YONGSUNG KIM | RESEARCH STATEMENT

The goal of my research is to develop socio-technical models and intelligent systems that transform how people and their communities can advance their individual and collective goals through people's busy lives. I am broadly interested in designing technologies that do not enforce rigid ways of working and collaborating but that instead surface flexible ways to advance goals of interest as opportunities arise. Towards this goal, my dissertation work designs, builds, and empirically evaluates intelligent systems that opportunistically recruit members in a community to make small, convenient contributions to flexibly and efficiently achieve collective goals through their existing routines. My research spans the fields of human-computer interaction (HCI), artificial intelligence (AI), social and crowd computing, and decision science.

Decades of advances in social and crowd computing have made it much easier for people to contribute to solving local and communal problems. This has led to volunteer-based services that recruit ordinary citizens to report local issues and collect data for science whenever and wherever they'd like, and to paid services that coordinate the efforts of gig workers to provide efficient transportation and delivery services on-demand. Despite successful applications, **a core remaining challenge is combining the benefits of *flexibility* in volunteer-based services with the benefits of *intelligent coordination* in paid services.** While the flexibility provided by volunteer-based services can attract more people to make convenient contributions, it can be challenging to reach desired collective goals with uncoordinated contributions. For example, a volunteer-based lost-and-found service may ask community members to look for a lost item anywhere they'd like along their route, but this ease and convenience makes it difficult for the service to ensure good search coverage. On the flip side, most gig workers have little flexibility in choosing when and what work to contribute to while they are "online," because the services that coordinate their contributions rely on their accepting the tasks they are given, without regard for their individual goals and situation.

My research addresses this core challenge by introducing principles, models, and intelligent systems for ***flexible coordination*, or ways to opportunistically engage people to contribute convenient actions that nevertheless are indirectly coordinated to achieve globally effective outcomes**. To enable this idea, I developed decision-theoretic frameworks that (1) model people's routines and activities, (2) account for individual and collective goals, and (3) find moments to connect people with opportunities where the performed actions are effectively coordinated to advance desired goals. Unlike existing approaches that can either provide individual flexibility or efficiently achieve system goals but not both, my approach simultaneously reasons about both the needs of potential helpers and the goals of the system to generate solutions that are convenient for helpers and globally effective for the system.

My dissertation uses the idea of flexible coordination to advances the design of *on-the-go crowdsourcing systems* that elicit convenient contributions that seamlessly fit into people's existing routines to solve local, communal problems, such as delivering packages or finding lost items within a community. I detail in this statement two mechanisms for flexible coordination my work contributes: (1) Opportunistic Hit-or-Wait, a decision-theoretic framework for *individual-level coordination* that surfaces opportunities for helpers to make valuable, convenient, and coordinated contributions on-the-fly to improve the quality of service; and (2) Opportunistic Supply Management, a decision-theoretic framework for *community-level coordination* that optimizes the balance between the experience of volunteers and the goals of the system. Beyond describing the design and implementation of these frameworks, I will share findings from *simulation studies* and *real-world deployment studies* that illustrate how flexible coordination allows people to effectively help others in their community while still attending to their own needs and goals.

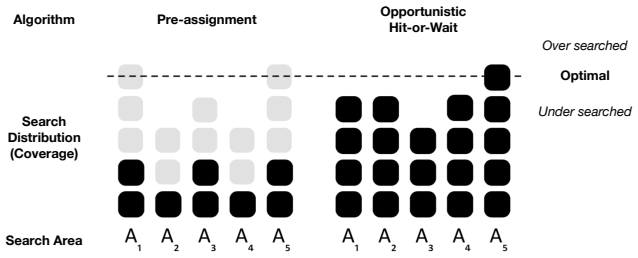

Figure 1: An illustration of search distribution coverage achieved through a baseline approach that pre-determines which tasks to surface to a helper vs. our Hit-or-Wait approach. The baseline approach can lead to many missed opportunities (grey boxes) as people may or may not reach pre-assigned tasks. Hit-or-Wait achieves near-optimal outcomes by avoiding missed opportunities while still directing contributions to where they are most needed.

