ISSN 2186-7437

NII Shonan Meeting Report

No.154

Human-in-the-loop Big Data and AI: Connecting Theories and Practices for a Better Future of Work

Sihem Amer-Yahia
Senjuti Basu Roy
Lei Chen
Atsuyuki Morishima

September 9-12, 2019

National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-Ku, Tokyo, Japan

Human-in-the-loop Big Data and AI: Connecting Theories and Practices for a Better Future of Work

Organizers:
Sihem Amer-Yahia (CNRS, France)
Senjuti Basu Roy (New Jersey Institute of Technology, USA)
Lei Chen (HKUST, China)
Atsuyuki Morishima (University of Tsukuba, Japan)

September 9-12, 2019

1. Overview of the Meeting

1.1 Background

"Human-in-the-loop" is a term to denote approaches that combine the power of humans and machines to solve problems. In the future, all of our works will be a part of human-in-the-loop systems and performed by a hybrid workforce comprising humans and machines. Two key fuels there are bigdata and AI; bigdata allows us to find better divisions of labor among workers, and AI agents will play many roles in the systems, such as workers, knowledge extractors and coordinators. Today, human-in-the-loop has been not only a hot research topic but also proven to be effective in many real-world applications, such as production design, citizen science, online job markets, and natural disaster response.

However, there is still work to be done to adapt methodologies and transfer research results across disciplines and to real-world applications. One of the reasons is that researchers tend to work in silos and focus on a specific aspect of the human-in-the-loop systems. The universally accepted human-in-the-loop system architecture is yet to be agreed upon. It usually contains languages, data management systems, data processing systems, and worker-resource managers. We stress that applications are a major component of the human-in-the-loop system architecture, since otherwise the system would not match real-world applications thereby hindering usability.

We note that we should focus more on worker's perspectives in our future of work. Most of the current "human-in-the-loop" researches have focused on the requesters' perspectives such as the cost and time for solving their problems. We believe that the human-in-the-loop systems should be redefined as the "Future of Work" systems, that must deal with human factors for addressing ELSI (Ethical, Legal, Social Issues), such as fairness and transparency.

An important outstanding question in our future of work is how AI will influence future workforces – a common fear is that the machines will fully replace the need of human workers. Instead, we believe that a hybrid workforce will be the future of work comprising of a diverse group of AI-powered machines and humans working together and achieving superior results compared to what either group could accomplish working alone. Given the current state of progress of human-in-the-loop research, we believe that the time is ripe to embark on a fundamental approach to connect research results to real-world problems in human-in-the-loop Big Data and AI for our better future of work.

1.2 Aims of the meeting

The main purpose of this NII Shonan meeting is to bring together researchers from the multidisciplinary fields of human-in-the-loop Big Data and AI as well as practitioners, to connect the latest research findings to real-world questions and examine its impact to future of work. The outputs are a series of reports and vision papers that show the state of the art of this area, and a coherent guideline in the future of work system research, including the reference architecture and research questions. We also establish a community formed by core members to develop (possible) future standards and platforms.

1.3 Outcome

We developed a set of written materials that contain the state of the art, key issues, research questions, and the reference architecture. The summary is given in Section 3. Those results will be published as an ACM SIGMOD blog, a SIGMOD Record vision paper and several other papers.

2. Meeting Schedule

2.1 Overview

The meeting started with talks by organizers followed by a mixture of lightning talks, working groups, and talks by the groups. Working groups create written outcomes and exchange the results after each session. The talks by participants and groups served as inputs for working groups in next sessions. The steps are as follows.

1. Started with three groups that discuss issues from different perspectives. Each group developed scenarios in the future of work (data-centric work, other online work, and offline work), identified issues that need to be addressed. The three groups are as follows:

   [Cross-discipline research group] This group discusses research challenges in different areas and those that require cross-discipline research for the future of work systems.

[ELSI (Ethical, Legal, Social Issues) group] This group develops ELSI considerations and new opportunities in the future of work.
[Applications and Benchmarks group] This group discusses applications and then benchmarks that will be effective in evaluating platforms.

2. Based on the results of the three working groups, we identified seven core aims, referred to as Intellectual Challenges (ICs in short), and the key entities and relationships among them that form a future of work reference architecture, into which the seven ICs are connected. The identified ICs are as follows:

   [IC1: Capturing Human Characteristics and Capabilities] This IC explores how to augment the ability to understanding different types of human roles in future jobs, modeling the inherent uncertainty in human behavior by understanding their evolving characteristics and be able to propose jobs to them by adapting to their changing perceptions, needs, and skills.
   IC2[Stakeholder Requirement Specification] This IC investigates how to add the capabilities of all stakeholders involved in the new job ecosystem to specify their requirements and needs.
   IC3[Social Processes] This IC addresses the development of appropriate interaction methodologies to provide onboarding, socializing, and training for workers, as well as ways to delegate work between humans and humans and machines.
   IC4[Platform Ecosystems] This IC reflects on the design of Future of Work platforms, by emphasizing on Interoperability issues, interaction with the platform environment, communication between platforms, and also human-human / human-machine communication.
   IC5[Computation Capabilities] This IC investigates computational challenges to enable desired characteristics described in other ICS, including, but not limited to the design of personalized, adaptive utility functions for different stakeholders in FoW ecosystem, designing principled algorithms, as well as providing data management capability for effective data analysis.
   IC6[Benchmarks and Metrics] This IC explores Benchmarking and development of appropriate metrics to measure computing capabilities as well as human aspects such as satisfaction, human capital advancement, and equity.
   IC7[Ethics] Questions concerning the security of the data describing individuals, access to the data, how the data are being used, and consent to provide such data are some of many other ethical issues that are discussed in this IC.

