin. Massive open online courses (moocs): Trends and future perspectives. *EDU/CERI/CD/RD*, 2016(5), 2016.
- [76] Josip Music, Ivo Stancic, and Vlasta Zanchi. Is it possible to detect mobile phone user's attention based on accelerometer measurement of gait pattern? In *2013 IEEE symposium on computers and communications (ISCC)*, pages 522–527. IEEE, 2013.
- [77] Gary Natriello and Sanford M Dornbusch. *Teacher evaluative standards and student effort*. Longman Publishing Group, 1984.
- [78] David I Newble and Kerry Jaeger. The effect of assessments and examinations on the learning of medical students. *Medical education*, 17(3):165–171, 1983.
- [79] Xavier Ochoa and Marcelo Worsley. Augmenting learning analytics with multimodal sensory data. *Journal of Learning Analytics*, 3(2): 213–219, 2016.
- [80] Masaki Ogawa, Takuro Yonezawa, Jin Nakazawa, and Hideyuki Tokuda. Exploring user model of the city by using interactive public display application. In *Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers*, pages 1595–1598. ACM, 2015.

- [81] Claire O'Malley, Giasemi Vavoula, JP Glew, Josie Taylor, Mike Sharples, Paul Lefrere, Peter Lonsdale, Laura Naismith, and Jenny Waycott. Guidelines for learning/teaching/tutoring in a mobile environment. 2005.
- [82] Daniel FO Onah and JE Sinclair. Massive open online courses-an adaptive learning framework. In *9th International Technology, Education and Development Conference, IATED, Madrid, Spain*, pages 2–4, 2015.
- [83] Alexandra Papoutsaki, Patsorn Sangkloy, James Laskey, Nediya Daskalova, Jeff Huang, and James Hays. Webgazer: Scalable webcam eye tracking using user interactions. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence-IJCAI 2016*, 2016.
- [84] Alexandra Papoutsaki, James Laskey, and Jeff Huang. Searchgazer: Webcam eye tracking for remote studies of web search. In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval*, pages 17–26. ACM, 2017.
- [85] Laura Pappano. The year of the mooc. *The New York Times*, 2(12): 2012, 2012.
- [86] Zachary A Pardos, Steven Tang, Daniel Davis, and Christopher Vu Le. Enabling real-time adaptivity in moocs with a personalized next-step recommendation framework. In *Proceedings of the Fourth (2017) ACM Conference on Learning@ Scale*, pages 23–32. ACM, 2017.
- [87] Mariana Patru and Venkataraman Balaji. Making sense of moocs: A guide for policy-makers in developing countries, 2016.
- [88] Lloyd Peterson and Margaret Jean Peterson. Short-term retention of individual verbal items. *Journal of experimental psychology*, 58(3):193, 1959.
- [89] Phuong Pham and Jingtao Wang. Attentivelearner: improving mobile mooc learning via implicit heart rate tracking. In *International Conference on Artificial Intelligence in Education*, pages 367–376. Springer, 2015.
- [90] Phuong Pham and Jingtao Wang. Adaptive review for mobile mooc learning via implicit physiological signal sensing. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, pages 37–44. ACM, 2016.

- [91] Phuong Pham and Jingtao Wang. Predicting learners' emotions in mobile mooc learning via a multimodal intelligent tutor. In *International Conference on Intelligent Tutoring Systems*, pages 150–159. Springer, 2018.
- [92] Paul R Pintrich and Elisabeth V De Groot. Motivational and self-regulated learning components of classroom academic performance. *Journal of educational psychology*, 82(1):33, 1990.
- [93] Krenare Pireva, Ali Shariq Imran, and Fisnik Dalipi. User behaviour analysis on lms and mooc. In *2015 IEEE Conference on e-Learning, e-Management and e-Services (IC3e)*, pages 21–26. IEEE, 2015.
- [94] Lalita Rajasingham. Will mobile learning bring a paradigm shift in higher education? *Education Research International*, 2011.
- [95] Arti Ramesh, Dan Goldwasser, Bert Huang, Hal Daume III, and Lise Getoor. Uncovering hidden engagement patterns for predicting learner performance in moocs. In *Proceedings of the first ACM conference on Learning@ scale conference*, pages 157–158. ACM, 2014.
- [96] Keith Rayner. Eye movements in reading and information processing: 20 years of research. *Psychological bulletin*, 124(3):372–422, 1998.
- [97] Jan Renz, Thomas Staubitz, and Christoph Meinel. Mooc to go. *International Association for Development of the Information Society*, 2014.
- [98] Evan F Risko, Nicola Anderson, Amara Sarwal, Megan Engelhardt, and Alan Kingstone. Everyday attention: variation in mind wandering and memory in a lecture. *Applied Cognitive Psychology*, 26(2):234–242, 2012.
- [99] Tarmo Robal, Yue Zhao, Christoph Lofi, and Claudia Hauff. Intellieye: Enhancing mooc learners' video watching experience through real-time attention tracking. In *Proceedings of the 29th on Hypertext and Social Media*, pages 106–114. ACM, 2018.
- [100] Tarmo Robal, Yue Zhao, Christoph Lofi, and Claudia Hauff. Webcam-based attention tracking in online learning: A feasibility study. In *23rd International Conference on Intelligent User Interfaces*, pages 189–197. ACM, 2018.
- [101] Osvaldo Rodriguez. The concept of openness behind c and x-moocs (massive open online courses). *Open Praxis*, 5(1):67–73, 2013.

- [102] Ido Roll, Daniel M Russell, and Dragan Gašević. Learning at scale. *International Journal of Artificial Intelligence in Education*, pages 1–7, 2018.
- [103] Yigal Rosen, Ilia Rushkin, Rob Rubin, Liberty Munson, Andrew Ang, Gregory Weber, Glenn Lopez, and Dustin Tingley. The effects of adaptive learning in a massive open online course on learners’ skill development. In *Proceedings of the Fifth Annual ACM Conference on Learning at Scale*, page 6. ACM, 2018.
- [104] Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.
- [105] Chris Rust. The impact of assessment on student learning how can the research literature practically help to inform the development of departmental assessment strategies and learner-centred assessment practices? *Active learning in higher education*, 3(2):145–158, 2002.
- [106] Allison M Ryan and Helen Patrick. The classroom social environment and changes in adolescents’ motivation and engagement during middle school. *American Educational Research Journal*, 38(2):437–460, 2001.
- [107] Dario D Salvucci and Joseph H Goldberg. Identifying fixations and saccades in eye-tracking protocols. In *Proceedings of the 2000 Symposium on Eye Tracking Research & Applications*, pages 71–78. ACM, 2000.
- [108] Kay Sambell and Liz McDowell. The construction of the hidden curriculum: messages and meanings in the assessment of student learning. *Assessment & Evaluation in Higher Education*, 23(4):391–402, 1998.
- [109] Brahima Sanou. Ict facts and figures 2017. *International Telecommunication Union (ITU) Fact Sheet*, 2017.
- [110] Filipe Santos, Ana Almeida, Constantino Martins, Paulo Moura de Oliveira, and Ramiro Gonçalves. Hybrid tourism recommendation system based on functionality/accessibility levels. In *International Conference on Practical Applications of Agents and Multi-Agent Systems*, pages 221–228. Springer, 2017.
- [111] Karen Scouller. The influence of assessment method on students’ learning approaches: Multiple choice question examination versus assignment essay. *Higher Education*, 35(4):453–472, 1998.

- [112] Daniel T Seaton, Yoav Bergner, Isaac Chuang, Piotr Mitros, and David E Pritchard. Who does what in a massive open online course? *Communications of the ACM*, 57(4):58–65, 2014.
- [113] María Luisa Sein-Echaluce, Ángel Fidalgo-Blanco, Francisco J García-Peñalvo, and Miguel Ángel Conde. imooc platform: Adaptive moocs. In *International Conference on Learning and Collaboration Technologies*, pages 380–390. Springer, 2016.
- [114] Dhawal Shah. By the numbers: Moocs in 2017, 2018.
- [115] Kshitij Sharma, Patrick Jermann, and Pierre Dillenbourg. How students learn using moocs: An eye-tracking insight. Technical report, 2014.
- [116] Kshitij Sharma, Patrick Jermann, and Pierre Dillenbourg. Displaying teacher’s gaze in a mooc: Effects on students’ video navigation patterns. In *Design for Teaching and Learning in a Networked World*, pages 325–338. Springer, 2015.
- [117] Mike Sharples, Inmaculada Arnedillo-Sánchez, Marcelo Milrad, and Giasemi Vavoula. Mobile learning. In *Technology-enhanced learning*, pages 233–249. Springer, 2009.
- [118] Mike Sharples, Josie Taylor, and Giasemi Vavoula. A theory of learning for the mobile age. In *Medienbildung in neuen Kulturräumen*, pages 87–99. Springer, 2010.
- [119] Mike Sharples, Carlos Delgado Kloos, Yannis Dimitriadis, Serge Garlatti, and Marcus Specht. Mobile and accessible learning for moocs. *Journal of interactive media in education*, 2015(1), 2015.
- [120] Abtar Darsham Singh and Moustafa Hassan. In pursuit of smart learning environments for the 21st century. *Current and Critical Issues in Curriculum, Learning and Assessment*, 12, 2017.
- [121] Ellen A Skinner and Michael J Belmont. Motivation in the classroom: Reciprocal effects of teacher behavior and student engagement across the school year. *Journal of educational psychology*, 85(4):571, 1993.
- [122] Jonathan Smallwood and Jonathan W Schooler. The restless mind. *Psychological bulletin*, 132(6):946–958, 2006.

- [123] Jonathan Smallwood, Daniel J Fishman, and Jonathan W Schooler. Counting the cost of an absent mind: Mind wandering as an underrecognized influence on educational performance. *Psychonomic bulletin & review*, 14(2):230–236, 2007.
- [124] Benson R. Snyder. *The Hidden Curriculum*. Cambridge, Mass., MIT Press, 1970.
- [125] Nishikant Sonwalkar. The first adaptive mooc: A case study on pedagogy framework and scalable cloud architecture—part i. In *MOOCs Forum*, volume 1, pages 22–29. Mary Ann Liebert, Inc. 140 Huguenot Street, 3rd Floor New Rochelle, NY 10801 USA, 2013.
- [126] John Stuart and RJD Rutherford. Medical student concentration during lectures. *The lancet*, 312(8088):514–516, 1978.
- [127] Rohail Syed and Kevyn Collins-Thompson. Retrieval algorithms optimized for human learning. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 555–564. ACM, 2017.
- [128] Karl K Szpunar, Novall Y Khan, and Daniel L Schacter. Interpolated memory tests reduce mind wandering and improve learning of online lectures. *Proceedings of the National Academy of Sciences*, 110(16):6313–6317, 2013.
- [129] Bernardo Tabuenca, Stefaan Ternier, and Marcus Specht. Everyday patterns in lifelong learners to build personal learning ecologies. In *Proceedings of the 11th World Conference on Mobile and Contextual Learning*, pages 86–93, 2012.
- [130] Sajjad Taheri, Alexander Veidenbaum, Alexandru Nicolau, and Mohammad R Haghighat. Opencv.js: Computer vision processing for the web. Technical report, Center for Embedded and Cyber-Physical Systems, University of California, Irvine, 2017.
- [131] Luis Talavera and Elena Gaudioso. Mining student data to characterize similar behavior groups in unstructured collaboration spaces. In *Workshop on artificial intelligence in CSCL. 16th European conference on artificial intelligence*, pages 17–23, 2004.
- [132] Tengpeng Tan, Hong Zou, Chuansheng Chen, and Jin Luo. Mind wandering and the incubation effect in insight problem solving. *Creativity Research Journal*, 27(4):375–382, 2015.

- [133] UNESCO. Six ways to ensure higher education leaves no one behind. *ED/GEMR/MRT*, (30), 2017.
- [134] Christopher O Walker, Barbara A Greene, and Robert A Mansell. Identification with academics, intrinsic/extrinsic motivation, and self-efficacy as predictors of cognitive engagement. *Learning and individual differences*, 16(1):1–12, 2006.
- [135] Xiaoyue Wang, Abdullah Mueen, Hui Ding, Goce Trajcevski, Peter Scheuermann, and Eamonn Keogh. Experimental comparison of representation methods and distance measures for time series data. *Data Mining and Knowledge Discovery*, 26(2):275–309, 2013.
- [136] Xu Wang, Diyi Yang, Miaomiao Wen, Kenneth Koedinger, and Carolyn P Rosé. Investigating how student’s cognitive behavior in mooc discussion forums affect learning gains. *International Educational Data Mining Society*, 2015.
- [137] Miaomiao Wen and Carolyn Penstein Rosé. Identifying latent study habits by mining learner behavior patterns in massive open online courses. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, pages 1983–1986. ACM, 2014.
- [138] Miaomiao Wen, Diyi Yang, and Carolyn Rose. Sentiment analysis in mooc discussion forums: What does it tell us? In *Educational data mining 2014*. Citeseer, 2014.
- [139] Miaomiao Wen, Diyi Yang, and Carolyn Penstein Rosé. Linguistic reflections of student engagement in massive open online courses. In *Eighth International AAAI Conference on Weblogs and Social Media*, 2014.
- [140] Karen Wilson and James H Korn. Attention during lectures: Beyond ten minutes. *Teaching of Psychology*, 34(2):85–89, 2007.
- [141] Beverly Woolf, Winslow Bursleson, Ivon Arroyo, Toby Dragon, David Cooper, and Rosalind Picard. Affect-aware tutors: recognising and responding to student affect. *International Journal of Learning Technology*, 4(3-4):129–164, 2009.
- [142] Xiang Xiao and Jingtao Wang. Understanding and detecting divided attention in mobile mooc learning. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 2411–2415. ACM, 2017.