Individual-Level Coordination: Opportunistic Hit-or-Wait

Designing effective on-the-go crowdsourcing systems requires *individual-level coordination* mechanisms that promote contributions from a potential helper when it is convenient for them *and* where their contribution is most needed for advancing collective goals. For example, a community-based lost-and-found service may only ask community members to look for lost items along their existing route, but still try to ensure good search coverage by controlling when and where to ask for their help. Given uncertainty in participation and in a potential helper's future trajectories, it is challenging to decide when and which tasks to notify a helper about among the many tasking needs a potential helper may encounter. On one hand, waiting for the helper to encounter a task where their help is most needed is not useful if they might never come across the task on their route. On the other hand, sending a helper the first task they come across can miss out on opportunities for them to contribute to a more needed task that they will encounter later. This problem illustrates the general challenge in trying to provide flexibility and convenience while also attempting to coordinate and direct efforts to where they are most needed.

To overcome this challenge, I introduce *Opportunistic Hit-or-Wait* [4], a decision-theoretic framework that coordinates opportunistic contributions by deciding when and whether to engage a helper with a task. Hit-or-Wait models helpers' changing patterns of mobility and the value of their tasking contributions toward a collective goal. Using this model, it makes on-the-fly decisions about whether to notify a helper of a task right now or to wait for better opportunities in the future. To do this, I model a sequence of Hit-or-Wait decisions with a Markov Decision Process (MDP) that considers the value of notifying a helper of a current task now (i.e. hitting) with the expected value from waiting and making subsequent hit-or-wait decisions across possible locations and tasks the helper may later encounter. Unlike existing task assignment mechanisms that unnecessarily pre-determines who should do what tasks—which results in missed opportunities—our approach generates strategies that are custom-tailored to a helper's changing availability, which effectively preserves volunteer flexibility while still reaching effective outcomes; see Figure 1.

Results from a simulation study and a field deployment study (N=25) of community-based lost-and-found showed that Hit-or-Wait collected small, opportunistic contributions along helpers' routes that nevertheless achieve near-optimal outcomes when compared to an optimal solution that has full knowledge of their trajectories. Specifically, our real-world deployment results show that Hit-or-Wait achieves this by making effective wait decisions to direct contributions to where they are most needed and avoids prematurely engaging helpers with less valued tasks. **This result highlights how in-**

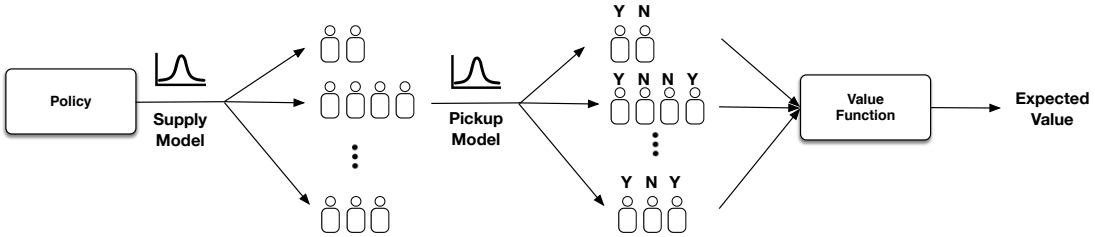

Figure 2: An illustration of how the opportunistic supply management framework simulates the possible unfoldings of adopting a task notification policy to determine its value for achieving intended system goals and desired helper experiences. Supply management uses this workflow to evaluate the value of each policy to determine the optimal policy.

telligent systems like Hit-or-Wait can provide flexibility to individuals *and* still achieve desired collective goals. In other words, by constantly monitoring potential helpers' current and future availability, and accounting for both helper and system needs, individual-level mechanisms for flexible coordination such as Hit-or-Wait can take helpers' routines and goals into account to promote convenient contributions where they are most useful for advancing collective goals.

