

Intent-Aware Interactive Internet of Things for Enhanced Collaborative Ambient Intelligence

Chaoran Huang and Lina Yao , University of New South Wales, Sydney, NSW, 2052, Australia

Xianzhi Wang , University of Technology Sydney, Ultimo, NSW, 2007, Australia

Quan Z. Sheng , Macquarie University, Macquarie Park, NSW, 2109, Australia

Schahram Dustdar , TU Wien, 1040, Vienna, Austria

Zhongjie Wang  and Xiaofei Xu, Harbin Institute of Technology, Harbin, 150001, China

The Internet of Things (IoT) enables the connection of a broad range of artifacts with advanced sensory technologies and produces massive amounts of data to support ambient intelligence. While the potential of IoT systems is widely recognized, little work has demonstrated a system with the ability to execute autonomously in the real world. Inspired by the success of robotics in specialized IoT environments, we propose an end-to-end solution for a generic, interactive ambient intelligence system where robots can assist humans in conducting activities in IoT-enabled smart homes. We evaluate the solution using implementations of public benchmarks on open-source platforms and use several activities to demonstrate the effectiveness of the proposed solution in real life.

The Internet of Things (IoT) promises to integrate digital and physical worlds by connecting artifacts and building networks of them. Traditionally, IoT devices work on their own, due to incompatible techniques or protocols. The Web of Things (WoT)^{a,b} provides IoT applications with standardized descriptions of actions, events, and properties of things to extend Web protocols, thereby supporting broader interactions between smart objects. With plenty of commercially available IoT products, recent research and industry have successfully instrumented connected things in our daily lives. However, these

connected things still cannot work fully autonomously but require human intervention.⁹ The existing research primarily focuses on recommending the next actions to take, yet cannot control devices directly to execute those actions. Inspired by multiagent systems, Internet of Robotic Things (IoRT) enables robots to gain awareness of the environment and adapt their actions. However, the limited related research focuses on specialized scenarios,¹ such as workplaces and smart cities; it still faces challenges for conducting physical interactions for daily tasks. There remains a significant gap to bring robotics into home applications to support streamlined physical human-machine interactions.

We envision a scenario (Figure 1) to demonstrate our goal toward ambient intelligence and to motivate this study. Bob is preparing a bowl of cereal as breakfast, and usually, he finishes the meal with a cup of tea. The system detects a series of events in the kitchen at the time, such as movements of a cabinet door, a bowl, the fridge door, cereal containers and such, thereby recognizing that Bob is preparing cereal as breakfast. Based on Bob's daily routine, the system envisages a

^a[Online]. Available: <https://www.w3.org/TR/wot-thing-description/>

^b[Online]. Available: <https://www.w3.org/TR/wot-architecture/>




FIGURE 1. An envisioned scenario of ambient intelligence for smart homes.

tea to finish the breakfast and commands a robot with a 3D depth camera to move to the kitchen. The robot then identifies a mug and teaspoon on the bench, confirming the predicted intent of Bob, followed by the robot coordinating a robotic arm on the bench to make a tea for Bob while the cereal is being consumed.

In this work, we propose a novel framework named Intent-aware Interactive Internet of Things for implementing ambient intelligence to facilitate seamless collaborations between humans, smart objects, and robots based on a unified IoT platform.² Our framework consists of four tasks: 1) fusing sensory data from robots and smart homes to infer human intent based on past and current human behaviors; 2) combining the inferred intent with observations of the robot (via a 3D camera) to recognize semantic context and prompt physical interactions; 3) defining objectives for the robot; and 4) commanding the robot to take actions toward these objectives. Our main contributions are summarized as follows:

- 1) We propose an end-to-end solution for interactive ambient intelligence, where robots (also called robotic assistants) assist humans in conducting physical activities in IoT-enabled smart homes.
- 2) We engineer generic Web-based descriptions to facilitate interactions of smart-home devices. They enable the system to not only decide but also execute actions in the physical world.

- 3) We train reinforced learning-based robotic assistants to exploit successful and failure experience. We also apply a policy continuation strategy and a Hindsight Experience Replay method to expedite the learning process.
- 4) We implement the system with the Deepbot framework in the OpenAI Gym^{c,19} environment to demonstrate some typical activities in real life.