- [143] Jingteng Xue and Chang Wen Chen. Mobile video perception: New insights and adaptation strategies. *IEEE journal of selected topics in signal processing*, 8(3):390–401, 2014.
- [144] Thorsten O Zander, Christian Kothe, Sabine Jatzhev, and Matti Gaertner. Enhancing human-computer interaction with input from active and passive brain-computer interfaces. In *Brain-computer interfaces*, pages 181–199. Springer, 2010.
- [145] Yue Zhao, Dan Davis, Guanliang Chen, Christoph Lofi, Claudia Hauff, and Geert-Jan Houben. Certificate achievement unlocked: How does mooc learners’ behaviour change? In *Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization*, pages 83–88. ACM, 2017.
- [146] Yue Zhao, Christoph Lofi, and Claudia Hauff. EC-TEL 2017 companion webpage. https://yue-zhao.github.io/MWDET_Project/, 2017.
- [147] Yue Zhao, Christoph Lofi, and Claudia Hauff. Scalable mind-wandering detection for moocs: A webcam-based approach. In *European Conference on Technology Enhanced Learning*, pages 330–344. Springer, 2017.
- [148] Yue Zhao, Tarmo Robal, Christoph Lofi, and Claudia Hauff. Can I have a Mooc2Go, please? on the viability of mobile vs. stationary learning. In *European Conference on Technology Enhanced Learning*, pages 101–115. Springer, 2018.
- [149] Yue Zhao, Tarmo Robal, Christoph Lofi, and Claudia Hauff. Stationary vs. non-stationary mobile learning in moocs. In *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization*, pages 299–303. ACM, 2018.
- [150] Barry J Zimmerman. Self-regulated learning and academic achievement: An overview. *Educational psychologist*, 25(1):3–17, 1990.

List of Figures

1.1	The completion rates of MOOCs provided by Delft University of Technology on edX from 2013 to 2017. Some MOOCs ran repeatedly in these years.	3
2.1	Overview of the fraction of scores that learners can earn in each unit. The passing threshold for SE is $\tau_{pass} = 58\%$, while it is $\tau_{pass} = 60\%$ for the other three MOOCs. Best viewed in color.	18
2.2	Total number of certificate earners (i.e. “passers”) at the end of each unit.	18
2.3	Quiz score distribution of FP and DA: passers are binned according to their passing unit. Rows represent groups of passers, columns represent one particular unit. Red plots show groups of passers that reached the passing threshold in a previous unit.	25
2.4	Quiz score distribution of SE and SEW: passers are binned according to their passing unit. Rows represent groups of passers, columns represent one particular unit. Red plots show groups of passers that reached the passing threshold in a previous unit.	26
2.5	Video consumption distribution of FP and DA: passers are binned according to their passing unit. Rows represent groups of passers, columns represent a particular unit. Red plots show groups of passers that reached the passing threshold in a previous unit.	27

2.6	Video consumption distribution of SE and SEW: passers are binned according to their passing unit. Rows represent groups of passers, columns represent a particular unit. Red plots show groups of passers that reached the passing threshold in a previous unit.	28
2.7	K-means clustering of learners normalized quiz score feature vectors for the first three unit-n passers groups (in SEW, learners' scores can reach τ_{pass} already in Unit 4). The cluster label in each graph shows the number of passers in each cluster. The vertical red line indicates the unit in which passers reached the passing threshold. The shaded areas around the lines show the upper (99%) and lower (70%) confidence bounds. Best viewed in color.	30
3.1	An example of a participant's learning progress. In this example, <i>video duration (VD)</i> is $V_1 + V_2 + V_3 + V_4$, <i>initial video watching duration</i> is $V_1 + V_2$, <i>video rewinding duration (VRD)</i> is $V_3 + V_4$, <i>question duration (QD)</i> is $Q_1 + V_3 + Q_2 + V_4 + Q_3$, and <i>question answering duration</i> is $Q_1 + Q_2 + Q_3$	44
3.2	The time participants spend on difference activities in <i>StaSc</i> and <i>MovSc</i>	48
3.3	Linear correlation coefficient between participants' learning performance, interactions and their perceived workload as measured through the NASA TLX form. The x-axis label also shows the average score of each workload dimension across our participants.	49
4.1	Gaze heatmaps of two study participants over a 30-second interval	57
4.2	Overview of the processing pipeline	60
4.3	An example inattention report	61
4.4	Overview of the reported inattention reports across the MOOC videos. Due to the randomized video order in the experiment, we partition the results according to whether the video was shown first ("order 1") or last ("order 2"). The video time displays the number of seconds since the start of the video. . .	66

5.1	Face fitting model generated by <i>CLM</i> . This example shows a common face fitting error due to hand positioning.	79
5.2	Face Boundary Box detected by tracking.js	80
5.3	Opening screen of the user study	82
5.4	Example task “cue card” of the user study.	83
6.1	<i>IntelliEye</i> ’s high-level architecture. The profiling and logger modules are always active; the attention tracking and alerting modules are only enabled if supported setup is detected and learner has granted access to webcam feed.	95
6.2	<i>IntelliEye</i> welcome screen.	101
6.3	<i>IntelliEye</i> ’s video player interface (arrow) embedded in the edX video player widget. The red hue around the video player is the visual alert we experiment with.	102
6.4	Distribution of video sessions and unique learners. A learner may be listed in more than one session type.	104

List of Tables

2.1	Overview of the summative assessment type(s) and average passing threshold τ_{pass} of all 46 computer science & programming MOOCs (in English, geared at undergraduates) open for enrolment on the edX platform on October 15, 2016. Assessments types include multiple choice (MC), fill-in-the-blank (FB), code submissions (CS), peer reviews (PR) and discussions (D). The column Early Passing shows the number of courses learners can pass before the final assessment is released.	14
2.2	Overview of the four MOOC datasets analyzed. #Questions refers to graded questions in MOOCs. #Attempts refers to how many attempts a learner has for each graded question. #Engaged refers to learners who watched at least one lecture video or answered one question. Compl Rate refers to the completion rate of a MOOC.	17
2.3	Overview of the five pre-passing behavior patterns for normalized quiz scores.	22
2.4	Overview of the six post-passing behavior patterns for normalized quiz scores.	22
2.5	Overview of the fraction of early passers falling into each course pattern (combination of pre- and post-passing pattern). P_n represents unit-n passers, shown in brackets (table header) are the number of passers in each group. Marked in bold are all values $\geq 10\%$. Each column sums up to 100%; a – indicates 0%.	32

3.1	Overview of our mini-MOOCs, the video length per MOOC and the minimum/median/maximum of participants' prior knowledge test scores on the topics. The highest possible score per topic is 20.	41
3.2	Number of participants under different experimental conditions.	42
3.3	The average value and standard deviation of metrics on participants' learning gain, learning efficiency and interactions under different experimental variables. † indicates significance at $p < 0.1$ level. ‡ indicates significance at $p < 0.05$ level. ◊ indicates significance at $p < 0.01$ level.	46
4.1	Features leveraged in the detection of participants' inattention	63
4.2	Inattention detection results based on gaze data (G means global features and L means local features)	67
4.3	Statistics of detection results on individual participants ($P_{highest}$ shows the detection results of the participant with highest F1-measure, P_{lowest} with lowest)	68
4.4	Detection across the entire length of the video ($Part\ 1$ means the first half part of the video, and $Part\ 2$ means the second half part of the video)	69
4.5	Detection with model translation (i.e. using a model on a different video than it was trained on)	70
5.1	Tobii's delay between the start of a face-miss/likely-face-miss task and the first face-miss event. The data is averaged across all participants of a single task.	84
5.2	Overview of the impact of the participants' background on <i>TJS</i> 's and <i>Tobii</i> 's accuracy.	84

5.3 Overview of all fifty benchmark tasks, and the accuracy (in %) of *CLM*, *TJS* and *Tobii* averaged across the 20 participants in our user study. For “(Likely) Face-Miss“ tasks we report the percentage of detected face misses, i.e. the eye tracker flags frames as not containing a face. For “Face-Hit” tasks we report the percentage of detected face hits. A higher percentage indicate a better performance. The best performance per task is shown in bold. † Note that for tasks Q_1 and Q_2 *Tobii*’s camera was not covered and the detection reflects participants’ hand moving through *Tobii*’s camera viewport to cover the webcam on the top bezel of the experimental laptop; for task Q_{43} there was no gaze to detect for *Tobii*. 85

6.1 Learners exposed to *IntelliEye*. Shown is the number of learners: (i) in each alert type condition, (ii) with at least one session with supported setup, (iii) who used *IntelliEye* at least once, and (iv) not accepting *IntelliEye*. 103

6.2 Learner attributes partitioned according to the use of *IntelliEye* (choices made on welcome page are not considered in grouping). Only learners with at least one supported video session are considered. * indicates Student’s t-test significance at $p < 0.05$ level. † and ‡ indicate Mann-Whitney U test significance at $p < 0.05$ and $p < 0.01$ levels respectively. 106

6.3 *IntelliEye* usage transition probabilities between subsequent video sessions; E=Enabled, D=Disabled, EF=Enabled Forever, DF=Disabled Forever. 106

6.4 Number of sessions with *IntelliEye* initially enabled grouped by the time it is switched off in the session. 107

6.5 Reasons provided for disabling *IntelliEye* forever. 108

6.6 Overview of the impact of *IntelliEye* on learners’ behaviors. There are 37 (pause), 27 (auditory) and 41 (visual) learners in each group. † indicates significance at $p < 0.05$ level between the first half and the second half of the *IntelliEye* sessions (Mann-Whitney U test). * indicates significance at $p < 0.05$ level between the marked group and the visual alert group (Mann-Whitney U test). 109

Summary

Learning Analytics Technology to Understand Learner Behavioral Engagement in MOOCs

As one of the most prominent examples of technology-enhanced learning, massive open online courses (MOOCs) have attracted extensive attention of learners, educators, and researchers since 2012. However, a low completion rate is a ubiquitous and severe problem in MOOCs, which means that only a small portion of learners got scores higher than or equal to the course requirements in MOOCs. Learner engagement is commonly presumed to be highly related to the completion rates of MOOCs.

In traditional classrooms, learner engagement can be observed by experienced educators. They can keep learner engagement by adjusting the course content and the way they teach. However, educators cannot observe learner engagement in MOOC learning the same way they usually do in traditional classrooms, while many learners lack skills to keep their engagement by themselves, which leads to high dropout rates of MOOCs. To observe learner engagement in MOOCs and provide learners feedback about their learning progress, learning analytics technology has been used by educators and researchers on MOOC platforms.

Learner engagement is usually investigated in three dimensions: behavioral engagement, emotional engagement, and cognitive engagement. In this thesis, we focus on using learning analytics technology to understand learner behavioral engagement in MOOCs. While many activities related to learner emotional engagement and cognitive engagement happen outside of MOOC platforms, most activities related to learner behavioral engagement are captured by the technology of MOOC platforms. Specifically, we study learner behavioral engagement on three time scales: throughout a course, in a learning session, and in a short period of time.

First, we explore learner behavioral engagement throughout a course based on learning analytics technology with large-scale trace data. To investigate the change of learner behavior after clinching a passing grade, we define a set of pre-passing and post-passing behavior patterns in our study. We present a data-driven approach which analyzes trace data from four thousand learners whose scores met the course requirements and find a certain subset of learners who heavily reduced their engagement in question answering after clinching a passing grade. Our study suggests that the course structure and grading schema of MOOCs should be designed to assign a certificate to learners only when they display mastery of an entire course subject.

Second, to investigate learner behavioral engagement in mobile learning sessions, we measure the impact of divided engagement and real-world environments on learner performance and interactions. We conducted a study which requires learners to have mobile learning sessions while sitting in the lab and walking on campus. To measure the impact of multitasking and divide attention, the trace data of learners and their answers in the questionnaire are analyzed. We find that learning on-the-go contributed to lowered learning performance and learners show different time arrangement in video watching and question answering while walking with learning.

Third, we investigate learner behavioral engagement in a short period of time. Specifically, we focus on tracking learner attention during video watching. Many MOOCs are centered around video lectures and learners can easily lose their attention while watching videos. If learner inattention can be detected automatically and in real-time, interventions can be provided to MOOC learners once they are being disengaged. We first propose an eye-tracking based method and our lab study indicates that it is possible to deploy a large-scale application of the webcam-based inattention detection in MOOCs. To avoid a high detection lag, low accuracy, and the complexity of design and maintenance in the eye-tracking method, we propose another method with face-tracking. We deploy our face-tracking based inattention detection method as a widget IntelliEye in real MOOCs. Through the deployment of IntelliEye, we find that most learners have capable setups to run our widget and one-third of them were willing to use it. Based on analyzing learner trace data, we observe high levels of learner inattention and their adaption toward our attention tracking technology.

Samenvatting

Learning Analytics-technologie om de betrokkenheid van leerlingen in MOOC's te begrijpen

Als een van de meest prominente voorbeelden van door technologie ondersteund leren, hebben massale open online cursussen (MOOC's) sinds 2012 uitgebreid de aandacht getrokken van studenten, docenten en onderzoekers. Een lage ratio van voltooiing is echter een alomtegenwoordig en serieus probleem bij MOOC's en dat betekent dat slechts een klein deel van de leerlingen een score behaalt die ten minste gelijk is aan de cursusvereisten van de MOOC. De betrokkenheid van de leerlingen wordt in het algemeen verondersteld sterk gerelateerd te zijn aan de ratio van voltooiing van MOOC's.