Community-Level Coordination: Opportunistic Supply Management

On-the-go crowdsourcing systems also need *community-level coordination* mechanisms that optimize the balance between system efficiency *and* the experience of helpers. For example, a community-based package delivery service may try to deliver packages in a timely manner while avoiding overdisrupting or overburdening the helpers that it recruits. Given uncertainty in helpers' availability and their willingness to accept tasks based on their individual situations, it is challenging to set an effective *task notification policy*, or a set of conditions that determine when, where, and whom to notify of tasks. An effective system must manage the tradeoffs imposed by being too aggressive in recruitment—which can be overly disruptive and result in a low task pickup rate—and being too restrictive in recruitment—which can involve too few helpers, overburden the ones that are involved, and leave a disproportionately large number of task demands unfulfilled [6]. This problem illustrates the general challenge in recruiting helpers across a community in ways that preserve flexibility and provide good experiences for helpers, while still achieving the intended goals of the system.

To overcome this challenge, I introduce *Opportunistic Supply Management* [5], a general decision-theoretic framework for modeling and optimizing the choice of task notification policies. Supply Management integrates models that describe how task notification policies affect the available supply of helpers and their likelihood to accept tasks, and how that in turn affects system efficiency and helpers' experience. Using these models, Supply Management simulates the possible outcomes that may result from adopting a task notification policy and choose an optimal policy for a given situation (or set of situations) that best achieves intended system goals and desired helper experiences in expectation; see Figure 2. With this approach, Supply Management can reason about how the world will unfold, take into consideration people's availability and willingness to help, and devise custom-tailored strategies that adapt to changing situations without ever imposing on what each individual must do. Unlike existing task recommendation mechanisms that only consider how each individual can best contribute to the system, our approach considers how to best leverage helper efforts across the community to best meet system goals in ways that still ensure good helper experiences by not overburdening or disrupting potential helpers.

Results from a simulation study and a field deployment study (N=26) of community-based task delivery showed that Supply Management set task notification policies that effectively optimized the balance between task completion and helper disruption that are custom-tailored to the changing levels of demand and the availability of potential helpers. Moreover, our real-world deployment results show that Supply Management can encode and effectively promote specific goals and values a community may hold, such as emphasizing the need to avoid overdisrupting helpers while still attempting to effectively complete tasks. **This result highlights how intelligent systems like Supply Management can find effective 'goldilocks zones' that advance and balance multiple goals that a community cares about.** In other words, by reasoning about tasking demands and the availability of potential helpers, community-level mechanisms for flexible coordination such as Supply Management provide principled ways for shaping how a community can flexibly engage with tasks in ways that fully account for its impact on the service that the community provides and on the experience of helpers who provide that service.

Research Agenda

I am driven to empower people and communities to achieve their goals with the help of intelligent systems that address the need for coordination and the need for flexibility. In communities, workplaces, and people's personal lives, I am interested in developing better ways for intelligent systems and people to work together that transform how we get things done, help one another, and generally balance multiple, competing goals (e.g., individual and collective goals; short-term and long-term goals) that demand our attention across our busy lives. While there are many research directions that interest me, I describe a few concrete directions below.

Opportunistic Help-Seeking and Help-Giving in Workplaces. While my research in on-the-go crowdsourcing developed ways to flexibly and effectively advance collective goals through people's *physical routines* (e.g., see also [2, 1]), in future work I am particularly interested in developing ways to advance help-seeking and collaboration in workplaces through people's *work routines.* Flexible coordination mechanisms for the workplace might help identify opportunistic moments for helping others and for collaborating on key tasks, while still ensuring that people can be productive and work towards their individual goals. I expect that many of the frameworks I developed for on-the-go crowdsourcing can translate and be used in the work setting, but where models of people's mobility patterns and routines are replaced with models of people's work schedules, routines, and needs. Once developed, these models can be used by individual-level coordination mechanisms to flexibly decides when to best recruit a person to seek or provide help (or to continue attending to their own tasks and goals), and by community-level coordination mechanisms to decide how to best orchestrate work and collaborations across a team or organization.

