Kasthuri Jayarajah | Research Statement

The proliferation of ubiquitous technologies such as the Internet of Things, personal wearable sensors, and autonomous vehicles, enable smarter and efficient communities. A wide variety of rich sensing modalities available on these devices such as vision, motion, and physiology of the persons, opens interesting possibilities; coordination amongst a group of traffic cameras at an intersection for improved situation awareness; robotic teammates such as indoor drones that respond to and proactively assist human partners in safety critical environments; smart buildings that dynamically adjust lighting and heating control based on its occupants' routines. In my work, I build **real-time, human-centered, cyber-physical systems that address challenges in Mobile Computing, Artificial Intelligence of Things, and Human-Autonomy Teaming.** Due to the multidisciplinary nature of my work, I'm versed in (a) prototyping systems for resource-constrained edge platforms, wearable devices and autonomous vehicles, and (b) conducting survey-based user studies as well as field studies involving technology-mediated behavioral interventions.

As my colleagues and I articulate in recent articles [1, 2], realizing such human-centered systems operating at the edge of the network is often challenged by two factors: (a) formidable energy and compute constraints restrict the types of processing loads these hardware platforms can handle, especially for perception tasks, and (b) not accounting for individual and collective behaviors of people who interact with such systems can lead to suboptimal performance. Therefore, my research approach involves either enhancing the decision-making fidelity, or significantly reducing the resource demands, of such systems, revolving around the following themes:

- **Wearable-based implicit guidance for autonomous agents** that adapt runtime behavior [1,3],
- **Collaborative perception** where networked cameras cooperate for improved performance through lightweight state sharing [2,4,5], and,
- **Human behavioral characterization** for building predictive cyber-physical systems [6-12]**.**

My long-term research goal is to push the frontiers in context-aware computing through collaborative human-machine and machine-machine systems. For example, I envision Future Work environments and Human(s)–Swarm teams that adapt to individual and collective human intent, preferences, performances and nuances, through a combination of wearable, infrastructural sensors and actuators. Below, I describe my current and past efforts towards realizing this future vision.

1. Wearable-based Guidance for Attention–Sharing, Human–Drone Teams

In this work, I explore the concept of wearable-based "passive guidance" from human(s) in the team in helping direct machines towards maximizing the effectiveness of human-machine combined sensing objectives. While communication in such teams is imperative for successful team performance, direct communication (e.g., verbally) can be distracting to humans under high cognition workloads (e.g., first responders involved in search and rescue). Thus, "implicit coordination" where machines are able to synchronize with their human-teammates without explicit intervention has its advantages. Our vision for such ***implicitly interactive, attention-sharing human--machine teams [1]*** is motivated by two salient trends: the availability of (a) rich sensor modalities such as fine-grained motion, gaze, gestures and physiological sensing on off-the-shelf wearable devices allow for continuous, passive monitoring of nonverbal, implicit behavioral cues, and (b) tiny autonomous drones with both sensing and on-board compute resources that are also safe to operate in close proximity to humans. In ongoing work, I explore scenarios where partnering drones efficiently scan environments and assist humans, overcoming challenges owing to their size, weight, and power (SWAP) constraints, through ***gaze-based uncertainty predictions, assisted spatial awareness[1] and efficient resource***

[1] Short video demonstration of a drone synchronized with the human's motion: https://www.youtube.com/watch?v=INQULM5csMU

orchestration. Coordinated scanning includes two goals: (1) *complementary* - achieve wider coverage where the machine is able to scan regions where the human is not paying attention to, and as a team, achieve faster and efficient scanning, and (2) *collaborative* - higher resolution scanning and inference where a human requires more accurate visibility and augmentation from the machine. In both cases, machines require to know where the human is paying attention to, and not just what is within the "visible" range. Through this work, I'm developing techniques to gauge human visual attention using a combination of wearable technologies that allow for accurate tracking of eye movements (e.g., using noisy EOG signals) in the presence of motion artifacts (determined using on-body inertial and EEG signals) and investigate techniques for continuous, light-weight exchange of attention information. The drone then adapts its attentional focus and/or resolution, on-the-fly, to synchronize with their human-teammates' intent.

