

# Report on the WSDM 2020 Workshop on State-based User Modelling (SUM'20)

Sahan Bulathwela  
University College London,  
London, UK  
*m.bulathwela@ucl.ac.uk*

María Pérez-Ortiz  
University College London,  
London, UK

Rishabh Mehrotra  
Spotify Research,  
London, UK

Davor Orlic  
Knowledge 4 All Foundation,  
London, UK

Colin de la Higuera  
Nantes University,  
Nantes, France

John Shawe-Taylor  
University College London,  
London, UK

Emine Yilmaz  
University College London,  
London, UK

## Abstract

The SUM'20 workshop was held at the 13th ACM International WSDM Conference on Web Search and Data Mining (WSDM 2020) in Houston, Texas. The purpose of the workshop was to stimulate the research community to explore open challenges in building systems that can capture the user's state, context and goals, as well as effectively use these for leveraging intelligent user-centric systems in a wide range of applications. The workshop incorporated different plenary sessions and contributed talks. The workshop website and proceedings are available at <https://www.k4all.org/event/wsdmsum20>.

## 1 Introduction

Recently, there has been a significant interest in applications that require intelligent algorithms that can capture, keep track and utilise the user's state and context. Capturing and effectively utilising the user's states, context and goals is becoming a timely challenge for successfully leveraging intelligent user-centric systems in different machine learning and information retrieval applications [Trippas et al., 2019]. Examples of such systems are conversational agents [Vtyurina et al., 2017], intelligent assistants [Belkadi et al., 2020, Yan and Selker, 2000], educational [Piech

---

et al., 2015] and contextual information retrieval systems [Syed and Collins-Thompson, 2017], recommender/match-making systems [Minka et al., 2018, Adomavicius and Tuzhilin, 2011, Unger et al., 2016, Christakopoulou et al., 2016, Yan et al., 2016], and advertising systems [Graepel et al., 2010], all of which rely on identifying the user’s state/context in order to model the user’s preferences or features, provide the most relevant information and assist users in achieving their goals.

For example, due to their interactive nature, search and recommender systems are notoriously hard to develop and evaluate, since they involve a multi-step decision-making process, where a stream of interactions occurs between the user and the system. Traditionally, to make the problem tractable, the interactions are often viewed as independent. However, these systems could greatly benefit from considering and leveraging dynamic user states [Zhao et al., 2019], e.g. such as spatial and temporal user information [Costa and Dolog, 2019].

In recent years, several interactive machine learning techniques (including bandit based and reinforcement learning based techniques) have been applied for interaction tasks, as these approaches allow systems to learn and continuously change the strategy based on user feedback and other type of rewards. Even though many such systems significantly benefit from having a representation of the user’s state (a large part of the related work has focused on the representational view of the user), there has been very limited work towards building intelligent learning mechanisms that can be used to identify, represent and update the state of the user (e.g. sequence-aware and latent state-aware systems). For example, one of the great current open challenges is inferring the implicit state of the user in a real-time scenario, which requires modeling complex, partially observable and dynamic factors, which are crucial for many different applications (education, health, advertising, affective computing, etc.).

Because of this, devising information retrieval systems that can keep track of the user’s state and make recommendations based on this has been listed as one of the grand challenges of information retrieval that needs to be tackled during the next few years [Culpepper et al., 2018]. In this context, we believe it was timely to organize a workshop that revisited the problem of designing and evaluating state-aware user-centric systems and made sure the community, spanning academic and industrial backgrounds, worked together to tackle these challenges.

## 2 Scope and Objectives

The workshop aim was to facilitate a forum for discussing and leveraging intelligent methods for representing and modelling a user’s state and its transition in an effective manner. The workshop focused on several open questions:

- (i) **State-aware user representations:** How to design accurate state-aware dynamic representations of the user and embed that information in predictive models?
- (ii) **Task understanding and supporting user tasks:** How to identify the specific sub-task a user is trying to accomplish and design task-aware intelligent systems (e.g. recommender systems) that assist users in each step of their goals?
- (iii) **State-aware evaluation:** How to develop metrics and evaluation techniques that model and understand user states, so as to provide more sensitive and enhanced evaluation of these

---

user-centric systems?