In traditionele klaslokalen kan de betrokkenheid van leerlingen worden waargenomen door ervaren docenten. Ze kunnen de betrokkenheid van de leerlingen behouden door de inhoud van de cursus aan te passen en de manier waarop ze lesgeven. Docenten kunnen echter de betrokkenheid van leerlingen bij MOOC's niet zien zoals ze dat in traditionele klaslokalen doen, terwijl veel leerlingen de vaardigheden missen om zelf hun betrokkenheid te behouden en dat leidt tot hoge uitvalcijfers voor MOOC's. Om de betrokkenheid van leerlingen bij MOOC's te observeren en leerlingen feedback te geven over hun leerproces, wordt "learning analytics"-technologie gebruikt door docenten en onderzoekers op MOOC-platforms.

De betrokkenheid van de leerling wordt meestal onderzocht in drie dimensies: gedragsmatige betrokkenheid, emotionele betrokkenheid en cognitieve betrokkenheid. In dit proefschrift richten we ons op het gebruik van learning analytics-technologie om de betrokkenheid van leerlingen in MOOC's te begrijpen. Terwijl veel activiteiten met betrekking tot emotionele betrokkenheid en cognitieve betrokkenheid van leerlingen plaatsvinden buiten MOOC-platforms, worden de meeste activiteiten met betrekking

tot het gedrag van de leerlingen vastgelegd door de technologie van MOOC-platforms. We bestuderen daarom met name de gedragsinteractie van de leerlingen, op drie tijdschalen: gedurende een cursus, in een leersessie en in een korte tijdsperiode.

Ten eerste onderzoeken we de betrokkenheid van leerlingen tijdens een cursus op basis van learning analytics-technologie met grootschalige log-gegevens. Om de verandering van het gedrag van de leerling te onderzoeken na het behalen van een voldoende cijfer, definiëren we een reeks gedragspatronen voorafgaand aan het passeren en na het passeren van onze studie. We presenteren een gegevensgestuurde aanpak die log-gegevens van vierduizend leerlingen analyseert waarvan de scores voldeden aan de cursusvereisten en vinden een bepaalde subset van leerlingen die hun betrokkenheid bij het beantwoorden van vragen sterk hebben verminderd na het behalen van een voldoende cijfer. Onze studie suggereert dat de cursusstructuur en het beoordelingsschema van MOOC's zodanig moeten zijn ontworpen dat een certificaat alleen aan studenten wordt toegekend wanneer zij het volledige vak beheersen.

Ten tweede, om de betrokkenheid van leerlingen in mobiele leersessies te onderzoeken, meten we de impact van gedeelde betrokkenheid en realistische omgevingen op leerprestaties en interacties. We hebben een onderzoek uitgevoerd waarbij leerlingen verplicht zijn om mobiele leersessies te volgen terwijl ze in het lab zitten en op de campus lopen. Om de impact van multitasking te meten en de aandacht te verdelen, worden de log-gegevens van leerlingen en hun antwoorden in de vragenlijst geanalyseerd. We merken dat leren "on-the-go" heeft bijgedragen aan verlaagde leerprestaties en dat leerlingen verschillende tijdsbestedingen hebben bij het kijken naar video's en het beantwoorden van vragen tijdens het "lopende" leren.

Ten derde onderzoeken we het gedrag van leerlingen in een korte periode van tijd. We concentreren ons met name op het volgen van de aandacht van leerlingen tijdens het kijken naar video's. Veel MOOC's zijn gecentreerd rond videocolleges en leerlingen kunnen gemakkelijk hun aandacht verliezen tijdens het kijken naar video's. Als de onoplettendheid van de leerlingen automatisch en in actuele tijd kan worden gedetecteerd, kunnen interventies aan MOOC-leerlingen worden gegeven zodra ze de betrokkenheid verliezen. We stellen eerst een op eye-tracking gebaseerde methode voor en onze laboratoriumstudie geeft aan dat het mogelijk is om een grootschalige toepassing van de op webcams gebaseerde onoplettendheidsdetectie in MOOC's te implementeren. Om hoge detectie-vertraging, lage nauwkeurigheid en de complexiteit van ontwerp en onderhoud in de eye-tracking methode te voorkomen,

stellen we een andere methode voor met face-tracking. We gebruiken onze face-tracking gebaseerde onoplettendheidsdetectiemethode als een widget IntelliEye in echte MOOC's. Door de inzet van IntelliEye ontdekken we dat de meeste leerlingen in staat zijn om onze widget uit te voeren en dat een derde van hen bereid is om het te gebruiken. Gebaseerd op het analyseren van de log-data, zien we een hoge mate van onoplettendheid van de leerling en zien we zijn aanpassing aan onze technologie voor het volgen van de aandacht.

Curriculum Vitae

Yue Zhao is currently a Ph.D. student in Web Information Systems (WIS) group, Delft University of Technology. He is supervised by Prof.dr.ir Geert-Jan Houben, Dr. Claudia Hauff, and Dr. Christoph Lofi. His research focuses on Learning Analytics Technology in MOOC learning.

Yue Zhao was born in Zhengzhou, China on March 28, 1989. He obtained his Bachelor degree and Master degree in Computer Science from National University of Defense Technology, Changsha, China in 2011 and 2013 respectively. As a Ph.D. student at Delft University of Technology, he was sponsored by China Scholarship Council (CSC) from September 2014 to September 2018.

Publications

1. Yue Zhao, Tarmo Robal, Christoph Lofi and Claudia Hauff. “Can I have a MOOC2Go, please? On the viability of mobile vs. stationary learning.” In 13th European Conference on Technology Enhanced Learning, pp. 101-115. Springer, Cham, 2018. (*Best Paper Nomination*)
2. Yue Zhao, Tarmo Robal, Christoph Lofi and Claudia Hauff. “Stationary vs. non-stationary mobile learning in MOOCs.” In Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization, pp. 299-303. ACM, 2018.
3. Tarmo Robal*, Yue Zhao*, Christoph Lofi and Claudia Hauff. “IntelliEye: enhancing MOOC learners’ video watching experience through real-time attention tracking.” In Proceedings of the 29th on Hypertext and Social Media, pp. 106-114. ACM, 2018. (**: Equal Contribution*).

4. Tarmo Robal, Yue Zhao, Christoph Lofi and Claudia Hauff. “Webcam-based attention tracking in online learning: a feasibility study.” In 23rd International Conference on Intelligent User Interfaces, pp. 189-197. ACM, 2018.
5. Yue Zhao, Dan Davis, Guanliang Chen, Christoph Lofi, Claudia Hauff and Geert-Jan Houben. “Certificate achievement unlocked: how does MOOC learners’ behaviour change?.” In Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization, pp. 83-88. ACM, 2017.
6. Yue Zhao, Christoph Lofi, and Claudia Hauff. “Scalable mind-wandering detection for MOOCs: a webcam-based approach.” In 12th European Conference on Technology Enhanced Learning, pp. 330-344. Springer, Cham, 2017. (*Best Paper Nomination*)
7. Yue Zhao and Claudia Hauff. “Sub-document timestamping: a study on the content creation dynamics of web documents.” In International Conference on Theory and Practice of Digital Libraries, pp. 203-214. Springer, Cham, 2016.
8. Yue Zhao and Claudia Hauff. “Temporal query intent disambiguation using time-series data.” In Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval, pp. 1017-1020. ACM, 2016.
9. Yue Zhao and Claudia Hauff. “WIS@ the NTCIR-12 Temporalia-2 task.” In Proceedings of the 12th NTCIR Conference. 2016.
10. Yue Zhao and Claudia Hauff. “Sub-document timestamping of web documents.” In Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1023-1026. ACM, 2015.
11. Yue Zhao and Claudia Hauff. “Temporal information retrieval revisited: a focused study on the web.” In Proceedings of the 6th Symposium on Future Directions in Information Access, pp. 24-27. BCS Learning & Development Ltd., 2015.
12. Xiaoyong Li, Yijie Wang, Yue Zhao, Yuan Wang and Xiaoling Li. “GPS: a general framework for parallel queries over data streams in cloud.” In 2013 IEEE 10th International Conference on High Performance Computing and Communications & 2013 IEEE International

Conference on Embedded and Ubiquitous Computing, pp. 1139-1146. IEEE, 2013.

13. Yue Zhao, Yong Tang, Yijie Wang, Shuhui Chen. “Generating malware signature using transcoding from sequential data to amino acid sequence.” In 2013 International Conference on High Performance Computing & Simulation, pp. 266-272. IEEE, 2013.

SIKS Dissertation Series

Since 1998, all dissertations written by Ph.D. students who have conducted their research under auspices of a senior research fellow of the SIKS research school are published in the SIKS Dissertation Series.

- 2019-11** Yue Zhao (TUD), *Learning Analytics Technology to Understand Learner Behavioral Engagement in MOOCs*
- 2019-10** Qing Chuan Ye (EUR), *Multi-objective Optimization Methods for Allocation and Prediction*
- 2019-09** Fahimeh Alizadeh Moghaddam (UVA), *Self-adaptation for energy efficiency in software systems*
- 2019-08** Frits de Nijs (TUD), *Resource-constrained Multi-agent Markov Decision Processes*
- 2019-07** Soude Fazeli (OUN), *Recommender Systems in Social Learning Platforms*
- 2019-06** Chris Dijkshoorn (VUA), *Niche sourcing for Improving Access to Linked Cultural Heritage Datasets*
- 2019-05** Sebastiaan van Zelst (TUE), *Process Mining with Streaming Data*
- 2019-04** Ridho Rahmadi (RUN), *Finding stable causal structures from clinical data*
- 2019-03** Eduardo Gonzalez Lopez de Murillas (TUE), *Process Mining on Databases: Extracting Event Data from Real Life Data Sources*
- 2019-02** Emmanuelle Beauxis-Aussalet (CWI, UU), *Statistics and Visualizations for Assessing Class Size Uncertainty*
- 2019-01** Rob van Eijk (UL), *Comparing and Aligning Process Representations*
- 2018-30** Wouter Beek (VU), *The "K" in "semantic web" stands for "knowledge": scaling semantics to the web*
- 2018-29** Yu Gu (UVT), *Emotion Recognition from Mandarin Speech*
- 2018-28** Christian Willemsse (UT), *Social Touch Technologies: How they feel and how they make you feel*
- 2018-27** Maikel Leemans (TUE), *Hierarchical Process Mining for Scalable Software Analysis*
- 2018-26** Roelof de Vries (UT), *Theory-Based And Tailor-Made: Motivational Messages for Behavior Change Technology*
- 2018-25** Riste Gligorov (VUA), *Serious Games in Audio-Visual Collections*
- 2018-24** Jered Vroon (UT), *Responsive Social Positioning Behaviour for Semi-Autonomous Telepresence Robots*
- 2018-23** Kim Schouten (EUR), *Semantics-driven Aspect-Based Sentiment Analysis*
- 2018-22** Eric Fernandes de Mello Araujo (VUA), *Contagious: Modeling the Spread of Behaviours, Perceptions and Emotions in Social Networks*
- 2018-21** Aad Sloomaker (OUN), *EMERGO: a generic platform for authoring and playing scenario-based serious games*
- 2018-20** Manxia Liu (RUN), *Time and Bayesian Networks*
- 2018-19** Minh Duc Pham (VUA), *Emergent relational schemas for RDF*
- 2018-18** Henriette Nakad (UL), *De Notaris en Private Rechtspraak*
- 2018-17** Jianpeng Zhang (TUE), *On Graph Sample Clustering*
- 2018-16** Jaebok Kim (UT), *Automatic recognition of engagement and emotion in a group of children*
- 2018-15** Naser Davarzani (UM), *Biomarker discovery in heart failure*
- 2018-14** Bart Joosten (UVT), *Detecting Social Signals with Spatiotemporal Gabor Filters*
- 2018-13** Seyed Amin Tabatabaei (VUA), *Using behavioral context in process mining: Exploring the added value of computational models for increasing the use of renewable energy in the residential sector*
- 2018-12** Xixi Lu (TUE), *Using behavioral context in process mining*
- 2018-11** Mahdi Sargolzaei (UVA), *Enabling Framework for Service-oriented Collaborative Networks*
- 2018-10** Julienka Mollee (VUA), *Moving forward: supporting physical activity behavior change through intelligent technology*
- 2018-09** Xu Xie (TUD), *Data Assimilation in Discrete Event Simulations*
- 2018-08** Rick Smetsers (RUN), *Advances in Model Learning for Software Systems*
- 2018-07** Jieting Luo (UU), *A formal account of opportunism in multi-agent systems*
- 2018-06** Dan Ionita (UT), *Model-Driven Information Security Risk Assessment of Socio-Technical Systems*
- 2018-05** Hugo Huurdeman (UVA), *Supporting the Complex Dynamics of the Information Seeking Process*
- 2018-04** Jordan Janeiro (TUD), *Flexible Coordination Support for Diagnosis Teams in Data-Centric Engineering Tasks*
- 2018-03** Steven Bosems (UT), *Causal Models For Well-Being: Knowledge Modeling, Model-Driven Development of Context-Aware Applications, and Behavior Prediction*
- 2018-02** Felix Mannhardt (TUE), *Multi-perspective Process Mining*
- 2018-01** Han van der Aa (VUA), *Comparing and Aligning Process Representations*
- 2017-48** Angel Suarez (OU), *Collaborative inquiry-based learning*
- 2017-47** Jie Yang (TUD), *Crowd Knowledge Creation Acceleration*
- 2017-46** Jan Schneider (OU), *Sensor-based Learning Support*
- 2017-45** Bas Testerink (UU), *Decentralized Runtime Norm Enforcement*
- 2017-44** Garm Lucassen (UU), *Understanding User Stories - Computational Linguistics in Agile Requirements Engineering*
- 2017-43** Maaikje de Boer (RUN), *Semantic Mapping in Video Retrieval*
- 2017-42** Elena Sokolova (RUN), *Causal discovery from mixed and missing data with applications on ADHD datasets*
- 2017-41** Adnan Manzoor (VUA), *Minding a Healthy Lifestyle: An Exploration of Mental Processes and a*