While the above direction focuses on efficiency for machines, in ongoing work, my work explores how machines can in turn ***augment human functions through continuously monitoring nonverbal cues and active intervention*** *[3]*. While the effectiveness of physiological sensing (such as via EEG) has been demonstrated with high-fidelity sensors, in lab and controlled settings, for tasks such as inferring emotion, the ability to infer nonverbal cues using commercial-off-the-shelf devices (e.g., the 5-channel Emotiv Insight) with far-degraded, sparse signals, is relatively under-explored. Thus, a key research question this work addresses is in reliably sensing such cues under sparsity by exploring functional connectivity and multimodal sensor fusion based solutions. My colleagues and I studied recall during navigation as an example case study where active intervention from machines can improve memory functions in people. In user studies conducted on desktop-based navigation tasks, we found that visually cueing participants of their EEG-derived, high attentional moments during their walks, increases their ability to recall the route significantly better (by up to 12%).

2. Collaborative Perception at the Edge

A variety of physical environments, including smart cities and tactical battlefield networks are increasingly being instrumented with large numbers of resource-constrained sensors and IoT devices (e.g., cameras, microphone arrays and environmental sensors). A rising recent trend involves executing inferencing pipelines (to perform increasingly complex tasks, such as object recognition or target localization), ***in-situ*** and in ***real time***, at such edge nodes. There are two salient features associated with these trends: (1) Sensors are often deployed with varying degrees of redundant coverage--e.g., cameras in buildings often have partially overlapping fields of view, implying that their sensed data are implicitly spatiotemporally correlated, and (2) inferencing increasingly involves the execution of computationally prohibitive machine learning (ML) pipelines (e.g., CNNs for image-based object detection). Executing such deep neural networks (DNNs) gives rise to well-known throughput bottlenecks and prohibitive energy consumption.

We introduce and explore the paradigm of **Collaborative Deep Intelligence that exploits sensor multiplicity** (e.g., a group of networked cameras with overlapping views) for performance benefits such as improved accuracy with minimal overhead of latency. We explore alternate designs for collaboration for the illustrative task of person detection on video sequences [2, 4, 5]: (1) probabilistic fusion of independent deep inferences at the output stage (CNMS) [5], (2) augmentation of the image channels of a reference camera (i.e., the typical RGB) with additional input from collaborator views' past inferences (CSSD) [5], and, more recently, (3) by **adapting the ML pipeline on-the-fly based on lightweight digest sharing between collaborating nodes** [2]**.** For the latter, we do so by (a) first employing a light-weight scene summarization technique that extracts knowledge from collaborators' views, from early hidden layers of the deep pipeline that is shared over the network and (b) injecting the digest back into the later layers of the reference camera's processing pipeline for improved accuracy. Compared to the former alternatives, this design incurs zero re-training of the deep networks, adds minimal overhead to the processing pipeline and the bandwidth required for sharing digests over the network and seamlessly rolls back to the default, non-collaborative operational mode when collaborators are not available or deemed unreliable. This work is thus motivated by a fundamental research question: ***to enable execution on mobile/embedded devices, is it possible to rely on collaboration to replicate the accuracy of very-deep, but high accuracy neural networks while incurring the computational expense only of shallower, older DNNs (with concomitant latency and energy benefits)?*** On benchmark datasets, we recover as high as 80% of the accuracy gain with as few as two collaborators. In ongoing work, I explore notions of **self-configuration and steerability** for collaborative perception systems where cameras learn to identify regions that benefit the most from collaboration, quantify the goodness of collaborators for selective collaboration, and over time, "steer" themselves towards optimal configurations.

