Research Statement

"Pervasive Sensing & Machine Intelligence"

Archan Misra
Professor, School of Information Systems,
Singapore Management University
Tel: (65) 6808-5202; Email: archanm@smu.edu.sg
Dated: Dec 28, 2019

1. Background/Overview

At SMU, I focus primarily on techniques that apply machine learning (ML) technologies on sensor data from personal devices (smartphones and wearables) and Internet-of-Thing (IoT) sensors to **infer human behavior, activities and events, as they happen in the physical world**. Such pervasive sense-making enables the emerging vision of "**edge machine intelligence**", where such ubiquitously deployed devices are increasingly capable of performing in-situ human-like inferences from a variety of structured and unstructured (e.g., video, audio) data streams.

In past years, I have focused on first *understanding and quantify commonplace daily lifestyle activities (e.g., shopping at a retail store, playing a game using a VR device or exercising in a gym), and then applying real-time and predictive insights over such activity histories to develop new urban sensing paradigms and services (e.g., crowdsourcing)*. More recently, I have begun to apply such analytics to capturing the operations and activities (of both workers and machinery) in **industrial | factory environments** and **smart spaces,** with a focus on significantly **decreasing the energy consumption** & **computing overhead** involved in such sensing and analytics.

This work has direct relevance to two emerging domains of information-enhanced operations:

- **Smart Cities,** where my work can enable deeper profiling of the *individual* and *collective interactions* performed by citizens with various municipal/urban services (such as transportation, garbage pickup, package delivery, retail shopping, etc.), which in turn can lead to improvements in the provision and operation of such services.
- **Industry 4.0**, where my work can enabler robust and accurate capture of data, about both worker actions/instructions and factory equipment, which in turn can assist real-time optimization of manufacturing operations and interactive, natural human-robot collaboration.

Within this broad theme, my research interests and accomplishments (chronologically highlighted on the next page) can be organized around four technical themes (illustrated in Figure 1):

A) **Mobile & Wearable-based Sensing & Analytics:** This research thread focuses on the judicious use of wearable and infrastructure-based (IoT) sensors, to derive an understanding of the "what, where, when and why" people do, as part of their daily lifestyles. Over the last few years, my interests in this space have migrated to studying the new opportunities that *wearable devices*, such as smartwatches and virtual reality (VR) displays, provide in capturing fine-grained individual activity and deeper insights into the conditions of the ambient environment (e.g., the queuing levels in a food court or the vibration patterns of factory machinery).

B) **IoT-based Machine Intelligence:** As an extension of the theme above, this relatively recent body of research looks at way to efficiently execute deep neural network (DNN) pipelines (which can provide human-like perception capability) on resource-constrained embedded and

IoT devices. Recent interest has focused on 'collaborative machine inferencing', where multiple such pervasive IoT nodes mutually share hidden-layer state of their own DNN pipelines to provide high accuracy but with lower energy consumption and execution latency.

C) **Mobile Crowdsourcing & Other Mobility-Based Urban Services:** This is now a relatively mature body of work, focused on harnessing the collective patterns of urban mobility and behavior to explore new urban services. My current primary interest is on "*urban mobile crowdsourcing*", where the mobility patterns of a large body of individuals is leveraged upon to perform a variety of location-specific tasks (e.g., checking on the cleanliness of a restroom or delivering a package), either in campus environments or at city-scale. Recent interests include developing mechanisms to satisfy individual-specific privacy preferences or provide personalized notifications to enhance greater user participation in such urban crowdsourcing scenarios.

D) **Socio-physical Analytics:** This is an exciting, recent direction of research that looks to combine the insights on physical behavior (often gathered via the techniques of theme "A" above) with analytics of content and interactions on social media channels (such as Twitter and Instagram). Concretely, my current interests are in *(a)* analyzing different social sensing feeds (e.g., Instagram images and Twitter feeds) jointly to detect, localize and characterize events, and thus enhance urban situational understanding; *(b)* combining mobility data (e.g., from buses or taxis) with such social media content to provide deeper policy-level insights into the vibrancy of neighborhoods and private businesses and the anticipated transportation demands.