- (iv) **Human in the loop:** How to include the human in the learning "loop" of these systems in an interactive fashion, making feedback an integral part of the modelling?
- (v) **User-aware systems:** How to capture latent user state information and incorporate it to achieve better human-machine intelligence?
- (vi) **Cognitive/contextual user understanding:** How to understand from a cognitive point of view how actions taken by users can influence their state/context?
- (vii) **State-aware machine learning:** How do the different machine learning algorithms that are able to identify and leverage user states compare for different applications?

To answer these questions, it is imperative for research communities from user modelling, cognitive science, intelligent tutoring systems, information retrieval, human computer interaction, psychology and other diverse fields to bring forward their inputs, shaping the future of many personalised machine learning applications and providing inspiration to building more powerful user modelling approaches.

### 3 Related Workshops

There have been several workshops that can be considered of relevance to the SUM'20 research themes. These comprise: Task Intelligence Workshop<sup>1</sup> (WSDM'19), LearnIR: Workshop on Learning from User Interactions<sup>2</sup> (WSDM'18), several workshops at the Conf. on User Modeling, Adaptation and Personalization<sup>3</sup> (UMAP'19), Where is the Human? Bridging the Gap Between AI and HCI Workshop<sup>4</sup> (CHI'19), Interactive Machine Learning and Semantic Information Retrieval<sup>5</sup> (ICML'17), The TREC Tasks Track 2016<sup>6</sup> (TREC'16), The Future of Interactive Machine Learning<sup>7</sup> (NIPS'16), ML for Education<sup>8</sup> (NIPS'16), Human Centered Machine Learning<sup>9</sup> (CHI'16), Intelligent Techniques for Web Personalization and Recommendation<sup>10</sup> (AAAI'14) and Context-Aware Recommender Systems<sup>11</sup> (RecSys'20), among others. These very recent workshops and discussion forums demonstrate that this topic is recognised as a pressing one by many leading conferences and research communities related to user modelling.

---

<sup>1</sup><https://dl.acm.org/citation.cfm?doid=3289600.3291374>

<sup>2</sup><http://rishabhmehtrotra.com/WSDM18-learnIR.pdf>

<sup>3</sup><http://www.cyprusconferences.org/umap2019/pages/workshops.html>

<sup>4</sup><https://veale.github.io/ai-hci-workshop/>

<sup>5</sup><https://icml.cc/Conferences/2017/Schedule?showEvent=6>

<sup>6</sup><http://trec.nist.gov/pubs/trec25/papers/Overview-T.pdf>

<sup>7</sup><https://neurips.cc/Conferences/2016/Schedule?showEvent=6224>

<sup>8</sup><https://neurips.cc/Conferences/2016/Schedule?showEvent=6241>

<sup>9</sup><http://hcml2016.goldsmithsdigital.com/program/>

<sup>10</sup><https://www.aaai.org/Press/Reports/Workshops/ws-13-11.php>

<sup>11</sup><https://cars-workshop.com/>

---

## 4 Workshop Contributions

Rather than focusing on different aspects of state-based user modelling in isolation (such as user interactions, task taxonomy spaces, etc.), SUM'20 encouraged tabling the challenges, findings and solutions relating to the multitude of different aspects of state-based user modelling in one forum where all attendees are active participants. Given the focus areas of the workshop, participants from academic institutions, industrial research labs as well as start-ups could benefit from participating in the workshop. More than 70 conference attendees registered to participate in SUM'20 upon announcing the workshop. The workshop featured a mix of invited keynotes and contributed talks by leading researchers from diverse backgrounds working in this theme. The primary objective of having a series of talks from researchers working in different aspects of state-based user modelling is to spark cross-domain collaborations between attendees to address the open problems in the field. Both the contributed papers and the slides of the keynotes are available at the workshop website<sup>12</sup>.