3. Dissertation Work: Human Behavior-driven Cyber-Physical Systems

A large part of my PhD research has focused on understanding human mobility and social contexts, both indoors and outdoors, to drive the design of human-centered systems. I used fine-grained, longitudinal mobility observations using a server-side, WiFi-based positioning system, indoors, as well as mass transit transactions, city-wide, to study human mobility patterns [6-9]. By adopting measures from network science, I developed a set of features that captures correlated movement characteristics, to quantify the strength-of-ties between different individuals [6]. As I explain next, I used such passively extracted "physical – social" relationships to reveal key insights into human behavior (e.g., mobility, predictability in routines [10], interruptibility [8], influence of social context [11,12], etc.) to build novel, proactive cyber-physical systems.

We designed **BuScope [9]**, a live mobility analytics platform for smart cities, that generates neighborhood-level insights on high volumes of streaming mobility data, with **O(secs) responsiveness**, to create novel and **"live" smart city services**. We base this work on the observation that a vast majority of public bus trips are predictable and driven by routine commuting patterns which manifests in two aspects: (a) individual-level regularity and (b) aggregate-level conformity. The platform is flexible enough to recompute the analytical insights, at both individual and bus-level specificity, very frequently for peak city-scale workloads—e.g., it incurs 17.33 msecs average latency to process each of ~270,000 boarding and alighting transactions generated by 221,217 commuters on 3777 buses, during a 30-minute peak period. Using realistic neighbourhood-scale simulation models, we show that look-ahead prediction of the number of disembarking passengers allows us to **dynamically pre-position shared, last-mile autonomous taxis** at different bus stops—this can reduce the waiting time experienced by a disembarking customer by 75%, to about 30 secs for single capacity vehicles.

In the past, studies have relied on the predictability of mobility to generate various urban insights. In a complementary effort, I studied **the ability to predict instances of *unpredictability in human mobility to support highly responsive smart building systems [10]***, specifically, to predict moments when individuals will default from their routine behavior. This work describes a framework to detect episodes of future anomalous mobility using an individual's current deviation from his/her routine mobility together with the anomalous behaviour of his/her social ties, with over 90% accuracy even at reasonably long look-ahead times of 4 hours. I demonstrate how proactive optimization of urban operations is possible through the use case of mobility-aware task assignments to crowd-workers on a smart campus where we found that, in comparison to workers with anomalous mobility behaviour, others achieve higher task completion rates.

My colleagues and I deployed the **LiveLabs ecosystem [11]** on campus to study human behaviour under varying hyper-contexts such as location and group, and a **novel class of behavioral experimentation**. The system consists of (a) a passive, server-side system for monitoring indoor mobility of thousands of users who connect to the university's WiFi infrastructure, (b) a suite of mobile applications running on opt-in participants' phones that capture interactions of its participants with their respective phones, and (3) a behavioural intervention engine that allows experimenters such as social scientists to use such contexts to actuate behavioural change and theorize new hypotheses. As an example of the latter, we designed a consumer marketing experiment on top of LiveLabs [12]. We showed that cohesiveness of the group a consumer is part of, at the time of receiving marketing promotions, influences their attitude towards those promotions.

4. Future Research Plans

In the future, I intend to broadly investigate the possibilities around ***dramatically increasing the operational efficiency of future urban environments through a combination of applied ML advances and human***

behavior-driven optimizations. I am excited to explore opportunities for optimizations in the following key problem areas that I believe are central to future cities:

(1) Next Generation Human-Swarm Teaming: With a paradigm shift towards edge-centric applications where low-powered IoT and autonomous devices sense and analyse data at real time, there's an increasing need for optimizing sensing and inference pipelines, especially under extreme SWAP constraints. While my present work explores re-configurations of the ML pipelines of cooperating devices, responding to the behavior of human partners “in-situation”, a next step in human-autonomy teaming is in designing autonomous partners that **continually learn and improve** through positive/negative, implicit reinforcements from human partners. I envision a next generation of teaming where autonomous agents gain fine-grained understanding of the human blueprint on navigating complex tasks (e.g., behavioral cloning of scouting) that they are able to exceed both human and state-of-the-art autonomous performance. A key outcome of this future vision is that increasingly a greater number of autonomous agents and fewer humans will be required for undertaking complex tasks in safety-critical scenarios (e.g., rescue operations in post-disaster response).