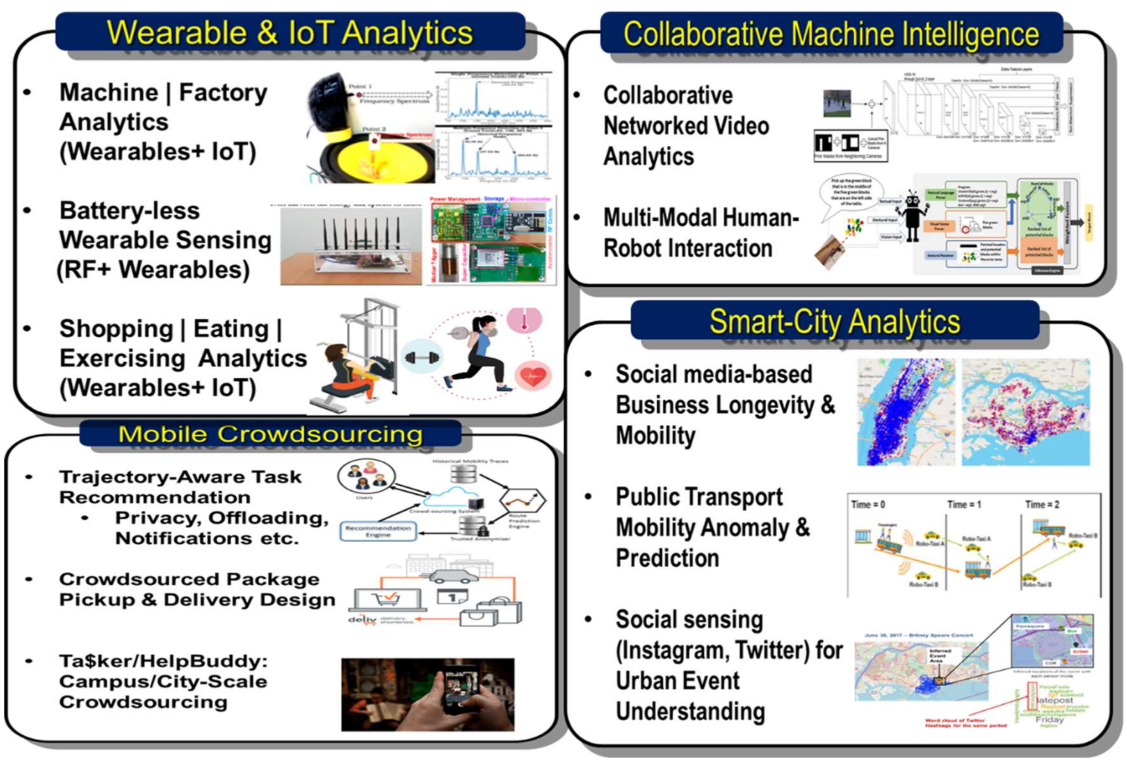

Figure 1. My Research themes and Activities

Broadly speaking, my work aligns with SMU's focus on addressing societal challenges related to "**Advancing Innovation & Technology**" & "**Managing for Sustainability**" areas of excellence and is carried out under the umbrella of 2 major research centers:

- **Living Analytics Research Centre (LARC)** (www.larc.smu.edu.sg) is an inter-disciplinary center for large-scale “Smart Nation” research that focuses on urban applications and services, primarily using online and transactional urban data.

- **Center for Applied Smart-Nation Analytics (CASA)** is a center that increases the translational impact of our school’s research, by embedding advanced technologies in novel “smart nation”-focused public-sector systems and applications.

2. Research Areas

Here is a more detailed enumeration of my research areas and activities.

2.1. Mobile & Wearable Sensing & Analytics

This body of work is based on the premise that sensors on personal mobile and wearable devices offer us an unprecedented capability to capture and understand an individual’s activities, as well as the state of the surrounding world. I’m primarily interested in developing the right analytics techniques, that can best fuse the sensor data available across a distributed set of mobile devices and backend cloud resources, and that can dramatically reduce the energy overheads for such battery-operated devices. The key challenge is to simultaneously support multiple perform metrics (e.g., high accuracy and low-energy overheads), while accommodating significant real-world behavioral diversity. Some of the problems and systems that I’ve worked on in the past 3 years include:

- **Eating & Shopping analytics.** This body of work looks at using wearable devices (such as a wrist-worn smartwatch) to gather fine-grained insights about an individual’s commonplace daily activities, such as shopping and eating. For eating analytics, we have built a fully-automated deployable system called *Annapurna* [http://is.gd/annapurna] [Sen:15] that combines gesture detection using accelerometer and gyroscope sensor data with automatic triggering of the smartwatch-embedded camera to capture a high-quality image of the food being eaten. I have also used both wearables and lightweight IoT devices (specifically BLE beacons) to capture shopper-product interactions in stores. Digitally capturing the shopper’s journey within the store can provide valuable information about her preferences and opinions about products, something that traditional point-of-sale transactional data cannot. Examples include (a) the *IRIS platform* [Radhakrishnan:16], which used a combination of smartphone and smartwatch sensor data to build a shopper’s profile based on inferring a shopper’s micro-gestural activities; and (b) more recently, the *I4S* technique ([Sen:18]) that combines multiple low-energy BLE beacons with smartwatch sensing to localize a shopper’s interactions with an accuracy of +-40cm.

- **Activity Recognition for Smart Gyms:** This body of work has been exploring the use of low-cost sensors to provide real-time capture and understanding of the exercise activities performed by individuals in gyms. In an initial body of work, we utilized a simple sensor device (consisting of an accelerometer and a magnetic sensor) mounted on the plates of a weight-stack machine, which provides information about the motion dynamics of an individual performing weight exercises. Our proposed **W8-Scope system** [Radhakrishnan:20] applies a set of innovative machine learning techniques on such sensor data to infer a variety of exercise-related attributes, such as *the amount of weight lifted, the type of exercise performed and any motion-related mistakes during such exercising*. More recently, we have extended our investigations to free-weight (dumbbell-based exercises). We explored [Radhakrishnan:19] the use of personal ‘earable’ devices (widely used by gym-goers) in providing personalized, quantified insights and feedback. As in-ear sensing by itself is often too weak to pick up exercise-driven motion dynamics, we propose a novel, low-cost system that can monitor multiple concurrent users by fusing data from (a) wireless earphones, equipped with inertial and physiological sensors and (b) inertial sensors attached to exercise equipment.

- **Passive, Battery-Less Wearable Sensing**: One of my long-term goals is to dramatically reduce the energy overheads of wearable sensing, such that, in the extreme case, such wearable sensors operate purely through *energy harvesting* and never need to be recharged! As an early and exciting example of this, we have developed WiWear [Vu:19], a wearable system that includes an accelerometer (to help track natural human gestures) based purely on WiFi energy harvesting. The key innovations include (i) the ability to trigger the wearable sensing functions only on demand (using a motion harvesting sensor), and (b) transmit dynamically beam-formed "power packets" from a multi-antenna AP, such that these packets can deliver significantly greater power (over 500 μW at distances of 3+m) than currently possible. As another exemplar of such passive, battery-less sensing, we have also shown [Jaiswal:18] how an RFID-powered accelerometer tag can be used, in conjunction with infrastructural triggers, to capture human activities indoors. **The future goal of this work is to combine such ultra-low energy wearable devices with infrastructural sensing modes (e.g., short range radars) to support battery-less pervasive sensing in multi-occupant environments (e.g., offices and gyms).**
- **Machine & Industrial Analytics:** This is a new and exciting area of work, where I've been investigating how a combination of wearable (e.g., smartglasses) and IoT devices (e.g., smart LED bulbs or BLE beacons) are innovatively combined to (a) sense fine-grained context (about humans, machinery or products) in industrial environments, such as factories and warehouses., or (b) enable new forms of subliminal communication in smart spaces (such as shopping malls and airport lounges). As an early example, in collaboration with researchers from TCS, I have developed techniques [Roy:18] to use a low-frame rate camera (one that may be mounted on a worker's protective smartglass), in combination with variable-frequency strobing of overhead LED lights, to infer multiple vibration frequencies of an industrial machine. This form of optical sensing offers the ability to derive the health of industrial machines using "visual remote sensing", without needing to mount any sensors on the machine itself. In recent, ongoing work, I am exploring the use of AI/ML techniques to support more robust "screen-camera communication", where additional information content is embedded in the visual content shown on public digital displays (e.g., at shopping malls) such that the information can be decoded by the camera sensor on smartphones/wearables while remaining below the threshold of human perception.

