

Preface

Research in artificial intelligence and data science is accelerating rapidly due to an unprecedented explosion in the amount of information on the web. In parallel, we noticed immense growth in the construction and utility of the knowledge network from Google, Netflix, NSF, and NIH. However, current methods risk an unsatisfactory ceiling of applicability due to shortcomings in bringing homogeneity between knowledge graphs, data mining, and deep learning. In this changing world, retrospective studies for building state-of-the-art AI and Data science systems have raised concerns on trust, traceability, and interactivity for prospective applications in healthcare, finance, and crisis response. We believe the paradigm of knowledge-infused mining and learning would account for both pieces of knowledge that accrue from domain expertise and guidance from physical models. Further, it will allow the community to design new evaluation strategies that assess robustness and fairness across all comparable state-of-the-art algorithms.

The Workshop on Knowledge-infused Mining and Learning for Social Impact (KiML 2020) was centered around the following thematic components: (a) Data Management: includes resource management, resource discovery across heterogeneous, and inconsistent data resources. (b) Data Usage: includes methods and systems for visualization, representations, reasoning, and interaction. (c) Evaluation: will bring together researchers involved at the intersection of databases, semantic web, information systems, and AI to create new approaches and tools to benefit a broad range of policymakers (e.g. mental health professions, education practitioners, emergency responders, and economists).

The workshop will bring together researchers and practitioners from both academia and industry who are interested in the creation and use of knowledge graphs in understanding online conversations on crisis response (e.g., COVID-19), public health (e.g., social network analysis for mental health insights), and finance (e.g., mining insights on the financial impact (recession, unemployment) of COVID-19 using twitter or organizational data). Additionally, we encourage researchers and practitioners from the areas of human-centered computing, interaction and reasoning, statistical relational mining and learning, intelligent agent systems, semantic social network analysis, deep graph learning, and recommendation systems.

The main program of KiML'20 consists of seven papers, selected out of thirteen submissions, covering topics related to knowledge-enabled feature elicitation, adversarial learning, crisis response, public health, and COVID-19. We sincerely thank the authors of the submissions as well as the attendees of the workshop. We wish to thank the members of our program committee for their help in selecting high-quality papers. Furthermore, we are grateful to Manuela Veloso, Sriraam Natarajan, Jose Ambite, and Pieter De Leenheer for giving keynote presentations on their recent work on Symbiotic Autonomy, Human Allied Probabilistic Learning, Biomedical Data Science, and Data Intelligence.

Manas Gaur, Alejandro Jaimes, Fatma Özcan, Srinivasan Parthasarathy,
Sameena Shah, Amit Sheth, and Biplav Srivastava
August 2020