3. Then, we formed the seven groups corresponding to the seven ICs, each of which focuses on the state of art and concrete challenges in each IC.

2.2 Detailed Schedule

Check-in Day: September 8 (Sun)

- 19:00 Welcome Banquet - Self-introduction according to the group list.

Day1: September 9 (Mon)

- 9:00-10:30 Opening
  Opening Speech (Organizers)
  Talks by the organizers
    - Sihem Amer-Yahia “Ethical, Legal and Social Issues of Human-in-the-loop Systems”
    - Atsuyuki Morishima “Future-of-Work Platforms”
    - Senjuti Basu Roy “Cross Disciplinary Research”
    - Lei’s talk was moved to the third day due to the delay by Typhoon
- 10:30-10:45 Break
- 10:45-12:30 Lightning talks
    - Andrea Wiggins “Balancing Efficiency and Engagement for Volunteers”
    - Koichiro Yoshida “CrowdWorks: Opportunities for Collaboration with Academia”
    - Jared Kenworthy “Automating the extraction of semantic meaning from text: small teams to *crowdsourcing*”
    - Raghav Rao “Information Categorization in Crowdsourced Crisis Mapping during the 2010 Haiti Earthquake: A Collective Sensemaking Approach”.
    - Munenari INOGUCHI "Expectations of Human-in-the-Loop in Disaster Response”
- 12:30- 13:30 Lunch
- 13:30-14:00 Demo (Marion Tommasi et al.)
- 14:00-15:30 Discussion by the three working groups
- 15:30-16:30 Lightning talks
    - Abhishek Dubey “Concerns of AI in Cyber-Physical Systems”
    - Sudeepa Roy “Making Database Query Answers Understandable with Explanations and Causality”
    - Saravanan Thirumuruganathan “Make *Crowdsourcing* Great Again!”
    - Emilie Hoareau “*Crowdsourcing* in Information Systems Research”
    - Gianluca Demartini. “Improving the Crowd Worker Experience”
- 16:30-18:00 Talks by each group (15 min each)
- 18:00 Dinner

Day2: September 10 (Tue)

- 9:00-10:30 Plenary Discussion

Identify seven intellectual challenges and form seven groups

- 10:30-11:00 Lightning talks
    - James Abello "Peeling Fix Points, Waves and Vases: Primitives for Massive Exploration"
    - David Gross Amblard "Towards Skill-aware, Verifiable and Higher-Order FoW (with a bit of Privacy)"
    - Pierre Senellart "Crowdsourcing: An Instance of Intensional Data Management"
    - Yunyao Li "Human-Machine Co-Creation of Explainable AI Models"
    - Dongwon Lee "Toward AI-Powered Workplace Where Humans and Machines Collaborate"

- 11:15-11:30 Break
- 11:30-12:00 Group photo session
- 12:00-13:30 Lunch
- 13:30-14:00 Demo (Shady Elbassuoni and Ria Borromeo)
- 14:00-15:40 Writing Session by the seven IC groups
- 15:40-16:00 Break
- 16:00-16:10 Talk on the current status of the Future of Work reference architecture (Paolo Papotti)
- 16:10-17:00 Writing Session by the seven IC groups
- 17:00-18:00 Presentation by each group (IC1 to IC7) followed by a talk on the current status of CrowdWorks, inc. (Koichiro Yoshida)
- 18:00 Dinner

Day3: September 11 (Wed)

- 9:00-9:45 Prenaly Discussion
  What should be written in the report by IC groups.
- 9:45-10:45 Writing Sessions by the IC groups
- 10:45-11:15 Break
- 11:45-12:00 Talks
    - Lei Chen "Human-powered Bigdata and AI: Applications and Benchmarks"
    - Raghav Rao "Management of Digital Evidence: A Police Perspective of Ethical Use of Body Worn Camera"
- 12:00-13:30 LUNCH
- 13:30-21:00 Excursion
    - Hokoku -ji Temple
    - Jomyo-ji with Japanese Tea Ceremony
    - Hachiman-gu Shrine

- Dinner at Japanese style restaurant “Minemoto”

Day4: September 12 (Thu)

- 9:30-10:30 Plenary Discussion
  Our next steps and publication venues
- 10:30-11:00 Break
- 11:00-12:00 Writing Sessions by the IC groups
- 12:00 Closing

3. Overview of Lighting Talks and Demos

Balancing Efficiency and Engagement for Volunteers

Andrea Wiggins, University of Nebraska

When the workers in the loop are volunteers, we can't rely on payment to motivate them. Of course, we need efficiency: we can't waste volunteers' time or they will stop contributing, putting system sustainability at risk. But we also need to support engagement through experience design to make the work worthwhile, which requires considering how incorporating AI changes workers' experiences. If applying machine learning to a task eliminates the "easy" work leaving only tasks that are increasingly difficult, it may discourage volunteers and sabotage the entire system. Plus the volunteers are smart and will figure out that they are working alongside machines; how do we manage this? Will it bias their work? Will they second guess themselves or be relieved to have machine support? We could see a boost to performance or a drop in engagement, which might change who participates and how, impacting data quality. We don't know how people will respond to this scenario, but the future of work will require that we find out so that we can balance sustainable and engaging experiences with maximizing system efficiency to create effective human-in-the-loop systems.