(2) Future of Work: While the success of smart environment deployments has traditionally been measured through hard metrics (e.g., wait time reduction of transportation systems), increasingly, there has been interest in softer metrics such as the quality of life of its residents, and productivity as a sign of the health of the economy. With an anticipated boom of personal wearables and IoT devices at workplaces, unobtrusive, continuous and privacy-preserving quantification of wellness and productivity measures, I believe, will be of critical interest for future cities. In addition to offering behavioral insights, a key application of technology would be in enabling in-situ interventions for promoting positive behavioural change – e.g., enabling virtual “casual” interactions in remote teams of white-collar office workers through continually sensed physiology, dynamically altered “personalized” workspaces (e.g., lighting, sound levels) based on implicitly learned productivity–preference relationships for desk-workers, etc.

References

[1] K. Jayarajah**,** A. Gangopadhyay and N. Waytowich. TagTeam: Towards Wearable-Assisted, Implicit Guidance for Human–Drone Teams, Proceedings of SmartWear’22 co-located with ACM MobiCom’22.

[2] K. Jayarajah**,** D. Wanniarachchige, T. Abdelzaher and A. Misra, ComAI: Enabling Lightweight, Collaborative Intelligence by Retrofitting Vision DNNs. Proceedings of the IEEE Conference on Computer Communications, 2022. **INFOCOM’22**.

[3] I. Ghosh, K. Jayarajah**,** N. Roy and N. Waytowich. Memento: Improving Working Memory Recall using Wearable-Assisted Cueing, Under Review, 2022.

[4] T. Abdelzaher, Y. Hao, K. Jayarajah**,** A. Misra, P. Skarin, S. Yao, D. Weerakoon and K. Arzen, Five Challenges in Cloud-Enabled Intelligence and Control. ACM Transactions on Internet Technologies (**ACM ToIT**), Vol. 3, Issue 1, Feb 2020.

[5] D. Weerakoon, K. Jayarajah, R. Tandriansyah and A. Misra, Resilient Collaborative Intelligence for Adversarial IoT Environments. 22nd International Conference on Information Fusion (FUSION), July 2019.

[6] Jayarajah, K.**,** A. Misra, X. W. Ruan, E. P. Lim. Event Detection: Exploiting Socio-Physical Interactions in Physical Spaces. IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, 2015. ***IEEE/ACM ASONAM’15.***

[7] Sen, R., Y. Lee, K. Jayarajah, A. Misra and R. K. Balan. GruMon: Fast and Accurate Group Monitoring for Heterogeneous Urban Spaces. 12th ACM Conference on Embedded Networked Sensor Systems, 2014. ***SenSys14.***

[8] Jayarajah, K.**,** Y. Lee, A. Misra, R. K. Balan. Need Accurate User Behavior? Pay Attention to Groups!. ACM International Joint Conference on Pervasive and Ubiquitous Computing, 2015. ***UbiComp15.***

[9] Meegahapola, L., T. Kandappu, K. Jayarajah, A. Misra, and L. Akoglu. BuSCOPE: Fusing Individual Aggregated Mobility Behavior for Live Smart City Services. 17th Int. Conference on Mobile Systems, Applications and Services. . ***MobiSys19.***

[10] Jayarajah, K**.** and A. Misra. Predicting Episodes of Non-Conformant Mobility in Indoor Environments. PACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT), Vol 2, Issue 4, Article 172, Dec 2018. ***UbiComp19.***

[11] Jayarajah, K., R. K. Balan., M. Radhakrishnan, A. Misra and Y. Lee. LiveLabs: Enabling In-Situ Behavioural Experiments. 14th International Conference on Mobile Systems, Applications and Services, 2016. ***MobiSys16.***

[12] R. Gupta, S. Mukherjee and K. Jayarajah**,** The Role of Group Cohesiveness in Targeted Mobile Promotions, **Journal of Business Research** (Special Issue on Mobile Applications), 2021. Impact Factor 4.87.