### 4.1 Invited Keynotes

The workshop included four keynote talks that covered different aspects of state-based user modelling from both academic and industrial perspectives. The SUM'20 keynote talks were delivered by the following speakers:

- Mounia Lalmas (Director of Research at Spotify)
- Eugene Agichtein (Winship Associate Professor at Emory University, USA)
- Ahmed Awadallah (Principal Research Manager at Microsoft Research)
- Emine Yilmaz (Professor at University College London and Amazon Scholar at Amazon)

**Engagement, metrics and recommenders, *Mounia Lalmas*** The first keynote of the workshop focused on different steps of developing state-aware recommendation system that optimise for the right user engagement metrics. The talk consisted of various relevant works and Lalmas' personal thoughts on how to achieve this goal. One of the main challenges in this front is to leverage knowledge about the online interaction of users to understand what engages them short-term and more importantly long-term [Hong and Lalmas, 2019]. A common way in which engagement is measured and understood is through the definition and development of metrics of user satisfaction, which can act as proxy of short-term user engagement. In the context of recommender systems, developing a better understanding of how users interact (implicit signals) with the system during their online session is important for developing metrics of user satisfaction [Lehmann et al., 2012]. The talk focused on how to develop metrics of user engagement, which recommender systems can optimize for. An important message was that, for recommender systems to work both in the short and the long-term, it is important to consider the heterogeneity of both user and content to formalise the notion of engagement, and in turn design the appropriate metrics to capture these and optimize for Anderson et al. [2020]. One way to achieve this is to follow these four steps:

---

<sup>12</sup><https://www.k4all.org/event/wsdmsum20>

---

1) Understanding intents; 2) optimizing for the right metric; 3) acting on segmentation; and 4) thinking about diversity.

**Measuring and improving user satisfaction with voice-based conversational systems, Eugene Agichtein** In this keynote, Agichtein presented the critical role understanding user satisfaction and user state plays in succeeding as a conversational agent. As spoken conversational assistants operate in increasingly complex domains, predicting user satisfaction in conversational systems has become critical. In particular, online satisfaction prediction (i.e., predicting satisfaction of the user with the system after each turn) could be used as a new proxy for implicit user feedback, and offers promising opportunities to create more responsive and effective conversational agents, which adapt to the user engagement with the agent [Choi et al., 2019]. Measuring immediate changes in user satisfaction also enables for natural user studies without explicit disruptions to user experience. Agichtein discussed his team progress on conversational satisfaction modeling, and applications to automatically evaluating the effects new agent features, e.g., conversation topic suggestion [Ahmadvand et al., 2019], design choices, e.g., new recommendation phrasing [Ahmadvand et al., 2020], and subtle VUI features, e.g., prosody modulation for agent responses [Choi and Agichtein].

**Learning from user interactions in productivity applications, Ahmed Awadallah** Information workers spend significant amounts of their time managing personal information such as communications, documents and tasks. A recent study reported that the average information worker spends an estimated 28 percent of the work week managing e-mail and nearly 20 percent looking for internal information or tracking down colleagues who can help with specific tasks. This highlights the potential value intelligent technologies could provide by realizing faster and more effective collaboration and information sharing. The use of artificial intelligence in online services has been steadily increasing. In this keynote, Awadallah presented relevant advancements made in the topic of office productivity improvement through his research work. Their aim is to create new intelligent experiences that can help us be more productive. Awadallah discussed three main directions, namely: 1) user interaction in communications (email intelligence) [Wang et al., 2019], 2) document recommendation [Xu et al., 2020] and 3) task management [White et al., 2019] where the use of state-aware intelligent systems is explored.

**Incorporating user state in task based information retrieval and recommendation of educational resources, Emine Yilmaz** In the final keynote of the SUM'20 workshop, Yilmaz presents two applications in which incorporating user state could be highly important: 1) task based information retrieval [Mehrotra and Yilmaz, 2017a] and 2) educational recommender systems [Bulathwela et al., 2020a]. Accurate representation of user state plays a critical role in understanding user needs and providing them with accurate information at the correct time. In the first part of the talk, Yilmaz described her own work on incorporating representations of user state to task based search engines and showed how one can use these representations to predict user needs [Mehrotra and Yilmaz, 2017b]. Yilmaz proceeded to cover recommendation of educational resources, describing a way of representing user state (skills) and describing methods that can be used to match learners with educational resources based on their skill level [Bulathwela et al., 2020b]. She concluded her talk by demonstrating X5Learn platform ([www.x5learn.org](http://www.x5learn.org)),

---

an intelligent learning platform where the state-based educational recommender is incorporated [Bulathwela et al., 2020c].