- Smart Environment to Provide Support for a Healthy Lifestyle*
- 2017-40** Altaf Hussain Abro (VUA), *Steer your Mind: Computational Exploration of Human Control in Relation to Emotions, Desires and Social Support For applications in human-aware support systems**
- 2017-39** Sara Ahmadi (RUN), *Exploiting properties of the human auditory system and compressive sensing methods to increase noise robustness in ASR*
- 2017-38** Alex Kayal (TUD), *Normative Social Applications*
- 2017-37** Alejandro Montes Garca (TUE), *WiBAF: A Within Browser Adaptation Framework that Enables Control over Privacy*
- 2017-36** Yuanhao Guo (UL), *Shape Analysis for Phenotype Characterisation from High-throughput Imaging*
- 2017-35** Martine de Vos (VU), *Interpreting natural science spreadsheets*
- 2017-34** Maren Scheffel (OUN), *The Evaluation Framework for Learning Analytics*
- 2017-33** Brigit van Loggem (OU), *Towards a Design Rationale for Software Documentation: A Model of Computer-Mediated Activity*
- 2017-32** Thaeer Samar (RUN), *Access to and Retrievability of Content in Web Archives*
- 2017-31** Ben Ruijl (UL), *Advances in computational methods for QFT calculations*
- 2017-30** Wilma Latuny (UVT), *The Power of Facial Expressions*
- 2017-29** Adel Alhuraibi (UVT), *From IT-BusinessStrategic Alignment to Performance: A Moderated Mediation Model of Social Innovation, and Enterprise Governance of IT*
- 2017-28** John Klein (VU), *Architecture Practices for Complex Contexts*
- 2017-27** Michiel Joosse (UT), *Investigating Positioning and Gaze Behaviors of Social Robots: People's Preferences, Perceptions and Behaviors*
- 2017-26** Merel Jung (UT), *Socially intelligent robots that understand and respond to human touch*
- 2017-25** Veruska Zamborlini (VU), *Knowledge Representation for Clinical Guidelines, with applications to Multimorbidity Analysis and Literature Search*
- 2017-24** Chang Wang (TUD), *Use of Affordances for Efficient Robot Learning*
- 2017-23** David Graus (UVA), *Entities of Interest—Discovery in Digital Traces*
- 2017-22** Sara Magliacane (VU), *Logics for causal inference under uncertainty*
- 2017-21** Jeroen Linszen (UT), *Meta Matters in Interactive Storytelling and Serious Gaming (A Play on Worlds)*
- 2017-20** Mohammadbashir Sedighi (TUD), *Fostering Engagement in Knowledge Sharing: The Role of Perceived Benefits, Costs and Visibility*
- 2017-19** Jeroen Vuurens (TUD), *Proximity of Terms, Texts and Semantic Vectors in Information Retrieval*
- 2017-18** Ridho Reinanda (UVA), *Entity Associations for Search*
- 2017-17** Daniel Dimov (UL), *Crowdsourced Online Dispute Resolution*
- 2017-16** Aleksandr Chuklin (UVA), *Understanding and Modeling Users of Modern Search Engines*
- 2017-15** Peter Berck, Radboud University (RUN), *Memory-Based Text Correction*
- 2017-14** Shoshannah Tekofsky (UvT), *You Are Who You Play You Are: Modelling Player Traits from Video Game Behavior*
- 2017-13** Gijs Huisman (UT), *Social Touch Technology - Extending the reach of social touch through haptic technology*
- 2017-12** Sander Leemans (TUE), *Robust Process Mining with Guarantees*
- 2017-11** Florian Kunnean (RUN), *Modelling patterns of time and emotion in Twitter #anticipointment*
- 2017-10** Robby van Delden (UT), *(Steering) Interactive Play Behavior*
- 2017-09** Dong Nguyen (UT), *Text as Social and Cultural Data: A Computational Perspective on Variation in Text*
- 2017-08** Rob Konijn (VU), *Detecting Interesting Differences: Data Mining in Health Insurance Data using Outlier Detection and Subgroup Discovery*
- 2017-07** Roel Bertens (UU), *Insight in Information: from Abstract to Anomaly*
- 2017-06** Damir Vandic (EUR), *Intelligent Information Systems for Web Product Search*
- 2017-05** Mahdieh Shadi (UVA), *Collaboration Behavior*
- 2017-04** Mrunal Gawade (CWI), *MULTI-CORE PARALLELISM IN A COLUMN-STORE*
- 2017-03** Daniël Harold Telgen (UU), *Grid Manufacturing; A Cyber-Physical Approach with Autonomous Products and Reconfigurable Manufacturing Machines*
- 2017-02** Sjoerd Timmer (UU), *Designing and Understanding Forensic Bayesian Networks using Argumentation*
- 2017-01** Jan-Jaap Oerlemans (UL), *Investigating Cybercrime*
- 2016-50** Yan Wang (UVT), *The Bridge of Dreams: Towards a Method for Operational Performance Alignment in IT-enabled Service Supply Chains*
- 2016-49** Gleb Polevoy (TUD), *Participation and Interaction in Projects. A Game-Theoretic Analysis*
- 2016-48** Tanja Buttler (TUD), *Collecting Lessons Learned*
- 2016-47** Christina Weber (UL), *Real-time foresight - Preparedness for dynamic innovation networks*
- 2016-46** Jorge Gallego Perez (UT), *Robots to Make you Happy*
- 2016-45** Bram van de Laar (UT), *Eaperienicing Brain-Computer Interface Control*
- 2016-44** Thibault Sellam (UVA), *Automatic Assistants for Database Exploration*
- 2016-43** Saskia Koldijk (RUN), *Context-Aware Support for Stress Self-Management: From Theory to Practice*
- 2016-42** Spyros Martzoukos (UVA), *Combinatorial and Compositional Aspects of Bilingual Aligned Corpora*
- 2016-41** Thomas King (TUD), *Governing Governance: A Formal Framework for Analysing Institutional Design and Enactment Governance*
- 2016-40** Christian Detweiler (TUD), *Accounting for Values in Design*
- 2016-39** Merijn Bruijnes (UT), *Believable Suspect Agents; Response and Interpersonal Style Selection for an Artificial Suspect*
- 2016-38** Andrea Minuto (UT), *MATERIALS THAT MATTER - Smart Materials meet Art & Interaction Design*
- 2016-37** Giovanni Sileno (UvA), *Aligning Law and Action - a conceptual and computational inquiry*
- 2016-36** Daphne Karreman (UT), *Beyond R2D2: The design of nonverbal interaction behavior optimized for robot-specific morphologies*
- 2016-35** Zhaochun Ren (UVA), *Monitoring Social Media: Summarization, Classification and Recommendation*
- 2016-34** Dennis Schunselar (TUE), *Configurable Process Trees: Elicitation, Analysis, and Enactment*
- 2016-33** Peter Bloem (UVA), *Single Sample Statistics, exercises in learning from just one example*
- 2016-32** Elco Vriezepakolk (UT), *Assessing Telecommunication Service Availability Risks for Crisis Organisations*
- 2016-31** Mohammad Khelghati (UT), *Deep web content monitoring*
- 2016-30** Ruud Mattheij (UvT), *The Eyes Have It*
- 2016-29** Nicolas Hning (TUD), *Peak reduction in decentralised electricity systems -Markets and prices for flexible planning*
- 2016-28** Mingxin Zhang (TUD), *Large-scale Agent-based Social Simulation - A study on epidemic prediction and control*
- 2016-27** Wen Li (TUD), *Understanding Geo-spatial Information on Social Media*
- 2016-26** Dilhan Thilakarathne (VU), *In or Out of Control: Exploring Computational Models to Study the Role of Human Awareness and Control in Behavioural Choices, with Applications in Aviation and Energy Management Domains*
- 2016-25** Julia Kiseleva (TU/e), *Using Contextual Information to Understand Searching and Browsing Behavior*
- 2016-24** Brend Wanders (UT), *Repurposing and Probabilistic Integration of Data; An Iterative and data model independent approach*
- 2016-23** Fei Cai (UVA), *Query Auto Completion in Information Retrieval*
- 2016-22** Grace Lewis (VU), *Software Architecture Strategies for Cyber-Foraging Systems*
- 2016-21** Alejandro Moreno Clieri (UT), *From Traditional to Interactive Playspaces: Automatic Analysis of Player Behavior in the Interactive Tag Playground*

- 2016-20** Daan Odijk (UVA), *Context & Semantics in News & Web Search*
- 2016-19** Julia Efreanova (Tu/e), *Mining Social Structures from Geographical Data*
- 2016-18** Albert Meroo Peuela (VU), *Refining Statistical Data on the Web*
- 2016-17** Berend Weel (VU), *Towards Embodied Evolution of Robot Organisms*
- 2016-16** Guangliang Li (UVA), *Socially Intelligent Autonomous Agents that Learn from Human Reward*
- 2016-15** Steffen Michels (RUN), *Hybrid Probabilistic Logics - Theoretical Aspects, Algorithms and Experiments*
- 2016-14** Ravi Khadka (UU), *Revisiting Legacy Software System Modernization*
- 2016-13** Nana Baah Gyan (VU), *The Web, Speech Technologies and Rural Development in West Africa - An ICT4D Approach*
- 2016-12** Max Knobbout (UU), *Logics for Modelling and Verifying Normative Multi-Agent Systems*
- 2016-11** Anne Schuth (UVA), *Search Engines that Learn from Their Users*
- 2016-10** George Karafotias (VUA), *Parameter Control for Evolutionary Algorithms*
- 2016-09** Archana Nottamkandath (VU), *Trusting Crowdsourced Information on Cultural Artefacts*
- 2016-08** Matje van de Camp (TiU), *A Link to the Past: Constructing Historical Social Networks from Unstructured Data*
- 2016-07** Jeroen de Man (VU), *Measuring and modeling negative emotions for virtual training*
- 2016-06** Michel Wilson (TUD), *Robust scheduling in an uncertain environment*
- 2016-05** Evgeny Sherkhonov (UVA), *Expanded Acyclic Queries: Containment and an Application in Explaining Missing Answers*
- 2016-04** Laurens Rietveld (VU), *Publishing and Consuming Linked Data*
- 2016-03** Maya Sappelli (RUN), *Knowledge Work in Context: User Centered Knowledge Worker Support*
- 2016-02** Michiel Christiaan Meulendijk (UU), *Optimizing medication reviews through decision support: prescribing a better pill to swallow*
- 2016-01** Syed Saiden Abbas (RUN), *Recognition of Shapes by Humans and Machines*
- 2015-35** Jungxao Xu (TUD), *Affective Body Language of Humanoid Robots: Perception and Effects in Human Robot Interaction*
- 2015-34** Victor de Graaf (UT), *Gesocial Recommender Systems*
- 2015-33** Frederik Schadd (TUD), *Ontology Mapping with Auxiliary Resources*
- 2015-32** Jerome Gard (UL), *Corporate Venture Management in SMEs*
- 2015-31** Yakup Koç (TUD), *On the robustness of Power Grids*
- 2015-30** Kiavash Bahreini (OU), *Real-time Multimodal Emotion Recognition in E-Learning*
- 2015-29** Hendrik Baier (UM), *Monte-Carlo Tree Search Enhancements for One-Player and Two-Player Domains*
- 2015-28** Janet Bagorogozo (TiU), *KNOWLEDGE MANAGEMENT AND HIGH PERFORMANCE; The Uganda Financial Institutions Model for HPO*
- 2015-27** Sándor Héman (CWI), *Updating compressed column stores*
- 2015-26** Alexander Hogenboom (EUR), *Sentiment Analysis of Text Guided by Semantics and Structure*
- 2015-25** Steven Woudenberg (UU), *Bayesian Tools for Early Disease Detection*
- 2015-24** Richard Berendsen (UVA), *Finding People, Papers, and Posts: Vertical Search Algorithms and Evaluation*
- 2015-23** Luit Gazendam (VU), *Cataloguer Support in Cultural Heritage*
- 2015-22** Zheming Zhu (UT), *Co-occurrence Rate Networks*
- 2015-21** Sibren Fetter (OUN), *Using Peer-Support to Expand and Stabilize Online Learning*
- 2015-20** Lois Vanhée (UU), *Using Culture and Values to Support Flexible Coordination*
- 2015-19** Bernardo Tabuenca (OUN), *Ubiquitous Technology for Lifelong Learners*
- 2015-18** Holger Pirk (CWI), *Waste Not, Want Not! - Managing Relational Data in Asymmetric Memories*
- 2015-17** André van Cleef (UT), *Physical and Digital Security Mechanisms: Properties, Combinations and Trade-offs*
- 2015-16** Changyun Wei (UT), *Cognitive Coordination for Cooperative Multi-Robot Teamwork*
- 2015-15** Klaas Andries de Graaf (VU), *Ontology-based Software Architecture Documentation*
- 2015-14** Bart van Straalen (UT), *A cognitive approach to modeling bad news conversations*
- 2015-13** Giuseppe Procaccianti (VU), *Energy-Efficient Software*
- 2015-12** Julie M. Birkholz (VU), *Modi Operandi of Social Network Dynamics: The Effect of Context on Scientific Collaboration Networks*
- 2015-11** Yongming Luo (TUE), *Designing algorithms for big graph datasets: A study of computing bisimulation and joins*
- 2015-10** Henry Hermans (OUN), *OpenU: design of an integrated system to support lifelong learning*
- 2015-09** Randy Klaassen (UT), *HCI Perspectives on Behavior Change Support Systems*
- 2015-08** Jie Jiang (TUD), *Organizational Compliance: An agent-based model for designing and evaluating organizational interactions*
- 2015-07** Maria-Hendrike Peetz (UvA), *Time-Aware Online Reputation Analysis*
- 2015-06** Farideh Heidari (TUD), *Business Process Quality Computation - Computing Non-Functional Requirements to Improve Business Processes*
- 2015-05** Christoph Bösch (UT), *Cryptographically Enforced Search Pattern Hiding*
- 2015-04** Howard Spoelstra (OUN), *Collaborations in Open Learning Environments*
- 2015-03** Twan van Laarhoven (RUN), *Machine learning for network data*
- 2015-02** Faiza Bukhsh (UvT), *Smart auditing: Innovative Compliance Checking in Customs Controls*
- 2015-01** Niels Netten (UvA), *Machine Learning for Relevance of Information in Crisis Response*
- 2014-47** Shangsong Liang (UVA), *Fusion and Diversification in Information Retrieval*
- 2014-46** Ke Tao (TUD), *Social Web Data Analytics: Relevance, Redundancy, Diversity*
- 2014-45** Birgit Schmitz (OU), *Mobile Games for Learning: A Pattern-Based Approach*
- 2014-44** Paulien Meesters (UvT), *Intelligent Blaww. Intelligence-gestuurde politiezorg in gebiedsgebonden eenheden*
- 2014-43** Kevin Vlaanderen (UU), *Supporting Process Improvement using Method Increments*
- 2014-42** Carsten Eickhoff (CWI/TUD), *Contextual Multidimensional Relevance Models*
- 2014-41** Frederik Hogenboom (EUR), *Automated Detection of Financial Events in News Text*
- 2014-40** Walter Oboma (RUN), *A Framework for Knowledge Management Using ICT in Higher Education*
- 2014-39** Jasmina Maric (UvT), *Web Communities, Immigration and Social Capital*
- 2014-38** Danny Plass-Oude Bos (UT), *Making brain-computer interfaces better: improving usability through post-processing*
- 2014-37** Maral Dadvar (UT), *Experts and Machines United Against Cyberbullying*
- 2014-36** Joos Buijs (TUE), *Flexible Evolutionary Algorithms for Mining Structured Process Models*
- 2014-35** Joost van Oijen (UU), *Cognitive Agents in Virtual Worlds: A Middleware Design Approach*
- 2014-34** Christina Manteli (VU), *The Effect of Governance in Global Software Development: Analyzing Transactive Memory Systems*
- 2014-33** Tesfa Tegegne Asfaw (RUN), *Service Discovery in eHealth*
- 2014-32** Naser Ayat (UVA), *On Entity Resolution in Probabilistic Data*
- 2014-31** Leo van Moergestel (UU), *Agent Technology in Agile Multiparallel Manufacturing and Product Support*
- 2014-30** Peter de Kock Berenschot (UvT), *Anticipating Criminal Behaviour*
- 2014-29** Jaap Kabbedijk (UU), *Variability in Multi-Tenant Enterprise Software*
- 2014-28** Anna Chmielowiec (VU), *Decentralized k-Clique Matching*
- 2014-27** Rui Jorge Almeida (EUR), *Conditional Density Models Integrating Fuzzy and Probabilistic Representations of Uncertainty*
- 2014-26** Tim Baarslag (TUD), *What to Bid and When to Stop*