RELATED WORK

Smart-home-related studies have attracted enormous attention, thanks to the availability of networked smart devices and improved computing capability. Early studies in the field concern Web management, localization, tracking, or activity recognition. CASAS³ is one of the earliest smart-home experimental platforms that tracks users with preinstalled devices in apartments. Ruan *et al.*⁴ achieve device-free indoor human localization and tracking. Yao *et al.*^{5,6} demonstrated a unified management system that integrates monitors physical devices and finds relevant things according to past human interactions. Shemshadi *et al.*⁷ provide a framework for diversified relevance search in IoT, laying the foundation for IoT search engines. Their follow-up study^{2,21} reveals a real-time multilevel activity monitoring system for a personalized smart home, which

^c[Online]. Available: <https://gym.openai.com/>

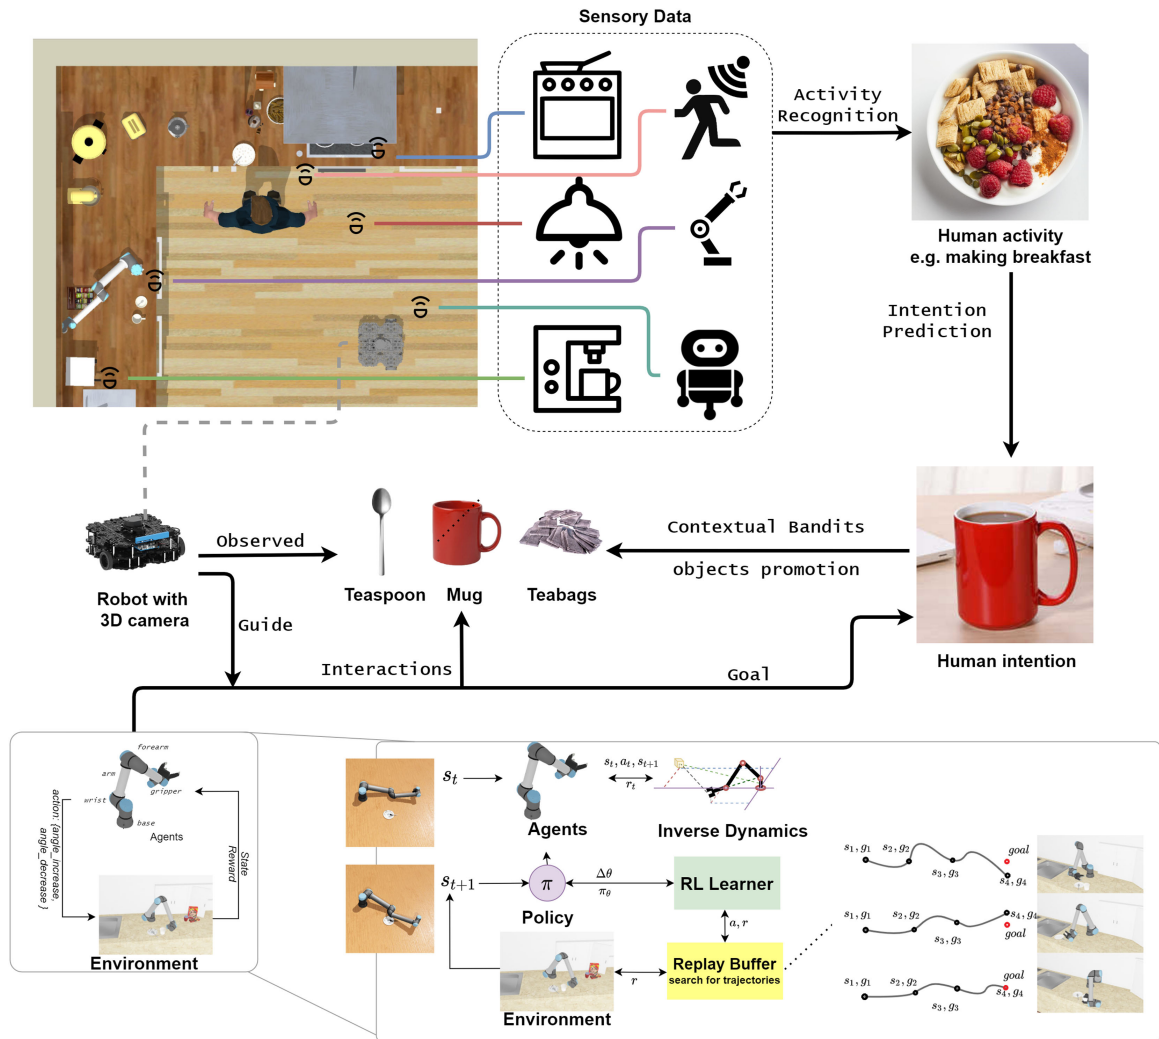


FIGURE 2. Workflow of the proposed system in our scenario (illustrated in Figure 1). First, the robot observes human activity and infers human's current behavior (i.e., making cereal breakfast) and intent (i.e., having a cup of tea) based on IoT sensory data. Then, the system figures out the objects to interact and works out the procedure for making tea. Finally, the robotic arm is guided to execute the procedure and make the predicted intent a reality (by preparing a cup of tea).

continuously tracks daily activities and conducts abnormal activity detection. These studies realize ambient intelligence in certain aspects or applications but are still far from being truly "smart" in the sense of directly assisting human activities. Currently, robots are usually dedicated to specialized tasks but they offer great flexibilities in terms of "taking actions" in a smart environment. The concept of Internet of Robotic Things (IoRT)¹ aims to integrate robotics with IoT networks. Mahieu *et al.*¹ investigated context-aware and personalized interactions on the IoRT. Vermesan *et al.*⁸ review IoRT and clarify some concepts with suggested architecture and applications. All these efforts form the base of our work.