## 4.2 Paper Contributions

We solicited submissions as full research papers that addressed issues related to the topic of state-based user modelling. The workshop contributions were reviewed by an international panel of programme committee members that brought in expertise from both academic and industrial fronts. Each contribution was reviewed by at least three reviewers on a double-blind basis. The programme committee consisted of (sorted alphabetically according to the last name) Sahan Bulathwela (University College London), Colin De La Higuera (University of Nantes), Shashank Gupta (Flipkart), Shangsong Liang (Sun Yat-sen University), Aldo Lipani (University College London), Marwa Mahmoud (University of Cambridge), Jarana Manotumruksa (University College London), Yin Cheng Ng (Man AHL), María Pérez-Ortiz (University College London), Javier Sánchez-Monedero (Cardiff University), Manisha Verma (Verizon Media) and Kyle Williams (Microsoft). The workshop accepted three contributions from a total number of five submissions (60% acceptance rate).

They are as follows:

- Daisuke Moriwaki, Komei Fujita, Shota Yasui, and Takahiro Hoshino, **Fatigue-Aware Ad Creative Selection** [Moriwaki et al., 2020]. Abstract of the paper: *"In online display advertising, selecting the most effective ad creative (ad image) for each impression is a crucial task for DSPs (Demand-Side Platforms) to fulfill their goals (click-through rate, number of conversions, revenue, and brand improvement). As widely recognized in the marketing literature, the effect of ad creative changes with the number of repetitive ad exposures. In this study, we propose an efficient and easy-to-implement ad creative selection algorithm that explicitly considers user's psychological status when selecting ad creatives. The proposed system was deployed in a real-world production environment and tested against the baseline algorithms. The results show superiority of the proposed algorithm"*.
- Behrooz Omidvar-Tehrani, Sruthi Viswanathan, Frederic Roulland, and Jean-Michel Renders, **SAGE: Interactive State-aware Point-of-Interest Recommendation** [Omidvar-Tehrani et al., 2020]. Abstract of the paper: *"Point-of-Interest (POI) recommendation is surfacing in many location-based services. User models are employed in these services to leverage historical check-ins and social links, and enable personalized and socialized POI recommendations. However these models often lack interactivity (incorporating user interactions) and state-awareness. This deficiency aggravates in cold start situations, where nearly no user information (historical check-ins and social graph) is available to generate effective recommendations. In this paper, we propose SAGE, an interactive state-aware POI recommendation system which tackles the aforementioned challenges by exploiting look-alike groups mined in public POI datasets, such as Foursquare and Yelp. SAGE reformulates the problem of POI recommendation as recommending explainable look-alike groups (and their POIs) which are in line with user's intent. SAGE frames the task of POI recommendation as an exploratory process where users interact with the system, and their interactions impact*

---

*the way look-alike groups are picked out. Moreover, SAGE defines and employs mindsets which capture the actual state of the user and enforce the semantics of POI interestingness. Our experiments show that SAGE is an effective approach to capture interactivity and contextuality for recommending relevant look-alike groups and their POIs which are oriented towards the user’s mindset.”*

- Alfonso White and Daniela M. Romano, **Scalable Psychological Momentum Estimation in E-sports** [White and Romano, 2020]. Abstract of the paper: *”The world of competitive Esports and video gaming has seen and continues to experience steady growth in popularity and complexity. Correspondingly, more research on the topic is being published, ranging from social network analyses to the benchmarking of advanced artificial intelligence systems in playing against humans. In this paper, we present ongoing work on an intelligent agent recommendation engine that suggests actions to players in order to maximise success and enjoyment, both in the space of in-game choices, as well as decisions made around play session timing in the broader context. By leveraging temporal data and appropriate models, we show that a learned representation of player psychological momentum, and of tilt, can be used, in combination with player expertise, to achieve state-of-the-art performance in pre-and post-draft win prediction. Our progress toward fulfilling the potential for deriving optimal recommendations is documented.”*