- 2014-25** Martijn Lappenschaar (RUN), *New network models for the analysis of disease interaction*
- 2014-24** Davide Ceolin (VU), *Trusting Semi-structured Web Data*
- 2014-23** Eleftherios Sidirouros (UvA/CWI), *Space Efficient Indexes for the Big Data Era*
- 2014-22** Marieke Peeters (UU), *Personalized Educational Games - Developing agent-supported scenario-based training*
- 2014-21** Kassidy Clark (TUD), *Negotiation and Monitoring in Open Environments*
- 2014-20** Mena Habib (UT), *Named Entity Extraction and Disambiguation for Informal Text: The Missing Link*
- 2014-19** Vincius Ramos (TUE), *Adaptive Hypermedia Courses: Qualitative and Quantitative Evaluation and Tool Support*
- 2014-18** Mattijs Ghijsen (VU), *Methods and Models for the Design and Study of Dynamic Agent Organizations*
- 2014-17** Kathrin Dentler (VU), *Computing healthcare quality indicators automatically: Secondary Use of Patient Data and Semantic Interoperability*
- 2014-16** Krystyna Milián (VU), *Supporting trial recruitment and design by automatically interpreting eligibility criteria*
- 2014-15** Natalya Mogles (VU), *Agent-Based Analysis and Support of Human Functioning in Complex Socio-Technical Systems: Applications in Safety and Healthcare*
- 2014-14** Yangyang Shi (TUD), *Language Models With Meta-information*
- 2014-13** Arlette van Wissen (VU), *Agent-Based Support for Behavior Change: Models and Applications in Health and Safety Domains*
- 2014-12** Willem van Willigen (VU), *Look Ma, No Hands: Aspects of Autonomous Vehicle Control*
- 2014-11** Janneke van der Zwaan (TUD), *An Empathic Virtual Buddy for Social Support*
- 2014-10** Ivan Salvador Razo Zapata (VU), *Service Value Networks*
- 2014-09** Philip Jackson (UvT), *Toward Human-Level Artificial Intelligence: Representation and Computation of Meaning in Natural Language*
- 2014-08** Samur Araujo (TUD), *Data Integration over Distributed and Heterogeneous Data Endpoints*
- 2014-07** Arya Adriansyah (TUE), *Aligning Observed and Modeled Behavior*
- 2014-06** Damian Tamburri (VU), *Supporting Networked Software Development*
- 2014-05** Jurriaan van Reijssen (UU), *Knowledge Perspectives on Advancing Dynamic Capability*
- 2014-04** Hanna Jochmann-Mannak (UT), *Websites for children: search strategies and interface design - Three studies on children's search performance and evaluation*
- 2014-03** Sergio Raul Duarte Torres (UT), *Information Retrieval for Children: Search Behavior and Solutions*
- 2014-02** Fiona Tuliayo (RUN), *Combining System Dynamics with a Domain Modeling Method*
- 2014-01** Nicola Barile (UU), *Studies in Learning Monotone Models from Data*
- 2013-43** Marc Bron (UVA), *Exploration and Contextualization through Interaction and Concepts*
- 2013-42** Léon Planken (TUD), *Algorithms for Simple Temporal Reasoning*
- 2013-41** Jochem Liem (UVA), *Supporting the Conceptual Modelling of Dynamic Systems: A Knowledge Engineering Perspective on Qualitative Reasoning*
- 2013-40** Pim Nijssen (UM), *Monte-Carlo Tree Search for Multi-Player Games*
- 2013-39** Joop de Jong (TUD), *A Method for Enterprise Ontology based Design of Enterprise Information Systems*
- 2013-38** Eelco den Heijer (VU), *Autonomous Evolutionary Art*
- 2013-37** Dirk Börner (OUN), *Ambient Learning Displays*
- 2013-36** Than Lam Hoang (TUE), *Pattern Mining in Data Streams*
- 2013-35** Abdallah El Ali (UvA), *Minimal Mobile Human Computer Interaction*
- 2013-34** Kien Tjin-Kam-Jet (UT), *Distributed Deep Web Search*
- 2013-33** Qi Gao (TUD), *User Modeling and Personalization in the Microblogging Sphere*
- 2013-32** Kamakshi Rajagopal (OUN), *Networking For Learning; The role of Networking in a Lifelong Learner's Professional Development*
- 2013-31** Dinh Khoa Nguyen (UvT), *Blueprint Model and Language for Engineering Cloud Applications*
- 2013-30** Joyce Nakatumba (TUE), *Resource-Aware Business Process Management: Analysis and Support*
- 2013-29** Iwan de Kok (UT), *Listening Heads*
- 2013-28** Frans van der Sluis (UT), *When Complexity becomes Interesting: An Inquiry into the Information eXperience*
- 2013-27** Mohammad Huq (UT), *Inference-based Framework Managing Data Provenance*
- 2013-26** Alireza Zarghami (UT), *Architectural Support for Dynamic Homecare Service Provisioning*
- 2013-25** Agnieszka Anna Latoszek-Berendsen (UM), *Intention-based Decision Support. A new way of representing and implementing clinical guidelines in a Decision Support System*
- 2013-24** Haitham Bou Ammar (UM), *Automated Transfer in Reinforcement Learning*
- 2013-23** Patrício de Alencar Silva (UvT), *Value Activity Monitoring*
- 2013-22** Tom Claassen (RUN), *Causal Discovery and Logic*
- 2013-21** Sander Wubben (UvT), *Text-to-text generation by monolingual machine translation*
- 2013-20** Katja Hofmann (UvA), *Fast and Reliable Online Learning to Rank for Information Retrieval*
- 2013-19** Renze Steenhuisen (TUD), *Coordinated Multi-Agent Planning and Scheduling*
- 2013-18** Jeroen Janssens (UvT), *Outlier Selection and One-Class Classification*
- 2013-17** Koen Kok (VU), *The PowerMatcher: Smart Coordination for the Smart Electricity Grid*
- 2013-16** Eric Kok (UU), *Exploring the practical benefits of argumentation in multi-agent deliberation*
- 2013-15** Daniel Hennes (UM), *Multiagent Learning - Dynamic Games and Applications*
- 2013-14** Jafar Tanha (UVA), *Ensemble Approaches to Semi-Supervised Learning*
- 2013-13** Mohammad Safiri (UT), *Service Tailoring: User-centric creation of integrated IT-based homecare services to support independent living of elderly*
- 2013-12** Marian Razavian (VU), *Knowledge-driven Migration to Services*
- 2013-11** Evangelos Pournaras (TUD), *Multi-level Reconfigurable Self-organization in Overlay Services*
- 2013-10** Jeewanie Jayasinghe Arachchige (UvT), *A Unified Modeling Framework for Service Design*
- 2013-09** Fabio Gori (RUN), *Metagenomic Data Analysis: Computational Methods and Applications*
- 2013-08** Robbert-Jan Merk (VU), *Making enemies: cognitive modeling for opponent agents in fighter pilot simulators*
- 2013-07** Giel van Lankveld (UvT), *Quantifying Individual Player Differences*
- 2013-06** Romulo Goncalves (CWI), *The Data Cyclotron: Juggling Data and Queries for a Data Warehouse Audience*
- 2013-05** Dulce Pumareja (UT), *Groupware Requirements Evolutions Patterns*
- 2013-04** Chetan Yadati (TUD), *Coordinating autonomous planning and scheduling*
- 2013-03** Szymon Klarman (VU), *Reasoning with Contexts in Description Logics*
- 2013-02** Erietta Liarou (CWI), *MonetDB/DataCell: Leveraging the Column-store Database Technology for Efficient and Scalable Stream Processing*
- 2013-01** Viorel Milea (EUR), *News Analytics for Financial Decision Support*
- 2012-51** Jeroen de Jong (TUD), *Heuristics in Dynamic Scheduling; a practical framework with a case study in elevator dispatching*
- 2012-50** Steven van Kervel (TUD), *Ontology driven Enterprise Information Systems Engineering*
- 2012-49** Michael Kaisers (UM), *Learning against Learning - Evolutionary dynamics of reinforcement learning algorithms in strategic interactions*
- 2012-48** Jorn Bakker (TUE), *Handling Abrupt Changes in Evolving Time-series Data*
- 2012-47** Manos Tsagkias (UVA), *Mining Social Media: Tracking Content and Predicting Behavior*
- 2012-46** Simon Carter (UVA), *Exploration and Exploitation of Multilingual Data for Statistical Machine Translation*
- 2012-45** Benedikt Kratz (UvT), *A Model and Language for Business-aware Transactions*
- 2012-44** Anna Tordai (VU), *On Combining Alignment Techniques*