CrowdWorks: Opportunities for Collaboration with Academia

Koichiro Yoshida, CrowdWorks Inc.

Crowdworks is the largest crowdsourcing platform in Japan and we went public on the Tokyo Stock Exchange market in 2014. As the clients, over 400 thousand companies, such as Toyota, Honda, Sony, Panasonic, and so many Small businesses, are already registered. And as workers, over 3 million users are using our services now.

There are over 200 types of work that can be requested to a Crowd worker. Mainly, the clients are offering jobs, such as writing articles, designing, programming, engineering, online secretaries, and customer support. Who is registered on Crowd Works? 97% of users are individuals. Most of our workers are in their 30s and 40s, but our service is used by people of all ages, valued also by seniors and stay-at-home mothers. In terms of the contract amount, which means how much the job price is, the top contract amount is Programming and the second one is design. In terms of the number of contractors, the top number of contractors are designers and writers.

I am exploring potential collaborations between businesses and academia to bridge theory and practice in the future of work. I have five real-world crowdsourcing business issues. The 1st issue is ensuring the quality of the evaluations of users and clients. On any crowdsourcing platform, the quality of the work must be ensured, and low-quality users and clients must be identified. But, the accuracy, reliability, and objectivity of the evaluations are not still guaranteed. The 2nd issue is optimizing the order price. Can we trust the "Invisible Hand"? The more users who participate in the bidding, the more likely that the order price is driven downward. On the other hand, the cheapest bidder does not always leave the client fully satisfied. Clients are often more satisfied with the high-skilled workers, even if the price is much higher. The 3rd issue is preventing fraud and scam. In our business and many others, there are spammers and multi-level marketing companies that register as a client or worker and scam others to steal their money. We have built algorithms to detect such fraud users, but the scammers have developed ways to avoid detection. There are limits for humans to detect these users. The 4th issue is preventing money laundering. Similar to the previous issue, there are money launderers who actively work to steal money from our platform. They are different from MLMs or scammers in that they act as a good client for a long time -- building a good reputation across a few months -- then suddenly launder money by abusing the CrowdWorks advance payment system.

The 5th issue is balancing demand and supply on the platform. Every crowdsourcing platform always has a supply and demand. On our platform, the demand is higher than the supply. The reason is not completely known - maybe the contract prices are too low or there are not as many workers with the required skills for the contract work. With academia, I would like to explore ways in which AI can balance demand and supply. In the future, I imagine that when the client registers for a job, the AI will tell the client the rate of contract for that type of job and give suggestions to increase the rate (i.e. increase price, include more details about the job, etc.). I hope to work with each of you to discuss how current and future research can support businesses to resolve these issues.

Automating the extraction of semantic meaning from text: small teams to crowdsourcing

Jared Kenworthy, University of Texas at Arlington

In this presentation I discussed two related research challenges from a collaboration between psychology and computer science researchers. The research context was described as collaborative creativity, which is a research paradigm that brings groups of people together to come up with novel uses for various commonplace objects (e.g., coffee mug, shoelace, brick, etc.). One objective is to use a wide array of individual characteristics, personality dimensions, prior performance, as well as combinations of such variables in group settings, to predict the number of unique ideas generated. The second, more challenging objective, is to use a computational approach to tap the language of the ideas generated to evaluate the quality and novelty of the ideas so that computers/AI agents can facilitate human creativity and prevent inefficiency and process loss. This has proven, across other research domains, to be particularly difficult because of the known difficulties in deriving and using semantic meaning from human texts. We continue to work toward solving this problem and welcome solutions from the computational community.

Information Categorization in Crowdsourced Crisis Mapping during the 2010 Haiti Earthquake: A Collective Sensemaking Approach

Valecha, R., Oh, O., Rao, H. R. (2019). Information Categorization in Crowdsourced Crisis Mapping during the 2010 Haiti Earthquake: A Collective Sensemaking Approach. Working Paper, September 2019. ISCS, University of Texas at San Antonio.

Raghav Rao, University of Texas at San Antonio

Crisis mapping is the process of real-time collection and visualization of crisis data for humanitarian relief. Focusing on a real-life event, the 2010 Haiti earthquake, we explore the specific aspect of information categorization on a crisis mapping platform (known as Ushahidi). Information categorization is a process wherein messages from affected citizens (victims) are categorized by crowd volunteers (digital humanitarians) for use by crisis responders. In order to make categories actionable for onsite response and recovery efforts, first responders have to be confident that the categories are reliable. However, the reliability of information categorization has been questioned in the context of crowdsourced crisis mapping. This is because much of the victim reported information is ambiguous or incomplete and because crowd volunteers may not have the requisite training for processing such information. This leads to the following research questions: Are there features of citizen-reported crisis messages that can lead to reliable categorization? How can the capabilities of technology platforms help the online crowd volunteers in reliable categorization?