## 5 Discussion and Conclusions

The paper contributions to the workshop relate to how the user state can be learned and leveraged to improve intelligent systems devised for open problems in three different fields, namely 1) online advertising, 2) recommendation systems and 3) e-sports. A multitude of common grounds between many ideas tabled in the keynotes and the contributed talks were evident. Moriwaki et al. [2020] attempts to efficiently model fatigue and increase user satisfaction by diversifying the online advertisements shown to different individuals. This idea aligns well with understanding and optimising for the right user satisfaction metrics that Lalmas and Agichtein emphasised in their keynotes. Omidvar-Tehrani et al. [2020] showed how user preference states can be leveraged to provide point-of-interest recommendations. This application describes many ideas presented by all the keynote speakers in relation to improving recommendations using knowledge about the user’s state. White and Romano [2020] demonstrates the possibility of understanding the user’s state and improving player productivity in e-sports domain showing connections to ideas from the keynotes relating to improving user satisfaction and productivity. To summarise, the keynote talks and the contributed talks in union provided an excellent space to the participants to reflect and realise that the ideas central to state-based user modelling connect to each other across different application domains.

We believe that the WSDM 2020 Workshop on State-based User Modelling (SUM’20) had a rich and diverse set of contributions. WSDM was an excellent venue for the workshop – and we thank them for their assistance in running the workshop. This workshop has proven to be a much needed forum for bringing together practitioners from industry and academia working in user modelling.

---

**Acknowledgments** This work is conducted as part of the X5GON project funded from the EU’s Horizon 2020 grant No 761758 and partially funded by the EPSRC Fellowship titled ”Task Based Information Retrieval”, under grant No EP/P024289/1. This material is based upon work supported by, or in part by, the U.S. Army Research Laboratory and the U. S. Army Research Office, and by the U.K. Ministry of Defence and the U.K. Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/R013616/1.

## References

- Johanne R. Trippas, Damiano Spina, Falk Scholer, Ahmed Hassan Awadallah, Peter Bailey, Paul N. Bennett, Ryen W. White, Jonathan Liono, Yongli Ren, Flora D. Salim, and Mark Sanderson. Learning about work tasks to inform intelligent assistant design. In *Proceedings of the 2019 Conference on Human Information Interaction and Retrieval*, CHIIR ’19, 2019.
- Alexandra Vtyurina, Denis Savenkov, Eugene Agichtein, and Charles L. A. Clarke. Exploring conversational search with humans, assistants, and wizards. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, CHI EA ’17, pages 2187–2193, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-4656-6. doi: 10.1145/3027063.3053175. URL <http://doi.acm.org/10.1145/3027063.3053175>.
- Farouk Belkadi, Mohamed Anis Dhuieb, José Vicente Aguado, Florent Laroche, Alain Bernard, and Francisco Chinesta. Intelligent assistant system as a context-aware decision-making support for the workers of the future. *Computers & Industrial Engineering*, 139:105732, 2020. ISSN 0360-8352. doi: <https://doi.org/10.1016/j.cie.2019.02.046>. URL <http://www.sciencedirect.com/science/article/pii/S0360835219301330>.
- Hao Yan and Ted Selker. Context-aware office assistant. In *Proceedings of the 5th International Conference on Intelligent User Interfaces*, IUI ’00, page 276–279, New York, NY, USA, 2000. Association for Computing Machinery. ISBN 1581131348. doi: 10.1145/325737.325872. URL <https://doi.org/10.1145/325737.325872>.
- Chris Piech, Jonathan Bassen, Jonathan Huang, Surya Ganguli, Mehran Sahami, Leonidas J Guibas, and Jascha Sohl-Dickstein. Deep knowledge tracing. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 505–513. Curran Associates, Inc., 2015. URL <http://papers.nips.cc/paper/5654-deep-knowledge-tracing.pdf>.
- 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*, SIGIR ’17, pages 555–564, 2017.
- Tom Minka, Ryan Cleven, and Yordan Zaykov. Trueskill 2: An improved bayesian skill rating system. Technical Report MSR-TR-2018-8, Microsoft, 2018.
- Gediminas Adomavicius and Alexander Tuzhilin. *Context-Aware Recommender Systems*, pages 217–253. Springer US, Boston, MA, 2011. ISBN 978-0-387-85820-3. doi: 10.1007/978-0-387-85820-3\_7. URL [https://doi.org/10.1007/978-0-387-85820-3\\_7](https://doi.org/10.1007/978-0-387-85820-3_7).