- 2012-42** Dominique Verpoorten (OU), *Reflection Amplifiers in self-regulated Learning*
- 2012-41** Sebastian Kelle (OU), *Game Design Patterns for Learning*
- 2012-40** Agus Gunawan (UvT), *Information Access for SMEs in Indonesia*
- 2012-39** Hassan Fatemi (UT), *Risk-aware design of value and coordination networks*
- 2012-38** Selmar Smit (VU), *Parameter Tuning and Scientific Testing in Evolutionary Algorithms*
- 2012-37** Agnes Nakakawa (RUN), *A Collaboration Process for Enterprise Architecture Creation*
- 2012-36** Denis Ssebugwawo (RUN), *Analysis and Evaluation of Collaborative Modeling Processes*
- 2012-35** Evert Haasdijk (VU), *Never Too Old To Learn – On-line Evolution of Controllers in Swarm- and Modular Robotics*
- 2012-34** Pavol Jancura (RUN), *Evolutionary analysis in PPI networks and applications*
- 2012-33** Rory Sie (OUN), *Coalitions in Cooperation Networks (COCOON)*
- 2012-32** Wietske Visser (TUD), *Qualitative multi-criteria preference representation and reasoning*
- 2012-31** Emily Bagarukayo (RUN), *A Learning by Construction Approach for Higher Order Cognitive Skills Improvement, Building Capacity and Infrastructure*
- 2012-30** Alina Pommeranz (TUD), *Designing Human-Centered Systems for Reflective Decision Making*
- 2012-29** Almer Tigelaar (UT), *Peer-to-Peer Information Retrieval*
- 2012-28** Nancy Pascall (UvT), *Engendering Technology Empowering Women*
- 2012-27** Hayrettin Gurkok (UT), *Mind the Sheep! User Experience Evaluation & Brain-Computer Interface Games*
- 2012-26** Emile de Maat (UVA), *Making Sense of Legal Text*
- 2012-25** Silja Eckartz (UT), *Managing the Business Case Development in Inter-Organizational IT Projects: A Methodology and its Application*
- 2012-24** Laurens van der Werff (UT), *Evaluation of Noisy Transcripts for Spoken Document Retrieval*
- 2012-23** Christian Muehl (UT), *Toward Affective Brain-Computer Interfaces: Exploring the Neurophysiology of Affect during Human Media Interaction*
- 2012-22** Thijs Vis (UvT), *Intelligence, politie en veiligheidsdienst: verenigbare grootheden?*
- 2012-21** Roberto Cornacchia (TUD), *Querying Sparse Matrices for Information Retrieval*
- 2012-20** Ali Bahramisharif (RUN), *Covert Visual Spatial Attention, a Robust Paradigm for Brain-Computer Interfacing*
- 2012-19** Helen Schonenberg (TUE), *What's Next? Operational Support for Business Process Execution*
- 2012-18** Eltjo Poort (VU), *Improving Solution Architecting Practices*
- 2012-17** Amal Elgammal (UvT), *Towards a Comprehensive Framework for Business Process Compliance*
- 2012-16** Fiemke Both (VU), *Helping people by understanding them - Ambient Agents supporting task execution and depression treatment*
- 2012-15** Natalie van der Wal (VU), *Social Agents. Agent-Based Modelling of Integrated Internal and Social Dynamics of Cognitive and Affective Processes.*
- 2012-14** Evgeny Knutov (TUE), *Generic Adaptation Framework for Unifying Adaptive Web-based Systems*
- 2012-13** Suleman Shahid (UvT), *Fun and Face: Exploring non-verbal expressions of emotion during playful interactions.*
- 2012-12** Kees van der Sluijs (TUE), *Model Driven Design and Data Integration in Semantic Web Information Systems*
- 2012-11** J.C.B. Rantham Prabhakara (TUE), *Process Mining in the Large: Preprocessing, Discovery, and Diagnostics*
- 2012-10** David Smits (TUE), *Towards a Generic Distributed Adaptive Hypermedia Environment*
- 2012-09** Ricardo Neisse (UT), *Trust and Privacy Management Support for Context-Aware Service Platforms*
- 2012-08** Gerben de Vries (UVA), *Kernel Methods for Vessel Trajectories*
- 2012-07** Rianne van Lambalgen (VU), *When the Going Gets Tough: Exploring Agent-based Models of Human Performance under Demanding Conditions*
- 2012-06** Wolfgang Reinhardt (OU), *Awareness Support for Knowledge Workers in Research Networks*
- 2012-05** Marijn Plomp (UU), *Maturing Interorganizational Information Systems*
- 2012-04** Jurriaan Souer (UU), *Development of Content Management System-based Web Applications*
- 2012-03** Adam Vanya (VU), *Supporting Architecture Evolution by Mining Software Repositories*
- 2012-02** Muhammad Umair (VU), *Adaptivity, emotion, and Rationality in Human and Ambient Agent Models*
- 2012-01** Terry Kakeeto (UvT), *Relationship Marketing for SMEs in Uganda*
- 2011-49** Andreea Niculescu (UT), *Conversational interfaces for task-oriented spoken dialogues: design aspects influencing interaction quality*
- 2011-48** Mark Ter Maat (UT), *Response Selection and Turn-taking for a Sensitive Artificial Listening Agent*
- 2011-47** Azizi Bin Ab Aziz (VU), *Exploring Computational Models for Intelligent Support of Persons with Depression*
- 2011-46** Beibe Hu (TUD), *Towards Contextualized Information Delivery: A Rule-based Architecture for the Domain of Mobile Police Work*
- 2011-45** Herman Stehouwer (UvT), *Statistical Language Models for Alternative Sequence Selection*
- 2011-44** Boris Reuderink (UT), *Robust Brain-Computer Interfaces*
- 2011-43** Henk van der Schuur (UU), *Process Improvement through Software Operation Knowledge*
- 2011-42** Michal Sindlar (UU), *Explaining Behavior through Mental State Attribution*
- 2011-41** Luan Ibraimi (UT), *Cryptographically Enforced Distributed Data Access Control*
- 2011-40** Viktor Clerc (VU), *Architectural Knowledge Management in Global Software Development*
- 2011-39** Joost Westra (UU), *Organizing Adaptation using Agents in Serious Games*
- 2011-38** Nyree Lemmens (UM), *Bee-inspired Distributed Optimization*
- 2011-37** Adriana Burlutiu (RUN), *Machine Learning for Pairwise Data, Applications for Preference Learning and Supervised Network Inference*
- 2011-36** Erik van der Spek (UU), *Experiments in serious game design: a cognitive approach*
- 2011-35** Maaik Harbers (UU), *Explaining Agent Behavior in Virtual Training*
- 2011-34** Paolo Turrini (UU), *Strategic Reasoning in Interdependence: Logical and Game-theoretical Investigations*
- 2011-33** Tom van der Weide (UU), *Arguing to Motivate Decisions*
- 2011-32** Nees-Jan van Eck (EUR), *Methodological Advances in Bibliometric Mapping of Science*
- 2011-31** Ludo Waltman (EUR), *Computational and Game-Theoretic Approaches for Modeling Bounded Rationality*
- 2011-30** Egon van den Broek (UT), *Affective Signal Processing (ASP): Unraveling the mystery of emotions*
- 2011-29** Faisal Kamiran (TUE), *Discrimination-aware Classification*
- 2011-28** Rianne Kaptein (UVA), *Effective Focused Retrieval by Exploiting Query Context and Document Structure*
- 2011-27** Aniel Bhulai (VU), *Dynamic website optimization through autonomous management of design patterns*
- 2011-26** Matthijs Aart Pontier (VU), *Virtual Agents for Human Communication - Emotion Regulation and Involvement-Distance Trade-Offs in Embodied Conversational Agents and Robots*
- 2011-25** Syed Waqar ul Qounain Jaffry (VU), *Analysis and Validation of Models for Trust Dynamics*
- 2011-24** Herwin van Welbergen (UT), *Behavior Generation for Interpersonal Coordination with Virtual Humans On Specifying, Scheduling and Realizing Multimodal Virtual Human Behavior*
- 2011-23** Wouter Weerkamp (UVA), *Finding People and their Utterances in Social Media*
- 2011-22** Junte Zhang (UVA), *System Evaluation of Archival Description and Access*
- 2011-21** Linda Terlouw (TUD), *Modularization and Specification of Service-Oriented Systems*
- 2011-20** Qing Gu (VU), *Guiding service-oriented software engineering - A view-based approach*
- 2011-19** Ellen Rusman (OU), *The Mind's Eye on Personal Profiles*
- 2011-18** Mark Ponsen (UM), *Strategic Decision-Making in complex games*

- 2011-17** Jiyin He (UVA), *Exploring Topic Structure: Coherence, Diversity and Relatedness*
- 2011-16** Maarten Schadd (UM), *Selective Search in Games of Different Complexity*
- 2011-15** Marijn Koolen (UvA), *The Meaning of Structure: the Value of Link Evidence for Information Retrieval*
- 2011-14** Milan Lovric (EUR), *Behavioral Finance and Agent-Based Artificial Markets*
- 2011-13** Xiaoyu Mao (UvT), *Airport under Control. Multiagent Scheduling for Airport Ground Handling*
- 2011-12** Carmen Bratosin (TUE), *Grid Architecture for Distributed Process Mining*
- 2011-11** Dhaval Vyas (UT), *Designing for Awareness: An Experience-focused HCI Perspective*
- 2011-10** Bart Bogaert (UvT), *Cloud Content Contention*
- 2011-09** Tim de Jong (OU), *Contextualised Mobile Media for Learning*
- 2011-08** Nieske Vergunst (UU), *BDI-based Generation of Robust Task-Oriented Dialogues*
- 2011-07** Yujia Cao (UT), *Multimodal Information Presentation for High Load Human Computer Interaction*
- 2011-06** Yiwen Wang (TUE), *Semantically-Enhanced Recommendations in Cultural Heritage*
- 2011-05** Base van der Raadt (VU), *Enterprise Architecture Coming of Age - Increasing the Performance of an Emerging Discipline.*
- 2011-04** Hado van Hasselt (UU), *Insights in Reinforcement Learning; Formal analysis and empirical evaluation of temporal-difference learning algorithms*
- 2011-03** Jan Martijn van der Werf (TUE), *Compositional Design and Verification of Component-Based Information Systems*
- 2011-02** Nick Tinnemeier (UU), *Organizing Agent Organizations. Syntax and Operational Semantics of an Organization-Oriented Programming Language*
- 2011-01** Botond Cseke (RUN), *Variational Algorithms for Bayesian Inference in Latent Gaussian Models*
- 2010-53** Edgar Meij (UVA), *Combining Concepts and Language Models for Information Access*
- 2010-52** Peter-Paul van Maanen (VU), *Adaptive Support for Human-Computer Teams: Exploring the Use of Cognitive Models of Trust and Attention*
- 2010-51** Alia Khairia Amin (CWI), *Understanding and supporting information seeking tasks in multiple sources*
- 2010-50** Bouke Huurnink (UVA), *Search in Audiovisual Broadcast Archives*
- 2010-49** Jahn-Takeshi Saito (UM), *Solving difficult game positions*
- 2010-47** Chen Li (UT), *Mining Process Model Variants: Challenges, Techniques, Examples*
- 2010-46** Vincent Pijpers (VU), *e3Alignment: Exploring Inter-Organizational Business-ICT Alignment*
- 2010-45** Vasilios Andrikopoulos (UvT), *A theory and model for the evolution of software services*
- 2010-44** Pieter Bellekens (TUE), *An Approach towards Context-sensitive and User-adapted Access to Heterogeneous Data Sources, Illustrated in the Television Domain*
- 2010-43** Peter van Kranenburg (UU), *A Computational Approach to Content-Based Retrieval of Folk Song Melodies*
- 2010-42** Sybren de Kinderen (VU), *Needs-driven service bundling in a multi-supplier setting - the computational e3-service approach*
- 2010-41** Guillaume Chaslot (UM), *Monte-Carlo Tree Search*
- 2010-40** Mark van Assem (VU), *Converting and Integrating Vocabularies for the Semantic Web*
- 2010-39** Ghazanfar Farooq Siddiqui (VU), *Integrative modeling of emotions in virtual agents*
- 2010-38** Dirk Fahland (TUE), *From Scenarios to components*
- 2010-37** Niels Lohmann (TUE), *Correctness of services and their composition*
- 2010-36** Jose Janssen (OU), *Paving the Way for Lifelong Learning; Facilitating competence development through a learning path specification*
- 2010-35** Dolf Trieschnigg (UT), *Proof of Concept: Concept-based Biomedical Information Retrieval*
- 2010-34** Teduh Dirgahayu (UT), *Interaction Design in Service Compositions*
- 2010-33** Robin Aly (UT), *Modeling Representation Uncertainty in Concept-Based Multimedia Retrieval*
- 2010-32** Marcel Hiel (UvT), *An Adaptive Service Oriented Architecture: Automatically solving Interoperability Problems*
- 2010-31** Victor de Boer (UVA), *Ontology Enrichment from Heterogeneous Sources on the Web*
- 2010-30** Marieke van Erp (UvT), *Accessing Natural History - Discoveries in data cleaning, structuring, and retrieval*
- 2010-29** Stratos Idreos (CWI), *Database Cracking: Towards Auto-tuning Database Kernels*
- 2010-28** Arne Koopman (UU), *Characteristic Relational Patterns*
- 2010-27** Marten Voulon (UL), *Automatisch contracteren*
- 2010-26** Ying Zhang (CWI), *XRPC: Efficient Distributed Query Processing on Heterogeneous XQuery Engines*
- 2010-25** Zulfiqar Ali Memon (VU), *Modelling Human-Awareness for Ambient Agents: A Human Mindreading Perspective*
- 2010-24** Dmytro Tykhonov, *Designing Generic and Efficient Negotiation Strategies*
- 2010-23** Bas Steunebrink (UU), *The Logical Structure of Emotions*
- 2010-22** Michiel Hildebrand (CWI), *End-user Support for Access to Heterogeneous Linked Data*
- 2010-21** Harold van Heerde (UT), *Privacy-aware data management by means of data degradation*
- 2010-20** Ivo Swartjes (UT), *Whose Story Is It Anyway? How Improv Informs Agency and Authorship of Emergent Narrative*
- 2010-19** Henriette Cramer (UvA), *People's Responses to Autonomous and Adaptive Systems*
- 2010-18** Charlotte Gerritsen (VU), *Caught in the Act: Investigating Crime by Agent-Based Simulation*
- 2010-17** Spyros Kotoulas (VU), *Scalable Discovery of Networked Resources: Algorithms, Infrastructure, Applications*
- 2010-16** Sicco Verwer (TUD), *Efficient Identification of Timed Automata, theory and practice*
- 2010-15** Lianne Bodenstaff (UT), *Managing Dependency Relations in Inter-Organizational Models*
- 2010-14** Sander van Splunter (VU), *Automated Web Service Reconfiguration*
- 2010-13** Gianluigi Folino (RUN), *High Performance Data Mining using Bio-inspired techniques*
- 2010-12** Susan van den Braak (UU), *Sensemaking software for crime analysis*
- 2010-11** Adriaan Ter Mors (TUD), *The world according to MARP: Multi-Agent Route Planning*
- 2010-10** Rebecca Ong (UL), *Mobile Communication and Protection of Children*
- 2010-09** Hugo Kielman (UL), *A Politiele gegevensverwerking en Privacy, Naar een effectieve waarborging*
- 2010-08** Krzysztof Siewicz (UL), *Towards an Improved Regulatory Framework of Free Software. Protecting user freedoms in a world of software communities and eGovernments*
- 2010-07** Wim Fikkert (UT), *Gesture interaction at a Distance*
- 2010-06** Sander Bakkes (UvT), *Rapid Adaptation of Video Game AI*
- 2010-05** Claudia Hauff (UT), *Predicting the Effectiveness of Queries and Retrieval Systems*
- 2010-04** Olga Kulyk (UT), *Do You Know What I Know? Situational Awareness of Co-located Teams in Multidisplay Environments*
- 2010-03** Joost Geurts (CWI), *A Document Engineering Model and Processing Framework for Multimedia documents*
- 2010-02** Ingo Wassink (UT), *Work flows in Life Science*
- 2010-01** Matthijs van Leeuwen (UU), *Patterns that Matter*
- 2009-46** Loredana Afanasiev (UvA), *Querying XML: Benchmarks and Recursion*
- 2009-45** Jilles Vreeken (UU), *Making Pattern Mining Useful*
- 2009-44** Roberto Santana Tapia (UT), *Assessing Business-IT Alignment in Networked Organizations*
- 2009-43** Virginia Nunes Leal Franqueira (UT), *Finding Multi-step Attacks in Computer Networks using Heuristic Search and Mobile Ambients*
- 2009-42** Toine Bogers (UvT), *Recommender Systems for Social Bookmarking*
- 2009-41** Igor Berezhin (UvT), *Digital Analysis of Paintings*