2.2 Mobile Crowdsourcing

Over the last few years, mobile crowdsourcing, where a pool of at-will workers performs location-specific micro-tasks, has created disruptions in many urban services—including *transportation* (e.g., Uber) and *last-mile package delivery* (e.g., Amazon Flex). My research here is driven by a central question: Can such mobile crowdsourcing services be made more effective by better leveraging the predicted movement path and the behavioral preferences of workers? Some of my recent work in this area includes:

- **Trajectory-aware task recommendation.** All existing mobile crowd-sourcing platforms did not *personalize task recommendation* based on the movement of workers. At best, they allowed a worker to search for *nearby tasks*, close to the worker's current location. To address this situation, we pioneered a task recommendation strategy that maximizes the task completion rate while minimizing a worker's *detour* from her routine movement trajectory. We showed [Chen:14] that our proposed centrally-coordinated recommendation approach, which is an interesting variant of the orienteering problem, can result in higher worker productivity (higher rewards per unit detour), an overall higher task completion rate, and improved fairness, even when an individual worker's future movement is uncertain [Chen:15]. More recently, I've been exploring techniques to extend this paradigm to crowd-sourced *pickup-and-delivery tasks* (e.g., last mile package delivery), where task execution requires a worker to visit both (source, destination) locations. In ongoing/future work, we are investigating how the use of limited and differentiated information display (showing different workers different sets of such last-mile delivery tasks) can help better balance the desire for greater worker choice in task selection (which promotes worker satisfaction) and greater overall task execution efficiency (which increases a platform's profitability).

- **Experimental crowdsourcing**. Besides developing theoretical techniques for task recommendation, my work has also looked at developing *empirically-validated*, *behaviorally-driven* crowdsourcing mechanisms. To support such experimental investigations, we have developed and deployed *Ta$ker*, a campus-based experimental mobile crowdsourcing platform, on the SMU campus. Over a 4 year period, Ta$ker has had a loyal pool of around 1,000 student workers, who have completed over 150,000 reporting-oriented tasks, such as "checking on the cleanliness of restrooms" or "reporting on the stock availability in a vending machine." Ta$ker has enabled us to develop and empirically validate a variety of crowd-sourcing related technologies. For example, in [Kandappu:16], we demonstrated novel behavioral aspects, including a preference of workers for task bundles that minimize overall detour even if it results in lower per-task payout. Similarly, I have shown [Kandappu:17] that allowing workers to dynamically offload some of their tasks to designated friends can significantly improve the task completion rate. I have developed mechanisms [Kandappu:18] that support privacy-aware crowd-sourcing, by allowing workers to intelligently obfuscate their reported location trajectories without requiring any trusted 3rd party. More recently, I have shown [Kandappu:20] how appropriate context-aware notifications, reminding people of available tasks, can help not just improve the overall task acceptance/execution rates, but also shape the spatiotemporal attributes of such task execution.

  These ideas have had both academic and practical impact: *we have worked extensively with public agencies in Singapore to embed these crowd-sourcing concepts into a new, city-scale mobile crowdsourcing application (called HelpBuddy) that supports greater government-citizen and citizen-citizen engagement.* HelpBuddy has been successfully piloted with 8 different Singapore government agencies and a participant pool of ~4000 resident volunteers.

2.3 Smart City Analytics

The third aspect of my research involves *socio-physical analytics* for urban event detection and understanding. Here, I move beyond studying *individual-level activities and behavior* to using collective observations of such activities and behavior to understand events in urban spaces. The scale of these events can vary, from individual public venues (e.g., a college campus) to larger geographic areas such as city neighborhoods.