- 
- Moshe Unger, Ariel Bar, Bracha Shapira, and Lior Rokach. Towards latent context-aware recommendation systems. *Knowledge-Based Systems*, 104:165 – 178, 2016. ISSN 0950-7051. doi: <https://doi.org/10.1016/j.knosys.2016.04.020>. URL <http://www.sciencedirect.com/science/article/pii/S0950705116300727>.
- Konstantina Christakopoulou, Filip Radlinski, and Katja Hofmann. Towards conversational recommender systems. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pages 815–824, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4232-2. doi: 10.1145/2939672.2939746. URL <http://doi.acm.org/10.1145/2939672.2939746>.
- Rui Yan, Yiping Song, and Hua Wu. Learning to respond with deep neural networks for retrieval-based human-computer conversation system. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '16, pages 55–64, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4069-4. doi: 10.1145/2911451.2911542. URL <http://doi.acm.org/10.1145/2911451.2911542>.
- Thore Graepel, Joaquin Quiñero Candela, Thomas Borchert, and Ralf Herbrich. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft’s bing search engine. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML'10, pages 13–20, USA, 2010. Omnipress. ISBN 978-1-60558-907-7. URL <http://dl.acm.org/citation.cfm?id=3104322.3104326>.
- Xiangyu Zhao, Long Xia, Jiliang Tang, and Dawei Yin. Deep reinforcement learning for search, recommendation, and online advertising: A survey. *SIGWEB Newsl.*, (Spring):4:1–4:15, 2019. doi: 10.1145/3320496.3320500.
- Felipe Soares da Costa and Peter Dolog. Collective embedding for neural context-aware recommender systems. In *Proceedings of the 13th ACM Conference on Recommender Systems*, RecSys '19, page 201–209, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450362436. doi: 10.1145/3298689.3347028. URL <https://doi.org/10.1145/3298689.3347028>.
- J Shane Culpepper, Fernando Diaz, and Mark D Smucker. Research frontiers in information retrieval: Report from the third strategic workshop on information retrieval in lorne (swirl 2018). In *ACM SIGIR Forum*, volume 52, pages 34–90. ACM, 2018.
- Liangjie Hong and Mounia Lalmas. Tutorial on online user engagement: Metrics and optimization. In *Companion Proceedings of The 2019 World Wide Web Conference*, WWW '19, 2019.
- Janette Lehmann, Mounia Lalmas, Elad Yom-Tov, and Georges Dupret. Models of user engagement. In *International conference on user modeling, adaptation, and personalization*, UMAP '12, 2012.
- A Anderson, L Maystre, R Mehrotra, I Anderson, and M Lalmas. Algorithmic effects on the diversity of consumption on spotify. In *Proceedings of The 2020 World Wide Web Conference*, WWW '20, 2020.