For investigating the reliability of categorization by crowd volunteers, we utilize agreement (majority vote or consensus) among volunteers and/or evaluators. We compare information categorization by Ushahidi volunteers during the event with post-

event categorization by different groups of evaluators – CrowdFlower volunteers and a Registered Nurse-led team. We identify categories that have varied degrees of evaluator agreement. Subsequently, we employ collective sensemaking as an overarching framework to investigate the drivers of agreement within information categorization. Collective sensemaking is the shared comprehension of crisis messages that is facilitated through an understanding of the duality of (a) the crisis context and (b) interaction of crowd volunteers with the crisis mapping platform. We develop a research model that characterizes agreements within information categorization in terms of social and situational cues, as well as information structuring and crisis mapping interactions (captured through posts within crisis reports). Based on an analysis of 1,459 crisis reports in the Ushahidi crisis mapping platform for the 2010 Haiti Earthquake, we find the cues that are positively associated with agreements among crowd volunteers and crowd evaluators. In addition, crowd interaction with the platform is also seen to have a positive relationship with agreements regarding information categorization.

KEYWORDS: Crisis mapping, Digital humanitarianism, Crowd-sourcing, Information categorization, Collective sensemaking, Ushahidi, Haiti earthquake, Social cues, Situational cues, Posting mechanism

Expectations of Human-in-the-Loop in Disaster Response

Munenari Inoguchi, University of Toyama

Firstly, I am from the field of practitioners. Once disaster occurs, I always visit to the affected area and monitor the activities of local responders and survivors. In affected area, I find out some issues which can be solved by ICT, and I develop some support tools on-site as a prototype system, and implement it to validate the efforts of ICT by myself. I know there are many kinds of advanced technology, methodology, knowledge and professions in the field of ICT, however there is no way to integrate them for solving actual problems occurring in affected site at disaster. My main purpose of joining to this meeting is to find out effective solution by Human-in-the-Loop for disaster response and damage detection at actual disaster.

In order to proceed effective disaster response, responders have to grab the actual damage situation for designing strategic plan. However, it always takes much time to do it at huge disaster. In my presentation, I introduce the actual situation about how long it took to confirm disaster damage at Tohoku Earthquake in 2011. In this case, it took 1 month to confirm human damage, and 6 months to confirm building damage. The reason why it took much time to confirm building damage, local responders inspected building damage by visiting each building one-by-one. Following this method, one inspection team can detect the damage of only 30 buildings per day. Against this situation, my research team established a web-GIS system in which remote users can judge each building is washed out or not by referring aerial photos taken after disaster. This was a kind of crowdsourcing. At that time, we have no knowledge to promote it effectively and efficiently, so it took about 3 months to accomplish judgement of building damage. After this activity, I met Prof. Morishima and learned the framework of crowdsourcing, then I found that crowdsourcing is effective and essential way to clarify the damage situation urgently after disaster occurrence.

In recent research, I tried to detect roof damage from images taken by drones in a case study of 2019 Yamagata-oki Earthquake. Murakami city in Niigata prefecture was affected by this earthquake, and many buildings were damaged. However, most of damage were concentrated on the roof of building. Then, we decided to detect roof damage by drones. Our research team designed the plan of drone flight, and operated drones to take pictures of damage roofs. We created orthophoto mosaic from those images, and we published it for Murakami city officers in cloud-based GIS platform. However, it took much time to arrange the environment because we have not enough experience of image processing. After accomplishment of this arrangement, they survey the roof damage of each building referring the orthophoto mosaic, and they understood that survivors were suffered by roof damage. Then, Murakami city decided to create new support program for roof damage relief. Just now, we challenge to monitor the progress in survivors' life reconstruction by taking images of roof damage periodically. In this, we try to utilize object detection method with deep-learning in AI technology to detect roof damage, which is represented by blue-sheet covered over damaged roof. The accuracy of blue-sheet detection is not so high now, then we will find the effective way with Human-in-the-Loop to improve teaching data and processing model.

KEYWORDS: Disaster Response, Damage Detection, Rational Decision Making

CrowdFlow, a high level language for crowdsourcing applications (Demo)

Marion Tommasi, INRIA

We introduce CrowdFlow, a high level language for complex crowdsourcing workflows based on collaborative data-centric workflows.
In this demo, I show the results of a workflow written with CrowdFlow and compiled for the platform Headwork. In this workflow, a worker can decide to answer a task or redistribute it to others if she cannot complete it herself.

Concerns of AI in Cyber-Physical Systems

Abhishek Dubey, Vanderbilt University

Cyber-Physical Systems (CPS) are used in many applications where they must perform complex tasks with a high degree of autonomy in uncertain environments. Traditional design flows based on domain knowledge and analytical models are often impractical for tasks such as perception, planning in uncertain environments, control with ill-defined objectives, etc. Machine learning based techniques have demonstrated good performance for such difficult tasks, leading to the introduction of Learning-Enabled Components (LEC) in CPS. Model based design techniques have been successful in the development of traditional CPS, and toolchains which apply these techniques to CPS with LECs are being actively developed. However, there are still several gaps in understanding the risks involved in the use of LEC based approaches. In this talk we examine the underlying differences between traditional CPS design and the CPS design with LEC. We also examine the problems of assuring the correctness of the LEC components and its impact on the overall safety of the system.

Making Database Query Answers Understandable with Explanations and Causality

Sudeepa Roy, Duke University

In recent times, data is considered synonymous with knowledge, profit, power, and entertainment, requiring development of new techniques to extract useful information and insights from data. In this talk, I will describe our work in intervention-based data analysis toward the goal of understanding data and ensuring interpretability of query answers for a broad range of users. First, I will talk about approaches to explaining query answers in terms of “intervention” (how changes in the input data changes the output of a query) and “counterbalances” (how an outlier can be explained by an outlier in the opposite direction). Then I will discuss how to facilitate understanding and exploration data and query answers with useful graphical interfaces. Finally, I will discuss how to go beyond correlation by causal inference from observational data from the Statistics literature, and how it benefits from techniques in data management.