- 2009-40** Stephan Raaijmakers (UvT), *Multinomial Language Learning: Investigations into the Geometry of Language*
- 2009-39** Christian Stahl (TUE, Humboldt-Universitaet zu Berlin), *Service Substitution - A Behavioral Approach Based on Petri Nets*
- 2009-38** Riina Vuorikari (OU), *Tags and self-organisation: a metadata ecology for learning resources in a multilingual context*
- 2009-37** Hendrik Drachler (OUN), *Navigation Support for Learners in Informal Learning Networks*
- 2009-36** Marco Kalz (OUN), *Placement Support for Learners in Learning Networks*
- 2009-35** Wouter Koelwijn (UL), *Privacy en Politiegegevens; Over geautomatiseerde normatieve informatie-uitwisseling*
- 2009-34** Inge van de Weerd (UU), *Advancing in Software Product Management: An Incremental Method Engineering Approach*
- 2009-33** Khiet Truong (UT), *How Does Real Affect Affect Affect Recognition In Speech?*
- 2009-32** Rik Farenhorst (VU) and Remco de Boer (VU), *Architectural Knowledge Management: Supporting Architects and Auditors*
- 2009-31** Sofiya Katrenko (UVA), *A Closer Look at Learning Relations from Text*
- 2009-30** Marcin Zukowski (CWI), *Balancing vectorized query execution with bandwidth-optimized storage*
- 2009-29** Stanislav Pokraev (UT), *Model-Driven Semantic Integration of Service-Oriented Applications*
- 2009-28** Sander Evers (UT), *Sensor Data Management with Probabilistic Models*
- 2009-27** Christian Glahn (OU), *Contextual Support of social Engagement and Reflection on the Web*
- 2009-26** Fernando Koch (UU), *An Agent-Based Model for the Development of Intelligent Mobile Services*
- 2009-25** Alex van Ballegooij (CWI), *"RAM: Array Database Management through Relational Mapping"*
- 2009-24** Annerieke Heuvelink (VUA), *Cognitive Models for Training Simulations*
- 2009-23** Peter Hofgesang (VU), *Modelling Web Usage in a Changing Environment*
- 2009-22** Pavel Serdyukov (UT), *Search For Expertise: Going beyond direct evidence*
- 2009-21** Stijn Vanderlooy (UM), *Ranking and Reliable Classification*
- 2009-20** Bob van der Vecht (UU), *Adjustable Autonomy: Controlling Influences on Decision Making*
- 2009-19** Valentin Robu (CWI), *Modeling Preferences, Strategic Reasoning and Collaboration in Agent-Mediated Electronic Markets*
- 2009-18** Fabian Groffen (CWI), *Armada, An Evolving Database System*
- 2009-17** Laurens van der Maaten (UvT), *Feature Extraction from Visual Data*
- 2009-16** Fritz Reul (UvT), *New Architectures in Computer Chess*
- 2009-15** Rinke Hoekstra (UVA), *Ontology Representation - Design Patterns and Ontologies that Make Sense*
- 2009-14** Maksym Korotkiy (VU), *From ontology-enabled services to service-enabled ontologies (making ontologies work in e-science with ONTO-SOA)*
- 2009-13** Steven de Jong (UM), *Fairness in Multi-Agent Systems*
- 2009-12** Peter Massuthe (TUE, Humboldt-Universitaet zu Berlin), *Operating Guidelines for Services*
- 2009-11** Alexander Boer (UVA), *Legal Theory, Sources of Law & the Semantic Web*
- 2009-10** Jan Wielemaker (UVA), *Logic programming for knowledge-intensive interactive applications*
- 2009-09** Benjamin Kanagwa (RUN), *Design, Discovery and Construction of Service-oriented Systems*
- 2009-08** Volker Nannen (VU), *Evolutionary Agent-Based Policy Analysis in Dynamic Environments*
- 2009-07** Ronald Poppe (UT), *Discriminative Vision-Based Recovery and Recognition of Human Motion*
- 2009-06** Muhammad Subianto (UU), *Understanding Classification*
- 2009-05** Sietse Overbeek (RUN), *Bridging Supply and Demand for Knowledge Intensive Tasks - Based on Knowledge, Cognition, and Quality*
- 2009-04** Josephine Nabukenya (RUN), *Improving the Quality of Organisational Policy Making using Collaborative Engineering*
- 2009-03** Hans Stol (UvT), *A Framework for Evidence-based Policy Making Using IT*
- 2009-02** Willem Robert van Hage (VU), *Evaluating Ontology-Alignment Techniques*
- 2009-01** Rasa Jurgelenaite (RUN), *Symmetric Causal Independence Models*
- 2008-35** Ben Torben Nielsen (UvT), *Dendritic morphologies: function shapes structure*
- 2008-34** Jeroen de Knijf (UU), *Studies in Frequent Tree Mining*
- 2008-33** Frank Terpstra (UVA), *Scientific Workflow Design; theoretical and practical issues*
- 2008-32** Trung H. Bui (UT), *Toward Affective Dialogue Management using Partially Observable Markov Decision Processes*
- 2008-31** Loes Braun (UM), *Pro-Active Medical Information Retrieval*
- 2008-30** Wouter van Atteveldt (VU), *Semantic Network Analysis: Techniques for Extracting, Representing and Querying Media Content*
- 2008-29** Dennis Reidsma (UT), *Annotations and Subjective Machines - Of Annotators, Embodied Agents, Users, and Other Humans*
- 2008-28** Ildiko Flesch (RUN), *On the Use of Independence Relations in Bayesian Networks*
- 2008-27** Hubert Vogten (OU), *Design and Implementation Strategies for IMS Learning Design*
- 2008-26** Marijn Huijbregts (UT), *Segmentation, Diarization and Speech Transcription: Surprise Data Unraveled*
- 2008-25** Geert Jonker (UU), *Efficient and Equitable Exchange in Air Traffic Management Plan Repair using Spender-signed Currency*
- 2008-24** Zharko Aleksovski (VU), *Using background knowledge in ontology matching*
- 2008-23** Stefan Visscher (UU), *Bayesian network models for the management of ventilator-associated pneumonia*
- 2008-22** Henk Koning (UU), *Communication of IT-Architecture*
- 2008-21** Krisztian Balog (UVA), *People Search in the Enterprise*
- 2008-20** Rex Arendsen (UVA), *Geen bericht, goed bericht. Een onderzoek naar de effecten van de introductie van elektronisch berichtenverkeer met de overheid op de administratieve lasten van bedrijven.*
- 2008-19** Henning Rode (UT), *From Document to Entity Retrieval: Improving Precision and Performance of Focused Text Search*
- 2008-18** Guido de Croon (UM), *Adaptive Active Vision*
- 2008-17** Martin Op 't Land (TUD), *Applying Architecture and Ontology to the Splitting and Allying of Enterprises*
- 2008-16** Henriette van Vugt (VU), *Embodied agents from a user's perspective*
- 2008-15** Martijn van Otterlo (UT), *The Logic of Adaptive Behavior: Knowledge Representation and Algorithms for the Markov Decision Process Framework in First-Order Domains.*
- 2008-14** Arthur van Bunningen (UT), *Context-Aware Querying; Better Answers with Less Effort*
- 2008-13** Caterina Carraciolo (UVA), *Topic Driven Access to Scientific Handbooks*
- 2008-12** Jozsef Farkas (RUN), *A Semiotically Oriented Cognitive Model of Knowledge Representation*
- 2008-11** Vera Kartseva (VU), *Designing Controls for Network Organizations: A Value-Based Approach*
- 2008-10** Wauter Bosma (UT), *Discourse oriented summarization*
- 2008-09** Christof van Nimwegen (UU), *The paradox of the guided user: assistance can be counter-effective*
- 2008-08** Janneke Bolt (UU), *Bayesian Networks: Aspects of Approximate Inference*
- 2008-07** Peter van Rosmalen (OU), *Supporting the tutor in the design and support of adaptive e-learning*
- 2008-06** Arjen Hommersom (RUN), *On the Application of Formal Methods to Clinical Guidelines, an Artificial Intelligence Perspective*
- 2008-05** Bela Mutschler (UT), *Modeling and simulating causal dependencies on process-aware information systems from a cost perspective*
- 2008-04** Ander de Keijzer (UT), *Management of Uncertain Data - towards unattended integration*
- 2008-03** Vera Hollink (UVA), *Optimizing hierarchical menus: a usage-based approach*

- 2008-02** Alexei Sharpanskykh (VU), *On Computer-Aided Methods for Modeling and Analysis of Organizations*
- 2008-01** Katalin Boer-Sorbán (EUR), *Agent-Based Simulation of Financial Markets: A modular, continuous-time approach*
- 2007-25** Joost Schalken (VU), *Empirical Investigations in Software Process Improvement*
- 2007-24** Georgina Ramírez Camps (CWI), *Structural Features in XML Retrieval*
- 2007-23** Peter Barna (TUE), *Specification of Application Logic in Web Information Systems*
- 2007-22** Zlatko Zlatev (UT), *Goal-oriented design of value and process models from patterns*
- 2007-21** Karianne Vermaas (UU), *Fast diffusion and broadening use: A research on residential adoption and usage of broadband internet in the Netherlands between 2001 and 2005*
- 2007-20** Slinger Jansen (UU), *Customer Configuration Updating in a Software Supply Network*
- 2007-19** David Levy (UM), *Intimate relationships with artificial partners*
- 2007-18** Bart Orriens (UvT), *On the development an management of adaptive business collaborations*
- 2007-17** Theodore Charitos (UU), *Reasoning with Dynamic Networks in Practice*
- 2007-16** Davide Grossi (UU), *Designing Invisible Handcuffs. Formal investigations in Institutions and Organizations for Multi-agent Systems*
- 2007-15** Joyca Lacroix (UM), *NIM: a Situated Computational Memory Model*
- 2007-14** Niek Bergboer (UM), *Context-Based Image Analysis*
- 2007-13** Rutger Rienks (UT), *Meetings in Smart Environments; Implications of Progressing Technology*
- 2007-12** Marcel van Gerven (RUN), *Bayesian Networks for Clinical Decision Support: A Rational Approach to Dynamic Decision-Making under Uncertainty*
- 2007-11** Natalia Stash (TUE), *Incorporating Cognitive/Learning Styles in a General-Purpose Adaptive Hypermedia System*
- 2007-10** Huib Alderweld (UU), *Autonomy vs. Conformity: an Institutional Perspective on Norms and Protocols*
- 2007-09** David Mobach (VU), *Agent-Based Mediated Service Negotiation*
- 2007-08** Mark Hoogendoorn (VU), *Modeling of Change in Multi-Agent Organizations*
- 2007-07** Natasa Jovanovic (UT), *To Whom It May Concern - Addressee Identification in Face-to-Face Meetings*
- 2007-06** Gilad Mishne (UVA), *Applied Text Analytics for Blogs*
- 2007-05** Bart Schermer (UL), *Software Agents, Surveillance, and the Right to Privacy: a Legislative Framework for Agent-enabled Surveillance*
- 2007-04** Jurriaan van Diggelen (UU), *Achieving Semantic Interoperability in Multi-agent Systems: a dialogue-based approach*
- 2007-03** Peter Mika (VU), *Social Networks and the Semantic Web*
- 2007-02** Wouter Teepe (RUG), *Reconciling Information Exchange and Confidentiality: A Formal Approach*
- 2007-01** Kees Leune (UvT), *Access Control and Service-Oriented Architectures*
- 2006-28** Borkur Sigurbjornsson (UVA), *Focused Information Access using XML Element Retrieval*
- 2006-27** Stefano Bocconi (CWI), *Vox Populi: generating video documentaries from semantically annotated media repositories*
- 2006-26** Vojkan Mihajlovic (UT), *Score Region Algebra: A Flexible Framework for Structured Information Retrieval*
- 2006-25** Madalina Drugan (UU), *Conditional likelihood MDL and Evolutionary MCMC*
- 2006-24** Laura Hollink (VU), *Semantic Annotation for Retrieval of Visual Resources*
- 2006-23** Ion Juvina (UU), *Development of Cognitive Model for Navigating on the Web*
- 2006-22** Paul de Vrieze (RUN), *Fundamentals of Adaptive Personalisation*
- 2006-21** Bas van Gils (RUN), *Aptness on the Web*
- 2006-20** Marina Velikova (UvT), *Monotone models for prediction in data mining*
- 2006-19** Birna van Riemsdijk (UU), *Cognitive Agent Programming: A Semantic Approach*
- 2006-18** Valentin Zhizhkun (UVA), *Graph transformation for Natural Language Processing*
- 2006-17** Stacey Nagata (UU), *User Assistance for Multitasking with Interruptions on a Mobile Device*
- 2006-16** Carsten Riggelsen (UU), *Approximation Methods for Efficient Learning of Bayesian Networks*
- 2006-15** Rainer Malik (UU), *CONAN: Text Mining in the Biomedical Domain*
- 2006-14** Johan Hoorn (VU), *Software Requirements: Update, Upgrade, Redesign - towards a Theory of Requirements Change*
- 2006-13** Henk-Jan Lebbink (UU), *Dialogue and Decision Games for Information Exchanging Agents*
- 2006-12** Bert Bongers (VU), *Interaction - Towards an e-cology of people, our technological environment, and the arts*
- 2006-11** Joeri van Ruth (UT), *Flattening Queries over Nested Data Types*
- 2006-10** Ronny Siebes (VU), *Semantic Routing in Peer-to-Peer Systems*
- 2006-09** Mohamed Wahdan (UM), *Automatic Formulation of the Auditor's Opinion*
- 2006-08** Elco Herder (UT), *Forward, Back and Home Again - Analyzing User Behavior on the Web*
- 2006-07** Marko Smiljanic (UT), *XML schema matching - balancing efficiency and effectiveness by means of clustering*
- 2006-06** Ziv Baida (VU), *Software-aided Service Bundling - Intelligent Methods & Tools for Graphical Service Modeling*
- 2006-05** Cees Pierik (UU), *Validation Techniques for Object-Oriented Proof Outlines*
- 2006-04** Marta Sabou (VU), *Building Web Service Ontologies*
- 2006-03** Noor Christoph (UVA), *The role of metacognitive skills in learning to solve problems*
- 2006-02** Cristina Chisalita (VU), *Contextual issues in the design and use of information technology in organizations*
- 2006-01** Samuil Angelov (TUE), *Foundations of B2B Electronic Contracting*
- 2005-21** Wijnand Derks (UT), *Improving Concurrency and Recovery in Database Systems by Exploiting Application Semantics*
- 2005-20** Cristina Coteanu (UL), *Cyber Consumer Law, State of the Art and Perspectives*
- 2005-19** Michel van Dartel (UM), *Situated Representation*
- 2005-18** Danielle Sent (UU), *Test-selection strategies for probabilistic networks*
- 2005-17** Boris Shishkov (TUD), *Software Specification Based on Re-usable Business Components*
- 2005-16** Joris Graaumann (UU), *Usability of XML Query Languages*
- 2005-15** Tibor Bosse (VU), *Analysis of the Dynamics of Cognitive Processes*
- 2005-14** Borys Omelayenko (VU), *Web-Service configuration on the Semantic Web; Exploring how semantics meets pragmatics*
- 2005-13** Fred Hamburg (UL), *Een Computermodel voor het Ondersteunen van Euthanasiebeslissingen*
- 2005-12** Csaba Boer (EUR), *Distributed Simulation in Industry*
- 2005-11** Elth Ogston (VU), *Agent Based Matchmaking and Clustering - A Decentralized Approach to Search*
- 2005-10** Anders Bouwer (UVA), *Explaining Behaviour: Using Qualitative Simulation in Interactive Learning Environments*
- 2005-09** Jeen Broekstra (VU), *Storage, Querying and Inferencing for Semantic Web Languages*
- 2005-08** Richard Vdovjak (TUE), *A Model-driven Approach for Building Distributed Ontology-based Web Applications*
- 2005-07** Flavius Frasinca (TUE), *Hypermedia Presentation Generation for Semantic Web Information Systems*
- 2005-06** Pieter Spronck (UM), *Adaptive Game AI*
- 2005-05** Gabriel Infante-Lopez (UVA), *Two-Level Probabilistic Grammars for Natural Language Parsing*
- 2005-04** Nirvana Meratnia (UT), *Towards Database Support for Moving Object data*
- 2005-03** Franc Grootjen (RUN), *A Pragmatic Approach to the Conceptualisation of Language*
- 2005-02** Erik van der Werf (UM), *AI techniques for the game of Go*
- 2005-01** Floor Verdenius (UVA), *Methodological Aspects of Designing Induction-Based Applications*