- **Social Media-based Business Survivability and Mobility**: In this area of work, I focus on using a combination of social media and urban mobility data to understand the impact of people's movement on neighborhood businesses and amenities (e.g., parking demand). As a concrete and innovative example, we collaborated with researchers from University of Cambridge to predict the 6-month survivability of individual retail businesses using a combination of Foursquare check-in data and aggregated taxi usage data. The key insight [D'Silva:18] is that social mobility data can be used to uncover features characterizing both the competitive profile of an individual business and the neighborhood in which it operates. We showed that we can use such features to predict such business survivability with an AUC (area-under-curve) value of 0.86. In follow-on work, we tried to systematically understand how land use impacts crowd flow and transportation demand to different areas of the city, and in turn, how the influx of people to an area (or lack thereof) can influence the viability of business entities in that area. In preliminary work, we applied unsupervised learning techniques to show how the temporal demand pattern for carparks was correlated to the land-use mix (the nature of local businesses in a particular area). Transforming the problem of predicting carpark temporal demand as a multi-class classification task, we showed [Jayarajah:18] that we can achieve of AUC of 0.84 by exploiting features related to land-use mix extracted from social media data. Our work lays the foundation for using such mobility and social media data to study other aspects of a neighborhood and its constituent businesses and residents.

- **Public Transport Mobility Anomaly & Prediction:** My work leverages on the increased availability of digitally captured data on individual commuting behavior (e.g., data of smart card tap-ins and tap-outs on buses and trains) to obtain improved predictive insights on both collective patterns of urban mobility and the neighborhood events that often underpin such behavior. In a recent work we demonstrate the power of applying real-time, predictive analytics on the smart-card generated trip data of millions of public bus commuters in Singapore, to create two novel and "live" smart city services. Our work [Meeghapola:19a] combines two aspects of urban mobility: (a) conformity: which reflects the predictability in the aggregated flow of commuters along bus routes, and (b) regularity: which captures the repeated trip patterns of each individual commuter. The resulting *BusScope* platform provides O(mins) lookahead into the number of disembarking passengers at neighborhood bus stops; it achieves over 85% accuracy in predicting such disembarkations at each bus stop. By moving driverless vehicles proactively to match this predicted demand, we can reduce wait times for disembarking passengers by over 75%. Similarly, by using outlier measures of currently operating buses, we can detect and spatiotemporally localize dynamic urban events, as much as 1.5 hours in advance, with a localization error of 450 meters. While individual-specific transaction records (such as smart card (tap-in, tap-out) data or taxi trip records) hold a wealth of information, these are often private data available only to the service provider (e.g., taxicab operator). In parallel work, we have explored the utility in harnessing publicly available, albeit noisy, transportation datasets, such as noisy "Estimated Time of Arrival" (ETA) records (commonly available to commuters through transit Apps or electronic signages). In [Meeghapola:19b], we demonstrate how to develop a clustering-cum-deduplication mechanism to accurately reconstruct the transit times of individual bus instances at different bus stops from such coarse-grained, anonymized bus ETA records, achieving precision/recall values of inferring such arrivals with an error of less than +-1 minute. As a practical application, we show how such accurate reconstruction can help provide more accurate predictions of future bus arrivals (with an error of less than +-20 seconds) compared to traditional offline estimation techniques based on historical records.
- **Social Sensing for Urban Event Understanding:** This work tries to discover unexpected or latent events in public venues, based on the individual and collective patterns of movement a at these locations and content posted on social media channels, such as Twitter or Instagram, to better detect urban events. To detect such events, I apply information-processing tools on both the *metadata* (e.g., the location tag of an Instagram post, or the total number of Tweets with a specific hashtag) and the *data/content* (e.g., the objects in an Instagram image). A recent and concrete exploration of such multi-modal sensing is [Jayarajah:16b], where we show how to consider a larger urban event, such as a marathon, and identify and localize *micro-events* (start and finish sequences) based on Instagram posts.

2.4 Collaborative Machine Intelligence

This is my newest, most-recently initiated research thread. The goal of this work, broadly, is to enable ubiquitous intelligence in a future world of connected sensing and computing devices, seamlessly embedded in our surroundings. In my vision, embedded IoT devices (or "things", such as assistive robots or agents) will be capable of human-like interactions with their environment, including speech recognition, vision, and gesture understanding. These capabilities, articulated in our vision of machine intelligence being a composable service for pervasive applications [Yao:19], will bring about such features as verbal device control, (soft) user authentication, and gesture-based human machine communication. The key challenge, of course, is to somehow simplify the execution of complex Deep Neural Network (DNN)-based inferencing pipelines (which represent the state-of-the-art in machine-based perception) so that they can adhere to the resource constraints of embedded devices.