Making Crowdsourcing Great Again!

Saravanan Thirumuruganathan, Qatar Computing Research Institute

Currently crowdsourcing employs a small fraction of the human workforce ranging in the hundreds of thousands. In the future, it is possible that the vast majority of the workforce will be involved in a crowdsourcing style platform. It is important to think about the design principles for these next-gen platforms. These crowdsourcing platforms must be inherently collaborative and customized to each task. Finally, it must also be locally installable on-premise so that it could be used by any organization. We are currently working on such systems for two domains: manual data cleaning and systematic review with promising results.

Crowdsourcing in Information Systems Research

Emilie Hoareau, University of Grenoble Alpes

As platform is one of the main drivers of working transformation, understanding current issues relating to crowdsourcing is a first step to handle forthcoming challenges of future of work. This presentation offers an overview of current crowdsourcing issues in the Information Systems field which belong to management science. Crowdsourcing studies in Information Systems address three main questions: what is crowdsourcing? Under which circumstances to outsource to the crowd? How to incentivize the crowd? What are the best ways to manage the crowd for creating value? What are the drawbacks of crowdsourcing? As existing studies aims to optimize crowdsourcing activities for all stakeholders, they also point out the growing risk of crowd exploitation. Envisioning the future of work demands therefore a deep understanding about the distribution of power between workers, requesters and platform. We assume that crowd exploitation can be overcome through ethical reflections with a greater focus on platform's responsibilities.

Improving the Crowd Worker Experience

Gianluca Demartini, University of Queensland

In this talk I have discussed recent research where we looked at ways to improve the crowd worker experience including 1) work on understanding the phenomenon of workers abandoning tasks after having completed work (such partially completed tasks result in unrewarded work as workers abandon tasks before completion), and 2) research on the power imbalance between workers and requesters where we provided workers with tools to be aware on which tasks their work quality is being evaluated by automatically detecting gold questions in crowdsourcing tasks.

Peeling Fix Points, Waves and Vases: Primitives for Massive Exploration

James Abello, Rutgers University

We present efficient algorithmic mechanisms to partition graphs with up to 1.8 billion edges into subgraphs which are fix points of degree peeling. For fixed points which are larger than a desired interactivity parameter we further decompose them with a novel linear time algorithm into what we call "Graph Waves and Fragments". Fix points, Waves and Fragments have visual representations that we call Vases and Trapezoid Forks. We illustrate these visual abstractions in 2D and 3D with a variety of publicly available data sets; these include social, web, and citation networks.

Towards Skill-aware, Verifiable and Higher-Order FoW (with a bit of Privacy)

David Gross Amblard, Rennes 1 University

Current crowdsourcing platforms (turning potentially into Future of Work platforms) are very limited: task are repetitive (micro-tasks), skills of participants, career path or promotion mechanisms are not or barely modeled. Moreover, complex tasks (beyond simple chaining of simple tasks) cannot be expressed. In this talk, we describe the concepts we are discussing within the HEADWORK project (ANR funding) that could overcome these limitations. For example, we propose that declarative workflows (in the spirit of Business Artifacts) should be at the center of the model, allowing complex jobs to be described, delegated, or redesigned by user themselves. We also present the HEADWORK prototype that illustrates these notions along with skill modeling and formal verification of job designs.
Project homepage: http://headwork.gforge.inria.fr/
Sourcecode homepage: https://gitlab.inria.fr/druid/headwork (ask for access if needed)

Crowdsourcing: An Instance of Intensional Data Management

Pierre Senellart, ENS, PSL University

Many sources of data are intensional in the sense that data is not directly available in extension, but access to data access has a cost (which can be a computational cost, a monetary cost, time, a privacy budget…). This is the case of crowdsourcing, where one needs to pay or incentivize workers and wait for tasks to be completed, in order to get access to data; but this also applies to a large variety of settings (the deep Web, complex automated processes, reasoning over data, etc.). Intensional data management is about taking into account the cost of data access while solving a user's knowledge need, by building a recursive, dynamic, adaptive knowledge acquisition plan that minimizes access cost, and provides probabilistic guarantees on the quality of the answer.

Human-Machine Co-Creation of Explainable AI Models

Yunyao Li, IBM Research-Almaden

While the role of humans is increasingly recognized in machine learning community, representation of and interaction with models in current human-in-the-loop machine learning (HITL-ML) approaches are too low-level and far-removed from human's conceptual models. In this talk, I present ongoing work in my team to support human-machine co-creation with learning and human-in-the-loop techniques. In particular, I will focus on three topics: (1) how to use machine learning to leverage crowdsourced work in effectively to achieve expert-level quality while minimizing expert workload; (2) SystemER: how to learn an explainable AI model with active learning and a declarative system; (3) HEIDL, a system supports human-machine co-creation by exposing the machine-learned model through high-level, explainable linguistic expressions. In all three, human's role is elevated from simply evaluating model predictions to interpreting and even updating the model logic directly by enabling interaction with rule predicates themselves. Raising the currency of interaction to such semantic levels calls for new interaction paradigms between humans and machines that result in improved productivity for model development process. Moreover, by involving humans in the process, the human-machine co-created models generalize better to unseen data as domain experts are able to instill their expertise by extrapolating from what has been learned by automated algorithms from few labelled data.

