

Task Intelligence Workshop @ WSDM 2019

Ahmed Hassan Awadallah
Microsoft Research
Redmond, WA, USA
hassanam@microsoft.com

Mark Sanderson
RMIT University
Melbourne, Australia
mark.sanderson@rmit.edu.au

Cathal Gurrin
Dublin City University
Dublin, Ireland
cathal.gurrin@dcu.ie

Ryen W. White
Microsoft Research
Redmond, WA, USA
ryenw@microsoft.com

ABSTRACT

The task intelligence workshop at the 2019 ACM Web Search and Data Mining (WSDM) conference comprised a mixture of research paper presentations, reports from data challenge participants, invited keynote(s) on broad topics related to tasks, and a workshop-wide discussion about task intelligence and its implications for system development.

CCS CONCEPTS

• **Human-centered computing**; • **Information systems** → *Information retrieval*; *Data mining*; • **Computing methodologies** → *Machine learning*;

KEYWORDS

Tasks; Task intelligence; Workshops

ACM Reference Format:

Ahmed Hassan Awadallah, Cathal Gurrin, Mark Sanderson, and Ryen W. White. 2019. Task Intelligence Workshop @ WSDM 2019. In *The Twelfth ACM International Conference on Web Search and Data Mining (WSDM '19)*, February 11–15, 2019, Melbourne, VIC, Australia. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3289600.3290997>

1 INTRODUCTION

Tasks (i.e., defined pieces of work, ranging in scope from specific (e.g., sending an email) to broad (e.g., planning a wedding)) are central to all aspects of information access and use. Task intelligence spans technologies and experiences to extract, understand, and support the completion of short- and long-term tasks. Helping users complete tasks is a key capability of search systems, digital assistants, and productivity applications. Extracting tasks from data is a core challenge in data mining and knowledge representation and draws on additional research from areas such as machine learning and natural language processing. Task can be mined from sources such as email corpora (e.g., commitments and requests [1]) and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WSDM '19, February 11–15, 2019, Melbourne, VIC, Australia

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5940-5/19/02...\$15.00

<https://doi.org/10.1145/3289600.3290997>

search engine log data [4]. Attributes of tasks, such as priority, duration, and progress toward completion, can also be inferred from data (e.g., [5]) and deployed in a broad range of application scenarios in software and services we use every day.

The workshop brings together researchers from academia, industry, and beyond for a half-day meeting on task intelligence at WSDM 2019. The event comprises a mixture of:

- Invited keynote(s) on broad topics related to tasks (e.g., supporting tasks in productivity applications or learning about tasks in search engines)
- Research paper presentations
- Reports from data challenge participants, including system demonstrations if available
- Workshop-wide discussions about task intelligence and its implications for system development

Task completion drives much of our engagement with computer systems. Tasks are therefore central to search and data mining research and applications. Imbuing machines with the ability to mine, model, and support tasks, i.e., possess task intelligence, is increasingly important. The workshop theme covers variety of research topics ranging from task understanding and support for retrieval systems to contextual recommendations and new interfaces to support task completion. This makes it relevant to both the web search and data mining communities at WSDM 2019. The anticipated outcomes from the workshop include:

- Insights from accepted research papers
- Findings from a data challenge on a publicly available dataset
- A prioritized list of research challenges emerging from discussions during the event
- Idea sharing and connections between participants

2 WORKSHOP CONTENT AND DATA CHALLENGE

2.1 Keynote Speakers

We invited speakers from across industry and academia. We are honored to have an impressive lineup of three fabulous keynotes:

- Ed H. Chi (Google AI) – Ed is a Principal Scientist at Google, leading a machine learning research team focused on recommendation systems and social computing research. He has launched significant improvements of recommenders for YouTube, Google Play Store, and Google+. Ed is known for

research on web and online social systems, and the effects of social signals on user behavior.

- Paul Bennett (Microsoft Research AI) – Paul is a Principal Researcher and manager of the Information and Data Sciences group in Microsoft Research AI. He is interested in the development, improvement, and analysis of machine learning methods with a focus on systems that can aid in the automatic analysis of natural language as components of adaptive systems or information retrieval systems. Paul’s current focus is on contextually intelligent assistants. He also maintains an active interest in contextual and personalized search, enriched information retrieval, active sampling and learning, hierarchical and large-scale classification, and human computation and preferences.
- Hongning Wang (University of Virginia) – Hongning is an assistant professor in the Department of Computer Science of University of Virginia. His research interests includes data mining, machine learning, and information retrieval, with a special emphasis on computational user behavior modeling.