- **Collaborative Networked Video Analytics:** Many IoT networks involve the deployment of resource-constrained sensors with varying degrees of redundancy/overlap (i.e., their data streams possess significant spatiotemporal correlation). To tackle the performance challenges, we advocate the vision of Collaborative IoT Intelligence, where the inferencing pipelines of multiple individual devices share features and "internal state" in real time with one another, allowing the devices to collectively both overcome their individual processing bottlenecks and improve their sense-making fidelity. In current work, we investigate such collaborative intelligence for the specific case of a multi-camera, campus-scale video sensing network, where the cameras are tasked with executing people counting algorithms to collectively provide a "live" view of the occupancy levels in different parts of the campus. We have developed [Weerakoon:19] two different & initial approaches for such collaborative analytics, one called CNMS which refines the statistical output from the underlying DNNs and the other called CSSD which develops entirely new DNN models to take advantage of such collaboration. Using a benchmark video monitoring dataset, we have shown [Weerakoon] that collaborative inferencing results in significantly higher accuracy (75.5% for CNMS and 82% for CSSD), compared to a non-cooperative baseline value of 68.03% achieved by the state-of-the-art SSD object detection DNN. In ongoing work, we have been developing techniques to reduce the execution latency and the network overheads of sharing such 'state' across a wireless infrastructure. Very specifically, we have shown how using selected feature maps (fMaps) from the early convolutional layers of the DNN pipeline can be used to (a) reduce the total processing time for each image frame by over 90%, thereby allowing such collaborative techniques to operate on high-frame rate (30fps or higher) video streams and (b) extract only a small set of features that need to be shared with other peer camera sensors, thereby achieving a significant reduction in the communication overheads of such collaborating cameras.
- **Multi-Modal Human-Robot Interaction**: Advances in sensing, machine learning and robotics are ushering in new forms of human-agent interaction that will transform multiple industries and businesses. Today's state-of-the-art human-agent interactions (such as text-based 'chatbots' or voice-based 'home assistants') are typically unimodal, utilizing a single sensing stream. I predict that increasing penetration of embedded sensors in consumer devices will soon permit significantly richer forms of multi-modal interactions between humans and agents, in real-world venues such as shopping malls, office receptions and manufacturing floors. For example, a worker on a factor floor should be able to instruct a robot to pick up a specific part, via natural interaction-based inputs (such as voice, vision and gestures). To achieve this, we will need advances in two dimensions: (a) improved multi-modal ML models that combine inputs from multiple perceptual sensor streams, and (b) reducing the computational complexity of such ML inferencing pipelines to make them executable on resource-contained devices. In ongoing, early work, we have demonstrated the possibility of using pointing gestures, a naturally-generated additional input modality, to improve the multimodal comprehension accuracy of human instructions to robotic agents for collaborative tasks. We have developed M2Gestic, a system that combines neural-based text parsing with a novel knowledge-graph traversal mechanism, over a multi-modal input of vision, natural language text and pointing. Via multiple studies related to a benchmark table top manipulation task, we show that (a) M2Gestic can achieve close-to-human performance in reasoning over

unambiguous verbal instructions, and (b) utilize pointing input (even with its inherent location uncertainty) to achieve a significant (~ 30%) accuracy improvement when verbal instructions are ambiguous. In future work, we shall be looking at ways to use collaboration across the different ML pipelines (one for each input modality) to dramatically improve the accuracy-vs-latency tradeoff for such machine inferencing tasks.

3. Future Research Directions and Interests

In the next 1-2 years, I will continue my broad focus on socio-physical analytics, especially focusing on the ability to combine *wearable* & *infrastructural IoT* data to obtain a deeper understanding of human activities and interactions in a variety of spaces with dramatically reduced energy overheads. I find three directions of work particularly promising.