Toward AI-Powered Workplace Where Humans and Machines Collaborate

Dongwon Lee, Pennsylvania State University

Among diverse challenges that future of work may face, in particular, I first argue that how to address the retraining of human workers be a critical issue in an AI-powered environment where humans and machines collaborate and compete. Then, I lay out a few intellectually interesting research questions that we need to solve.

Exploring Fairness of Ranking in Online Job Marketplaces (Demo)

Shady Elbassuoni, American University

In this talk, I have demonstrated FaiRank, an interactive system to explore fairness of ranking in online job marketplaces. FaiRank takes as input a set of individuals and their attributes, some of which are protected, and a scoring function, through which those individuals are ranked for jobs. It finds a partitioning of individuals on their protected attributes over which fairness of the scoring function is quantified.

My “Crowdsourcing Platform” (Demo)

Ria Borromeo, University of the Philippines Open University

A common practice in validating task assignment algorithms in crowdsourcing is by first asking a worker to perform all possible tasks. For example, if there are ten tasks, worker A should do all ten tasks. After the algorithm is executed, and an assignment is generated, the metrics for only the assigned tasks are computed. If the algorithm assigns tasks 1 to 5 to worker A, only metrics such as quality, cost, and latency are computed for tasks 1 to 5. While this practice is cost-effective when comparing different task assignments, it does not capture the actual circumstances when a task is completed. In this talk, I introduce a simple web application that takes in three inputs: a database of tasks from the Figure Eight Open Data Library, HTML task templates, and task assignment generated by an algorithm. It then allows workers to complete only the tasks assigned to them. The answers provided by workers, along with task metadata, are stored in an SQL database, which can be easily analyzed later on.

Management of Digital Evidence: A Police Perspective of Ethical Use of Body Worn Camera

Jaeung Lee, Jingguo Wang, Gerald Cliff, and H. Raghav Rao (September 2019). Working paper, Louisiana Tech University

Raghav Rao, University of Texas at San Antonio

Body Worn Camera (BWC) is an emerging Information Technology and System (IT) artifact that has recently started to be used in law enforcement. Although BWC can bring many benefits, it may also result in negative outcomes, such as loss of citizens' privacy or failure of proper management of digital evidence. In this research, we develop a research model that focuses on police officers' perspective about the ethical use of BWC, from the standpoint of organizational justice as well as risk and benefit of using BWC. The paper specifically develops hypotheses to test the relationship between three factors: 1) Work Environment, 2) Risk of using BWC and 3) Benefit of using BWC and an outcome variable, police perceptions of ethical use of BWC. In addition, it tests the moderating effects of work related uncertainty on the above three factors. Finally, relationships between organizational justice factors and Work Environment also are presented as hypotheses. The paper also develops three gray BWC scenarios that are not clearly ethical nor unethical. We apply a survey methodology to test the research model in the context of the gray scenarios. The results show that work motivation and risk of using BWC are negatively related with the police officers' perceptions about ethical use of BWC and these relationships were negatively moderated by work related uncertainty, while justice constructs are positively associated with perceived work motivation. Understanding such issues will assist in development of policies and help in the provision of actionable guidelines for BWC use.

KEYWORDS: Police Body Worn Camera, Perceptions about Ethical Use of IT Artifact, Emerging Technology

4. Identified Issues, Open Questions and a Reference Architecture - Imagine all the People and AI at the Future of Work

We envision Future of Work to be a new world where people are empowered by providing them with the ability to rely on AI machines in an on-demand fashion, and enabling continuous knowledge and skill acquisition and improvement with a variety of onboarding and training tools. Such an environment will allow everyone everywhere to get a job online and offline, train for a new job, and get help from a mix of people and AI machines. Ensuring portability between platforms and guaranteeing the protection of workers' rights, will play a major role in providing a rewarding and safe work environment to all. Freelancing platforms such as CrowdWorks in Japan, TaskRabbit and Fiverr in the USA, and Qapa and MisterTemp in France, and crowdsourcing platforms such as Crowd4U in Japan, Wirk, and Prolific Academic in Europe, and Amazon Mechanical Turk and Figure Eight in the USA, must rethink their design to be at the frontier of new FoW.

Today, people's relationship to work is changing as online job platforms are blurring the boundaries between physical and virtual workplaces. Prospective employees can find temporary jobs in the physical world (e.g., a plumber, an event organizer, a gardener, can offer their jobs online), or in the form of virtual gigs (e.g., logo design, web programmer). Job providers can hire one or many individuals to achieve a task. The same person can take on those roles at any point in time. An employer can be a regular citizen who needs to hire a plumber, a social scientist needed to conduct population studies to verify some theories, a data scientist needing to validate a new algorithm, a domain expert seeking to verify how much interest a new product generates. The diversity of needs has given rise to a variety of platforms, all of which act as intermediaries between job providers and job seekers. Platforms differ in their ability to manage physical and virtual jobs, in their support for onboarding, socializing, training, and credentialing for employees, in automating the matching between jobs and workers. They also differ in the tools they offer to workers and job providers to express needs and requirements, and in their compliance with labor-related regulations and their handling of ethical concerns.

Human Factors in the workplace.