Personal Virtual Assistants to Balance Short-Term and Long-Term Goals. My research thus far has focused on developing flexible coordination approaches for advancing and balancing individual and collective goals. My future work will also develop flexible coordination approaches for helping an individual advance and balance between their short-term and long-term goals (e.g., meeting an impending deadline *and* living a healthy life). Specifically, I am interested in developing intelligent personal virtual assistants that can reason about various short-term and long-term tradeoffs in the decisions that a person might make, and that can surface such reasoning to users to support their everyday lives. Towards this direction, my past work [3] studied the "privacy cost" of accidental information disclosure in notifications, which shed light on the importance of reasoning about the trade-offs between short-term benefits of immediate information access and long-term privacy cost (e.g., based on the content of the notification and the context in which the notification appeared). In

this vein, I am broadly interested in developing principles, models, and intelligent systems that can help people reason about such complex tradeoffs and their implications on the present and future.

Improving Human-AI Communication and Collaboration. My work on flexible coordination is part of a larger research agenda on advancing the design of intelligent systems that put people's needs and goals at the forefront. Making significant progress on this agenda will require intelligent systems that not only act on encoded goals, but that allow for better two-way communication between the human and the AI system. One direction I am interested in is setting clearer expectations [7] and surfacing the rationale behind complex decisions. For instance, as part of the Hit-or-Wait field deployment [4], we studied the extent to which visualizations can be used to communicate seemingly opaque notification decisions to helpers in ways that help them better grasp the value of their contributions. Another direction I am interested in is enabling more active communication between the human and the AI in the process of defining, refining, and changing goals over the course of an activity. For example, flexible coordination mechanisms may be extended to provide a richer language for people to communicate their goals to the system, and to actively shape how they'd like the intelligent system to notify them as their goals change. This can enable intelligent systems that can better serve the (actual) needs and goals of its users, while giving users more freedom and control over how its goals can be met with the support of intelligent systems.

References

[1] Kapil Garg, **Yongsung Kim**, Darren Gergle, and Haoqi Zhang. 2019. 4X: A Hybrid Approach for Scaffolding Data Collection and Interest in Low-Effort Participatory Sensing. In *Proceedings of the ACM on Human-Computer Interaction (CSCW '19)*.

[2] Katherine Lin, Henry Spindell, Scott Cambo, **Yongsung Kim**, and Haoqi Zhang. 2016. Habitsourcing: Sensing the Environment Through Immersive, Habit-Building Experiences. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*.

[3] **Yongsung Kim**, Adam Fourney, and Ece Kamar. 2019. Studying Preferences and Concerns About Information Disclosure in Email Notifications. In *The World Wide Web Conference (WWW '19)*.

[4] **Yongsung Kim**, Darren Gergle, and Haoqi Zhang. 2018. Hit-or-Wait: Coordinating Opportunistic Low-effort Contributions to Achieve Global Outcomes in On-the-go Crowdsourcing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*.

[5] **Yongsung Kim**, Darren Gergle, and Haoqi Zhang. 2020. Opportunistic Supply Management: A Decision-theoretic Approach to Balance Helper Needs and System Efficiency in On-the-go Crowdsourcing. Working paper.

[6] **Yongsung Kim**, Emily Harburg, Shana Azria, Aaron David Shaw, Elizabeth M Gerber, Darren Gergle, and Haoqi Zhang. 2016. Studying the Effects of Task Notification Policies on Participation and Outcomes in On-the-go Crowdsourcing. In *Proceedings, The Fourth AAAI Conference on Human Computation and Crowdsourcing (HCOMP '16)*. AAAI.

[7] **Yongsung Kim**, Aaron Shaw, Haoqi Zhang, and Elizabeth Gerber. 2017. Understanding Trust Amid Delays in Crowdfunding. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*.