- **Battery-free Wearable + IoT systems & sensing.** In recent years, there has been significant interest in *device-free localization and activity recognition* research, which uses technologies such as Wi-Fi and light to capture human movement. While promising, I believe that significantly better accuracy and energy-efficiency can be achieved by combining such device-free techniques with intermittent *battery-free wearable sensors.* Achieving this vision will require innovative designs of both embedded sensing platforms and sensor data analytics *at the edge.* This approach is likely to be useful for industrial IoT (*Industry 4.0*) technologies, and will likely involve the fusion of new forms of passive sensing technologies (e.g., indoor radar, 60 GHz WiFi) and ultra-low power wearable platforms. I have also been investigating the use of virtual-reality (VR) and augmented reality (AR) wearables for novel systems and applications. Early examples include the use of video+image cues, overlaid on an AR device, for infrastructure-independent wayfinding [Roy:17], and the combination of real-world sensing and VR devices to support experiential understanding of App usage under various impairments [Kim:18].
- **Collaborative Edge Inferencing for Smart City Operations**. Cities continue to roll out *smart city infrastructures*, with tens of millions of heterogeneous IoT devices, to support applications, such as dynamic bus route optimization and adaptive street lighting. To provide real-time services, and mitigate the backhaul data demands, it is important to develop "edge intelligence"—i.e., embed various types of inferencing capabilities on the resource-limited edge devices (e.g., cameras mounted on lampposts). To overcome the resulting latency and throughput bottlenecks, I shall develop an enhanced "cognitive edge" framework, where devices collaborate and share the intermediate state of machine learning pipelines with one another to improve system performance and resiliency to faults, while seeking to minimize the network bandwidth needed to support such real-time collaboration. These collaborative principles should also prove relevant to my work on human-robot/agent interaction, so as to enhance the agent's ability to perform sophisticated, energy-efficient comprehension of human instructions.
- **Socio-physical urban analytics**. I remain excited by the possibility of combining personal sensing (via mobile devices) with additional data from (1) urban informatics portals or (2) public social media sources. At present I am using urban transportation-related data sources (bus movement, parking garage occupancy, etc.) available from the Singapore Government's data.gov.sg portal, in combination with social media feeds (e.g., Twitter, Facebook and Foursquare) to profile the usage/viability of individual establishments, identify the impact of such usage (driven by overall land use) on aggregated transportation demand and create new spatiotemporal models of "retail competition".

Selected Publications and Outputs

[Chen:14] C. Chen, S-F. Cheng, A. Gunawan, **A. Misra**, K. Dasgupta and D. Chander, *TRACCS: Trajectory-Aware Coordinated Urban Crowd-Sourcing*, 2nd AAAI Conference on Human Computation and Crowdsourcing (HCOMP), November 2014.

[Chen:15] C. Chen, S-F. Cheng, H-C. Lau and **A. Misra**, *Towards City-scale Mobile Crowdsourcing: Task Recommendations under Trajectory Uncertainties*, International Joint Conference on Artificial Intelligence (IJCAI), July 2015.

[D'Silva:18] K. D'Silva, K. Jayarajah, A. Tassos, **A. Misra**, and C. Mascolo. *The Role of Urban Mobility on Retail Business Survival*, PACM IMWUT, Volume 2(3), September 2018.

[Jaiswal:18] D. Jaiswal, A. Gigie, T. Chakravorty, A. Ghose and **A. Misra**, *Table of Interest: Activity Recognition and Behaviour Analysis Using a BatteryLess Wearable Sensor*, Workshop on Wearable Systems and Applications (WearSys), June 2018.

[Jayarajah:15] K. Jayarajah, S. Yao, R. Mutharaju, **A. Misra**, G. de Mel, J. Skipper, T. Abdelzaher, M. Kolodny, *Social Signal Processing for Real-time Situational Understanding: a Vision and Approach*, 1st International Workshop on Social Sensing (SocialSens 2015), October 2015.

[Jayarajah:16b] K. Jayarajah and **A. Misra**, *Can Instagram Posts Help Characterize Urban Micro-Events?*, 19th International Conference on Information Fusion (Fusion), July 2016.

[Jayarajah:18] K. Jayarajah, A. Tan and **A. Misra**, *Exploiting the Interdependency of Land Use and Mobility for Urban Planning*, 7th International Workshop on Pervasive Urban Applications (PURBA), UbiComp'18 Adjunct Proceedings. [**Best Paper Award**]

[Kandappu:16] T. Kandappu, et al, *TASKer: Behavioral Insights via Campus-based Experimental Mobile Crowd-sourcing*, ACM International Joint Conference on Pervasive and Ubiquitous Computing (Ubicomp), September 2016.

[Kandappu:17] T. Kandappu, **A. Misra** and R. Daratan, *Collaboration Trumps Homophily in Urban Mobile Crowd-sourcing*, ACM Conference on Computer Supported Collaborative Work and Social Computing (CSCW), March 2017.