FoW will witness an evolution of humans from being mere "agents" or robots whose efforts and capabilities are used for the benefit of AI systems (broadly machines or business) to understanding psychological characteristics, attributes, skills, motivations, goals, etc. This will require capturing uncertainty in human behavior and individuals' evolving needs, and adapting offers to available demand. Additionally, the attitudes, values, opinions regarding the processes, policies, and outcomes (e.g., perceptions of justice, fairness, bias, etc.) will need to be assessed and considered in the design of FoW ecosystems.

As intelligent systems are increasingly powerful and pervasive in augmenting, supporting, and sometimes replacing human work, several shifts in human skills and related technologies will need to occur. Individual workers will need stronger technical skills and familiarity with systems tailored to their professions or work. Technologies will become more specialized, more closely integrated and interoperable, and will automate many otherwise general functions and "trivial" tasks, as well as taking over more sophisticated functions focused on aggregating, summarization, detecting or identifying opportunities for improvement in work products and processes, and making recommendations. This will leave workers with more time and systems support to focus on highly skilled aspects of their jobs, exercising and refining human-specific skills, such as empathy, and increasing the amount of specialized, highly-skilled work that they are able to handle by streamlining many supporting tasks and functions. Workers will also provide seamless feedback to automated systems as a part of the work that they do, such as through customizing and modifying standardized processes and plans, or augmenting records with direct observations and contextual information, which intelligent agents can then incorporate into future system operations. Workers will also take on a more supervisory role, both over their own work as well as the performance of intelligent systems that support their work or to which they delegate work, with their feedback providing corrective input that is used to continuously improve system performance.

Intellectual challenges of FoW.
It is safe to assume that the future of work will be increasingly technology-driven, with the real opportunity that human concerns are brought to the center of the design and

deployment of job platforms. It is therefore crucial to keep ahead of this trend by addressing the intellectual challenges that arise from this trend.

The first challenge relates to the ability to understand different types of human roles in future jobs, modeling the inherent uncertainty in human behavior by understanding their evolving characteristics and be able to propose jobs to them by adapting to their changing perceptions, needs, and skills (IC1). The ability for all stakeholders involved in the new job ecosystem to specify their requirements needs to be addressed by providing declarative and high level tools to express needs and expectations (IC2). The third challenge addresses the development of appropriate interaction methodologies to provide onboarding, socializing and training for workers, as well as ways to delegate work between humans and humans and machines (IC3: Social Processes). Bringing humans back to the frontier will require overhauling the design and engineering of online job platforms to enable the collection, storage, retrieval, analysis, and mining of a wide array of human data across different types of technology-driven work.

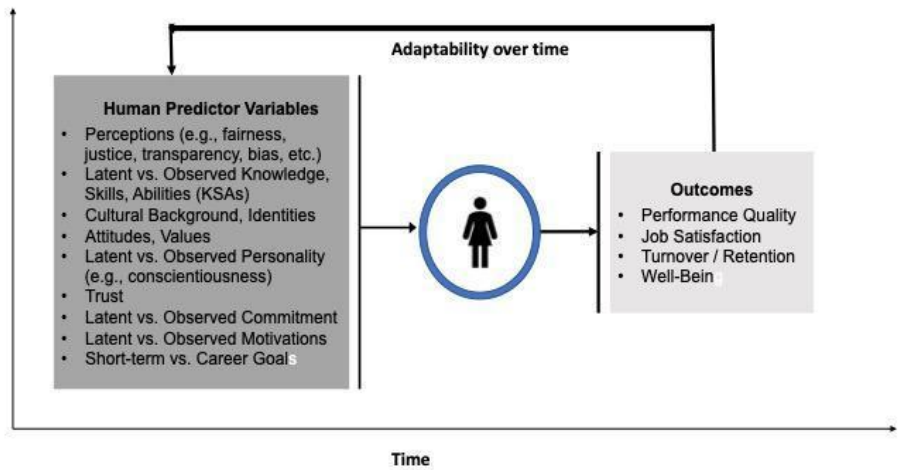

Figure 1. **A representation of human predictor variables and key outcome metrics in Workspace.**

A fair bit of engineering and testing will be needed to ensure the development of scalable and portable platforms and the integration of multi-stakeholder goals in efficient and effective ways (IC4 and IC5). Benchmarking and the development of appropriate metrics to measure computing capabilities as well as human aspects such as satisfaction, human capital advancement, and equity, is another intellectual opportunity (IC6). Bringing humans back to the frontier of future work will improve their quality of life, ensure better work performance, and positively shape long term social and economic outcomes of a society and a nation. However, if we are not cautious enough, there are some inevitable risks. Questions concerning the security of the data describing individuals, access to the data, how the data are being used, and consent to provide such data are among the many issues that must be considered (IC7).

FoW Architecture.
Human and AI workers operate the different components of the FoW architecture. Their characteristics are captured and updated over time (IC1). The requirements (IC2) come from all subjects, including regulators for policies, AI workers for job deployment and assignment, and human workers who are interested in specifying their needs (e.g., skill improvement and expected compensation).
All subjects in the architecture interact and generate human-human, human- machine to enable social processes (IC3), as well as machine-machine communication for interoperability and portability (IC4). The FoW agent, or platform, acts as