- 
- Jason Ingyu Choi, Ali Ahmadvand, and Eugene Agichtein. Offline and online satisfaction prediction in open-domain conversational systems. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, CIKM '19, 2019.
- Ali Ahmadvand, Harshita Sahijwani, Jason Ingyu Choi, and Eugene Agichtein. Concet: Entity-aware topic classification for open-domain conversational agents. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, CIKM '20, 2019.
- Ali Ahmadvand, Harshita Sahijwani, and Eugene Agichtein. Would you like to talk about sports now? towards contextual topic suggestion for open-domain conversational agents. In *Proceedings of the 2020 Conference on Human Information Interaction and Retrieval*, CHIIR '20, 2020.
- Jason Ingyu Choi and Eugene Agichtein. Quantifying the effects of prosody modulation on user engagement and satisfaction in conversational systems. In *Proceedings of the 2020 Conference on Human Information Interaction and Retrieval*, CHIIR '20.
- Wei Wang, Saghar Hosseini, Ahmed Hassan Awadallah, Paul N. Bennett, and Chris Quirk. Context-aware intent identification in email conversations. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR'19, 2019.
- Xuhai Xu, Ahmed Hassan Awadallah, Susan T Dumais, Farheen Omar, Bogdan Popp, Robert Routhwaite, and Farnaz Jahanbakhsh. Understanding user behavior for document recommendation. In *Proceedings of The 2020 World Wide Web Conference*, WWW '20, 2020.
- Ryen W. White, Ahmed Hassan Awadallah, and Robert Sim. Task completion detection: A study in the context of intelligent systems. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR'19, 2019.
- Rishabh Mehrotra and Emine Yilmaz. Extracting hierarchies of search tasks & subtasks via a bayesian nonparametric approach. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '17, 2017a.
- Sahan Bulathwela, Maria Perez-Ortiz, Emine Yilmaz, and John Shawe-Taylor. Towards an integrative educational recommender for lifelong learners. In *AAAI Conference on Artificial Intelligence*, AAAI '20, 2020a.
- Rishabh Mehrotra and Emine Yilmaz. Task embeddings: Learning query embeddings using task context. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, CIKM '17, 2017b.
- Sahan Bulathwela, Maria Perez-Ortiz, Emine Yilmaz, and John Shawe-Taylor. Truelearn: A family of bayesian algorithms to match lifelong learners to open educational resources. In *AAAI Conference on Artificial Intelligence*, AAAI '20, 2020b.
- Sahan Bulathwela, Stefan Kreitmayer, and María Pérez-Ortiz. What's in it for me? augmenting recommended learning resources with navigable annotations. In *Proceedings of the 25th International Conference on Intelligent User Interfaces Companion*, IUI '20, 2020c.

---

Daisuke Moriwaki, Komei Fujita, Shota Yasui, and Takahiro Hoshino. Fatigue-aware ad creative selection, 2020. Workshop on State-Based User Modelling (SUM '20) at 13th ACM International Conference on Web Search and Data Mining, [https://www.k4all.org/wp-content/uploads/2020/01/WSDMSUM20\\_paper\\_Fatigue\\_Aware\\_Ad\\_Creative\\_Selection.pdf](https://www.k4all.org/wp-content/uploads/2020/01/WSDMSUM20_paper_Fatigue_Aware_Ad_Creative_Selection.pdf).

Behrooz Omidvar-Tehrani, Sruthi Viswanathan, Frederic Roulland, and Jean-Michel Renders. Sage: Interactive state-aware point-of-interest recommendation, 2020. Workshop on State-Based User Modelling (SUM '20) at 13th ACM International Conference on Web Search and Data Mining, [https://www.k4all.org/wp-content/uploads/2020/01/WSDMSUM20\\_paper\\_SAGE\\_Interactive\\_State\\_aware-Point\\_of\\_Interest\\_Recommendation.pdf](https://www.k4all.org/wp-content/uploads/2020/01/WSDMSUM20_paper_SAGE_Interactive_State_aware-Point_of_Interest_Recommendation.pdf).

Alfonso White and Daniela M. Romano. Scalable psychological momentum estimation in e-sports, 2020. Workshop on State-Based User Modelling (SUM '20) at 13th ACM International Conference on Web Search and Data Mining, [https://www.k4all.org/wp-content/uploads/2020/01/WSDMSUM20\\_paper\\_Scalable\\_Psychological\\_Momentum\\_Forecasting\\_in\\_Esports.pdf](https://www.k4all.org/wp-content/uploads/2020/01/WSDMSUM20_paper_Scalable_Psychological_Momentum_Forecasting_in_Esports.pdf).