At the workshop, the speakers will cover a range of topics related to the main themes of the event, including areas such as search, assistance, productivity, personalization, and recommendation.

2.2 Accepted Papers

We sought submission of papers describing early-stage research on all aspects of task intelligence. In the call for papers, specific topics of interest included:

- Task extraction, e.g., mining tasks from data
- Task understanding, including:
 - Task attributes (duration, priority, etc.)
 - Task context (use contextual signals to understand tasks)
 - Task taxonomies and dependencies, incl. task ownership
- Task support, including:
 - Intelligent notifications (supporting task focus)
 - Intelligent interfaces (novel interfaces for task support)
 - Guided task support (help me do my task)
 - Contextual recommendations (assist me now)
 - Task-focused search (find information related to my task)
 - Scheduling (make time for my tasks)
- Task progress, including:
 - Inferring task completion
 - Summarizing completed tasks
- Task-based evaluation, e.g., gauging user satisfaction and utility from task support systems

Research papers were reviewed by a program committee of experts.

2.3 Task Intelligence Data Challenge

Associated with the Task Intelligence workshop was a data challenge, for which interested participants could download a rich personal data archive of lifelog data annotated with real-world tasks/activities. For the data challenge, the dataset employed for the Lifelog Search Challenge (LSC) [2], a participating workshop at the 2018 ACM International Conference on Multimedia Retrieval (ICMR) was used. The dataset consisted of 27 days of rich multi-modal lifelog data from one lifelogger organized into units of a minute duration. Specifically, the dataset included:

- Multimedia Content. Anonymized wearable camera images gathered using a Narrative Clip 2 at a frequency of about two images per minute and worn from breakfast to sleep. Associated with each image were a set of visual concepts (e.g. computer, car, glass) generated by the Microsoft Cognitive Services (computer vision API¹) [3].
- Biometric Data. 24 × 7 heart rate, galvanic skin response, calorie burn and steps, on a per-minute basis. In addition, daily blood pressure and blood glucose levels and weekly cholesterol and uric acid levels were recorded.
- Human Activity Data. The daily activity of the individual in terms of the semantic locations visited, physical activities (e.g., walking, running, standing) along with a time-stamped diet-log of all food consumed drinks taken, and a location record for every minute.
- Information Activities Data. Using the Loggerman app, the information creation and consumption activities were provided, which were organized into blacklist-filtered, sorted, document vectors representing every minute.
- Tasks. Each minute of the dataset was annotated with a single dominant activity label (e.g., washing clothes, preparing meals, shopping), which were analogous to real-world task annotations.

In all, five participants sought access to the dataset. At the time of writing, we were unsure of the nature of their dataset exploration. The results of the task challenge will be presented at the workshop.

3 CONCLUSIONS

Task intelligence is an important research area with broad practical implications: tasks are central to all aspects of our lives. The workshop presents an opportunity to bring together researchers with diverse interests and backgrounds for discussions on how technology can help with task completion. We hope that the event highlights the wealth of opportunity for impactful work on task intelligence and fosters future research collaborations between workshop attendees.

ACKNOWLEDGMENTS

The authors acknowledge the financial support of Science Foundation Ireland (SFI) under grant number SFI/12/RC/2289. Additional support was provided by the Cortana Intelligence Institute (CII).

REFERENCES

- [1] Simon Corston-Oliver, Eric Ringger, Michael Gamon, and Richard Campbell. 2004. Task-focused summarization of email. In *Proceedings of the ACL-04 Workshop: Text Summarization Branches Out*. 43–50.
- [2] C. Gurrin, K. Schoeffmann, H. Joho, B. Munzer, R. Albatal, F. Hopfgartner, L. Zhou, and D.-T. Dang-Nguyen. 2019. A test collection for interactive lifelog retrieval. <http://eprints.whiterose.ac.uk/137229/> © 2018 Springer Verlag.
- [3] Alessandro Del Sole. 2017. *Microsoft Computer Vision APIs Distilled: Getting Started with Cognitive Services*. Apress.
- [4] Hongning Wang, Yang Song, Ming-Wei Chang, Xiaodong He, Ryen W White, and Wei Chu. 2013. Learning to extract cross-session search tasks. In *Proceedings of the 22nd international conference on World Wide Web*. ACM, 1353–1364.
- [5] Ryen W. White and Ahmed Hassan Awadallah. 2019. Task duration estimation. In *Proceedings of the ACM WSDM Conference on Web Search and Data Mining*. ACM.

¹<https://azure.microsoft.com/en-us/services/cognitive-services/computer-vision/>