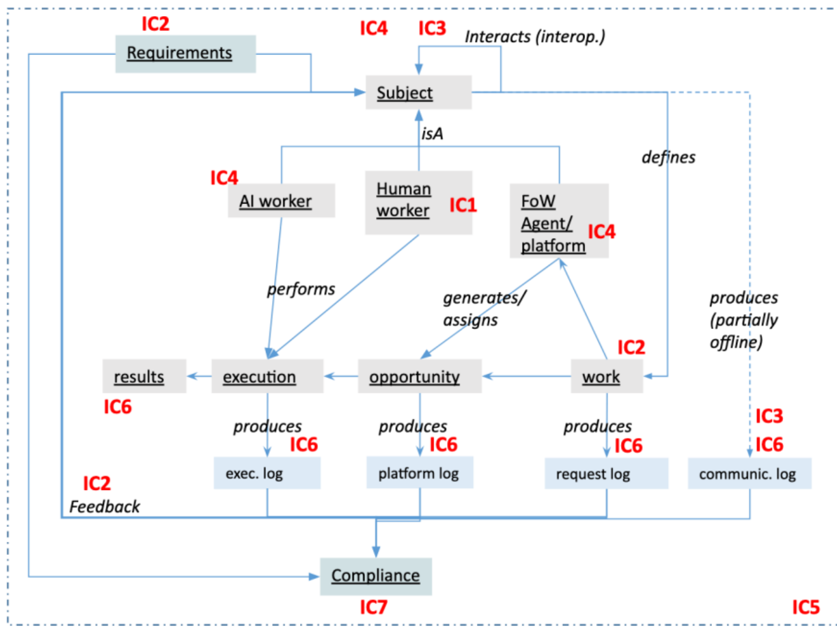

coordinator/scheduler. Any subject can define work through specifications (IC2). Work is handled by the platform, which generates and assigns opportunities to the workers. Opportunities are executed and generate contributions that can be used in benchmarks and metrics (IC6). All processes benefit from scalable computing capabilities (IC5). Processes are monitored to create logs of metadata used in 2 essential ICs: feedback to the subjects and verification of the compliance of the process wrt the requirements (IC7), and for benchmarking purposes (IC6).

Cross-Disciplinary Research.
The relevant research communities have been making some attempts to capitalize on human potential and improve workers' well-being. We believe it is now time for a deeper integration of what different research communities excel at. For example, CS research has focused on developing piecemeal solutions for the job design pipeline. It would benefit from approaches that will help to fully account for changing human characteristics, and to provide support for different needs and preferences of workers. Psychologists, sociologists, and organization management researchers, on the other hand, have conducted user experiments and surveys, interviews, and physical laboratory experiments that have tested and contributed to a range of social science theories. Their contributions would greatly benefit from computational communities to develop

quantitative models. It will be important for these different research communities to engage in collaboration and cross-talk in order to unveil and tap the full capabilities and potential of all subjects in FoW. Effective collaboration and genuine multi-disciplinary research will help approach the goal of optimizing AI systems while at the same time protecting and maximizing worker satisfaction, well-being, and performance potential.

Addressing the intellectual challenges of FoW will require integration and convergence of disciplines across computer science, engineering, education and workforce training, and social, behavioral, and economic sciences, and law. For example, psychology, management, and organization studies research has to understand to what degree do workers' perceived fairness and transparency affect the satisfaction of workers and/or employers? This question, and related ideas, can be tested using different research methods, such as broad user surveys and more controlled laboratory experiments. Similar methodologies can be employed across different platforms, adapting to the specifics of each case. Adapting organizational commitment framework to this work context may also be important. Computing communities, such as Data Management and AI researchers, on the other hand, have to focus on "human-centric" databases, with the goal of efficient storage, retrieval, and analysis of human-data that changes over time, as well as design algorithms that adapt to evolving needs and preferences of workers. Such a convergent perspective will be essential to understanding and shaping long-term social and economic outcomes, to fully exploit the potential of technology-driven future of work while maintaining a human-centric approach.

5. List of Participants

- James Abello Monedero, Rutgers University, USA
- Sihem Amer-Yahia, CNRS, France
- Senjuti Basu Roy, New Jersey Institute of Technology, USA
- Pierre Bourhis, CNRS, CRIStAL, France
- Lei Chen, HKUST, Hong Kong
- François Charoy Francois, University of Lorraine, Inria, CNRS, France
- Marina Danilevsky, IBM Research - Almaden, USA
- Gautam Das, University of Texas at Arlington, USA
- Gianluca Demartini, University of Queensland, Australia
- Abhishek Dubey, Vanderbilt University, USA
- Shady Elbassuoni, American University of Beirut, Lebanon
- David Gross-Amblard, Rennes 1 University, France
- Emilie Hoareau, University Grenoble Alpes, France
- Munenari Inoguchi, University of Toyama, Japan

- Jared Kenworthy, University of Texas at Arlington, USA
- Itaru Kitahara, University of Tsukuba, Japan
- Dongwon Lee, Pennsylvania State University, USA
- Yunyao Li, IBM Research - Almaden, USA
- Ria Mae Borromeo, University of the Philippines Open University, Philippines
- Atsuyuki Morishima, University of Tsukuba, Japan
- Paolo Papotti, EURECOM, France
- Raghav Rao, University of Texas at San Antonio, USA
- Sudeepa Roy, Duke University, USA
- Pierre Senellart, ENS, PSL University, France
- Keishi Tajima, Kyoto University, Japan
- Saravanan Thirumuruganathan, Qatar Computing Research Institute, Qatar
- Marion Tommasi, INRIA, France
- Kazutoshi Umemoto, The University of Tokyo, Japan
- Andrea Wiggins, University of Nebraska Omaha, USA
- Koichiro Yoshida, CrowdWorks Inc., Japan