[Kandappu:18] T. Kandappu, **A. Misra**, S-F. Cheng, R. Tandriansyah and H-C. Lau, *Obfuscation At-Source: Privacy in Context-Aware Mobile Crowd-Sourcing*, Proceedings of ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT), Vol. 2(1), March 2018.

[Kandappu:20] T. Kandappu, **A. Misra**, A. Mehrotra and M. Musolesi, *PokeME: Applying Context-Driven Notifications to Increase Worker Engagement in Mobile Crowd-sourcing*, to appear, ACM SIGIR Conference on Human Information Interaction & Retrieval (CHIIR), March 2020.

[Kim:18] W. Kim, K. Choo, Y. Lee, **A. Misra** and R. Balan, *Empath-D: VR-based Empathetic App Design for Accessibility*, ACM Mobisys, June 2018.

[Meeghapola:18a] L. Meeghapola, T. Kandappu, K. Jayarajah, L. Akoglu, S. Xiang and **A. Misra**, *BuSCOPE : Fusing Individual & Aggregated Mobility Behavior for "Live" Smart City Services*, ACM Mobisys, June 2019.

[Meeghapola:18b] L. Meeghapola, N. Athaide, K. Jayarajah, S. Xiang and **A. Misra**, *Inferring Accurate Bus Trajectories from Noisy Estimated Arrival Time Records*, ITSC'19, October 2019.

[Radhakrishnan:16] M. Radhakrishnan, S. Eswaran, **A. Misra**, D. Chander and K. Dasgupta, *IRIS: Tapping Wearable Sensing to Capture In-Store Retail Insights on Shoppers*, 14th IEEE International Conference on Pervasive Computing and Communication (PerCom), March 2016.

[Radhakrishan:19] M. Radhakrishnan and **A. Misra**, *Can Earables Support Effective User Engagement during Weight-Based Gym Exercises?*, EarComp 2019 Workshop (co-located with ACM UbiComp/ISWC), Sep 2019

[Radhakrishan:20] M. Radhakrishnan, **A. Misra** and R. Balan, *W8-Scope: Fine-Grained, Practical Monitoring of Weight Stack-based Exercises*, to appear, 18th IEEE International Conference on Pervasive Computing and Communication (PerCom), March 2020.

[Roy:17] Q. Roy, S. Perrault, S. Zhao, A. Vanniyar, R. Davis, V. Vechev, Y. Lee and **A. Misra**, *Follow-My-Lead: A Leader-Follower Approach to Visual Indoor Navigation*, ACM Conference on Human Factors in Computing Systems (CHI), May 2017.

[Roy:18] D. Roy, A. Ghose, T. Chakravarty, S. Mukherjee, A. Pal and **A. Misra**, *Analysing Multi-Point Multi-Frequency Machine Vibrations using Optical Sampling*, 1st International Workshop on Internet of People, Assistive Robots and ThingS (IoPARTS), June 2018.

[Sen:15] S. Sen, S. Vigneshwaran, **A. Misra**, R. Balan and Y. Lee, *The Case for Smartwatch-based Diet Monitoring*, 1st International Workshop on Workshop on Sensing Systems and Applications Using Wrist Worn Smart Devices (WristSense), 03/2015. (**Best Paper Award**)

[Sen:18] S. Sen, **A. Misra**, V. Subbaraju, K. Grover, M. Radhakrishnan, R. Balan and Y. Lee, *I4S: capturing shopper's in-store interactions,* ACM International Symposium on Wearable Computers (ISWC), October 2018.

[Tran:19] H.V. Tran, **A. Misra,** J. Xiong and R. Balan, *WiWear: Wearable Sensing via Directional WiFi Energy Harvesting*, to appear, 17th IEEE International Conference on Pervasive Computing and Communication (PerCom), March 2019.

[Weerakoon:19] 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.

[Yao:19] S. Yao, Y. Hao, Y. Zhao, A. Piao, H. Shao, D. Liu, S. Liu, S. Hu, D. Weerakoon, S. Jayarajah, **A. Misra** and T. Abdelzaher, *Eugene: Towards Deep Intelligence as a Service*, 39th IEEE International Conference on Distributed Computing Systems (ICDCS), July 2019.