- 2004-20** Madelon Evers (Nyenrode), *Learning from Design: facilitating multidisciplinary design teams*
- 2004-19** Thijs Westerveld (UT), *Using generative probabilistic models for multimedia retrieval*
- 2004-18** Vania Bessa Machado (UvA), *Supporting the Construction of Qualitative Knowledge Models*
- 2004-17** Mark Winands (UM), *Informed Search in Complex Games*
- 2004-16** Federico Divina (VU), *Hybrid Genetic Relational Search for Inductive Learning*
- 2004-15** Arno Knobbe (UU), *Multi-Relational Data Mining*
- 2004-14** Paul Harrenstein (UU), *Logic in Conflict. Logical Explorations in Strategic Equilibrium*
- 2004-13** Wojciech Jamroga (UT), *Using Multiple Models of Reality: On Agents who Know how to Play*
- 2004-12** The Duy Bui (UT), *Creating emotions and facial expressions for embodied agents*
- 2004-11** Michel Klein (VU), *Change Management for Distributed Ontologies*
- 2004-10** Suzanne Kabel (UVA), *Knowledge-rich indexing of learning-objects*
- 2004-09** Martin Caminada (VU), *For the Sake of the Argument: explorations into argument-based reasoning*
- 2004-08** Joop Verbeek (UM), *Politie en de Nieuwe Internationale Informatiemarkt, Grensregionale politieke gegevensuitwisseling en digitale expertise*
- 2004-07** Elise Boltjes (UM), *Voorbeeldig onderwijs: voorbeeldgestuurd onderwijs, een opstap naar abstract denken, vooral voor metsjes*
- 2004-06** Bart-Jan Hommes (TUD), *The Evaluation of Business Process Modeling Techniques*
- 2004-05** Viara Popova (EUR), *Knowledge discovery and monotonicity*
- 2004-04** Chris van Aart (UVA), *Organizational Principles for Multi-Agent Architectures*
- 2004-03** Perry Groot (VU), *A Theoretical and Empirical Analysis of Approximation in Symbolic Problem Solving*
- 2004-02** Lai Xu (UvT), *Monitoring Multi-party Contracts for E-business*
- 2004-01** Virginia Dignum (UU), *A Model for Organizational Interaction: Based on Agents, Founded in Logic*
- 2003-18** Levente Kocsis (UM), *Learning Search Decisions*
- 2003-17** David Jansen (UT), *Extensions of Statecharts with Probability, Time, and Stochastic Timing*
- 2003-16** Menzo Windhouwer (CWI), *Feature Grammar Systems - Incremental Maintenance of Indexes to Digital Media Warehouses*
- 2003-15** Mathijs de Weerdt (TUD), *Plan Merging in Multi-Agent Systems*
- 2003-14** Stijn Hoppenbrouwers (KUN), *Freezing Language: Conceptualisation Processes across ICT-Supported Organisations*
- 2003-13** Jeroen Donkers (UM), *Nosce Hostem - Searching with Opponent Models*
- 2003-12** Roeland Ordelman (UT), *Dutch speech recognition in multimedia information retrieval*
- 2003-11** Simon Keizer (UT), *Reasoning under Uncertainty in Natural Language Dialogue using Bayesian Networks*
- 2003-10** Andreas Lincke (UvT), *Electronic Business Negotiation: Some experimental studies on the interaction between medium, innovation context and culture*
- 2003-09** Rens Kortmann (UM), *The resolution of visually guided behaviour*
- 2003-08** Yongping Ran (UM), *Repair Based Scheduling*
- 2003-07** Machiel Jansen (UvA), *Formal Explorations of Knowledge Intensive Tasks*
- 2003-06** Boris van Schooten (UT), *Development and specification of virtual environments*
- 2003-05** Jos Lehmann (UVA), *Causation in Artificial Intelligence and Law - A modelling approach*
- 2003-04** Milan Petkovic (UT), *Content-Based Video Retrieval Supported by Database Technology*
- 2003-03** Martijn Schuemie (TUD), *Human-Computer Interaction and Presence in Virtual Reality Exposure Therapy*
- 2003-02** Jan Broersen (VU), *Modal Action Logics for Reasoning About Reactive Systems*
- 2003-01** Heiner Stuckenschmidt (VU), *Ontology-Based Information Sharing in Weakly Structured Environments*
- 2002-17** Stefan Manegold (UVA), *Understanding, Modeling, and Improving Main-Memory Database Performance*
- 2002-16** Pieter van Langen (VU), *The Anatomy of Design: Foundations, Models and Applications*
- 2002-15** Rik Eshuis (UT), *Semantics and Verification of UML Activity Diagrams for Workflow Modelling*
- 2002-14** Wieke de Vries (UU), *Agent Interaction: Abstract Approaches to Modelling, Programming and Verifying Multi-Agent Systems*
- 2002-13** Hongjing Wu (TUE), *A Reference Architecture for Adaptive Hypermedia Applications*
- 2002-12** Albrecht Schmidt (Uva), *Processing XML in Database Systems*
- 2002-11** Wouter C.A. Wijngaards (VU), *Agent Based Modelling of Dynamics: Biological and Organisational Applications*
- 2002-10** Brian Sheppard (UM), *Towards Perfect Play of Scrabble*
- 2002-09** Willem-Jan van den Heuvel (KUB), *Integrating Modern Business Applications with Objectified Legacy Systems*
- 2002-08** Jaap Gordijn (VU), *Value Based Requirements Engineering: Exploring Innovative E-Commerce Ideas*
- 2002-07** Peter Boncz (CWI), *Monet: A Next-Generation DBMS Kernel For Query-Intensive Applications*
- 2002-06** Laurens Mommers (UL), *Applied legal epistemology; Building a knowledge-based ontology of the legal domain*
- 2002-05** Radu Serban (VU), *The Private Cyberspace Modeling Electronic Environments inhabited by Privacy-concerned Agents*
- 2002-04** Juan Roberto Castelo Valdeuza (UU), *The Discrete Acyclic Digraph Markov Model in Data Mining*
- 2002-03** Henk Ernst Blok (UT), *Database Optimization Aspects for Information Retrieval*
- 2002-02** Roelof van Zwol (UT), *Modelling and searching web-based document collections*
- 2002-01** Nico Lassing (VU), *Architecture-Level Modifiability Analysis*
- 2001-11** Tom M. van Engers (VUA), *Knowledge Management: The Role of Mental Models in Business Systems Design*
- 2001-10** Maarten Sierhuis (UvA), *Modeling and Simulating Work Practice BRAHMS: a multiagent modeling and simulation language for work practice analysis and design*
- 2001-09** Pieter Jan 't Hoen (RUL), *Towards Distributed Development of Large Object-Oriented Models, Views of Packages as Classes*
- 2001-08** Pascal van Eck (VU), *A Compositional Semantic Structure for Multi-Agent Systems Dynamics.*
- 2001-07** Bastiaan Schonhage (VU), *Divva: Architectural Perspectives on Information Visualization*
- 2001-06** Martijn van Welie (VU), *Task-based User Interface Design*
- 2001-05** Jacco van Ossenbruggen (VU), *Processing Structured Hypermedia: A Matter of Style*
- 2001-04** Evgueni Smirnov (UM), *Conjunctive and Disjunctive Version Spaces with Instance-Based Boundary Sets*
- 2001-03** Maarten van Someren (UvA), *Learning as problem solving*
- 2001-02** Koen Hindriks (UU), *Agent Programming Languages: Programming with Mental Models*
- 2001-01** Silja Renooij (UU), *Qualitative Approaches to Quantifying Probabilistic Networks*
- 2000-11** Jonas Karlsson (CWI), *Scalable Distributed Data Structures for Database Management*
- 2000-10** Niels Nes (CWI), *Image Database Management System Design Considerations, Algorithms and Architecture*
- 2000-09** Florian Waas (CWI), *Principles of Probabilistic Query Optimization*
- 2000-08** Veerle Coupé (EUR), *Sensitivity Analysis of Decision-Theoretic Networks*
- 2000-07** Niels Peek (UU), *Decision-theoretic Planning of Clinical Patient Management*
- 2000-06** Rogier van Eijk (UU), *Programming Languages for Agent Communication*
- 2000-05** Ruud van der Pol (UM), *Knowledge-based Query Formulation in Information Retrieval.*
- 2000-04** Geert de Haan (VU), *ETAG, A Formal Model of Competence Knowledge for User Interface Design*
- 2000-03** Carolien M.T. Metselaar (UVA), *Sociaal-organisatorische gevolgen van kennistechnologie; een procesbenadering en actorperspectief.*

- 2000-02** Koen Holtman (TUE), *Prototyping of CMS Storage Management*
- 2000-01** Frank Niessink (VU), *Perspectives on Improving Software Maintenance*
- 1999-08** Jacques H.J. Lenting (UM), *Informed Gambling: Conception and Analysis of a Multi-Agent Mechanism for Discrete Reallocation.*
- 1999-07** David Spelt (UT), *Verification support for object database design*
- 1999-06** Niek J.E. Wijngaards (VU), *Re-design of compositional systems*
- 1999-05** Aldo de Moor (KUB), *Empowering Communities: A Method for the Legitimate User-Driven Specification of Network Information Systems*
- 1999-04** Jacques Penders (UM), *The practical Art of Moving Physical Objects*
- 1999-03** Don Beal (UM), *The Nature of Minimax Search*
- 1999-02** Rob Potharst (EUR), *Classification using decision trees and neural nets*
- 1999-01** Mark Sloof (VU), *Physiology of Quality Change Modelling; Automated modelling of Quality Change of Agricultural Products*
- 1998-05** E.W.Oskamp (RUL), *Computerondersteuning bij Straftoemeting*
- 1998-04** Dennis Breuker (UM), *Memory versus Search in Games*
- 1998-03** Ans Steuten (TUD), *A Contribution to the Linguistic Analysis of Business Conversations within the Language/Action Perspective*
- 1998-02** Floris Wiesman (UM), *Information Retrieval by Graphically Browsing Meta-Information*
- 1998-01** Johan van den Akker (CWI), *DEGAS - An Active, Temporal Database of Autonomous Objects*




TU Delft

ISBN 978-94-028-1462-0



9 789402 814620