METHOD

Our framework (Figure 2) explores ambient intelligence to enable autonomous robots to collaborate actively with humans in their daily activities in smart homes. The framework can infer human intent, compose actuation for interactions, and finally execute actions via robots.

Problem Setup

We first model human activities and develop a segmentation method for recognizing complex activities with awareness of concurring events and the subject intentions. Then, we finalize conscientious moves and

devise instructions for the robot assistants. We use the moving robot with a 3D stereo camera as the visual guide for other cooperative robot assistants to collaborate with the human. The 3D camera enables the robot to scan the environment, monitor the subject's activities, and infer the next move of the subject based on the usage history and basic rules preset or learned beforehand. After identifying interactable objects, the robot will receive step-by-step instructions by a smart and adaptive robotic control system.

Let $A \in \mathbb{R}^N$ be a set of activities, where $a_i \in A$ is the i th class of activity in the set. At time t , the system receives sensory input $x_t^A \in \mathbb{R}^d$, along with the predicted activity for this input:

$$x_t^A \rightarrow \hat{y}_t^A = a_i.$$

Suppose $Q_t^A \in R^q$ is the semantic description corresponded to the current input x_t^A . The system matches this description with the semantic keywords $Q_t^O \in R^p$ of the objects detected by the robot R_m . With the contexts and historical usage information (A_H), the system infers the future action of the human at time $t' + 1$, $y_{t'+1}^A = \hat{a}_j \in A_i$; accordingly, the action compositor $\alpha(\cdot)$ produces executable procedure with a series of interactions $\mathcal{I} \in R^M$ for the target to Q^A with objects to Q^O assumed that in a short time duration $Q_{t'}^O + 1 = Q_t^O = Q^O$, as illustrated below:

$$\{Q_t^A, Q^O, A_H\} \rightarrow \hat{a}_j$$

$$\mathcal{I} = \alpha(\hat{a}_j, Q^O).$$

Once receiving the interactions \mathcal{I} , the robot sees each interaction ι_m as a goal and learns to reach these goals.

We abstract all the connected and interactive things in a smart home environment as Web resources to support dynamic discovery, and perform hypermedia interactions as standardized Web interactions via normal Web protocols, such as HTTP. The device interactions are realized via Web services following WS-* standards, where JSON-based serialization help ease the implementation. Besides, we adapt a semantic description of Web-based Artifacts (based on Ricci *et al.*¹⁰) as the first-class abstraction for a clear integration and deployment of multiagent system (MAS).

Intent-Awareness With Connected Things

We aim to predict the next human activity and generate the corresponding semantics. We adopt a workflow inspired by Triboan *et al.*,¹¹ a semantic theory-

based approach for sensor event segmentation, to facilitate a semantic activity prediction.

Activity Modeling

The environmental context (EC) consists of human subject (H_n), location (L_m), ambient characteristic (AC_o), sensor characteristics (S_p), and interactable objects (Obj_q) of classes (C_l)

$$EC = \{H_n, L_m, AC_o, S_p, Obj_q\}.$$

Sensor Environment (SR) describes the semantic relationship (R_e) between sensor events and objects

$$\text{where } SR = a_n (R_e, EC) \rightarrow R_e \rightarrow SE$$

$SE = \text{instance}(R_e, S_p)$. Since the actual activities performed by human may not match the system's prior knowledge,¹¹ we additionally consider human preferences (Pre_{f_r}) below:

$$Pre_{f_r} = \text{instance}(R_e, a_n \cap Preference) \rightarrow R_e.$$

Semantic Decision

Based on the above relationships, we recognize activities using ontology-based semantic reasoning methods. Given a set of the streamed sensory events E^s and possible activity candidates A' , we construct an activity thread AT_i by conducting Terminology Box (T-Box) reasoning on regular, generic activities, and Assertion Box (A-Box) reasoning for user preferred activities:

$$A T_i = \{tBox[A', E^s], aBox[Pre_{f_r}[Pre_{f_r}', E^s]]\}.$$

We analyze the metadata of sensor event $e_m \in E^s$ for the corresponding entity (ET_K) to deduce candidate relationships with activities; the concurrent activities can be inferred during this process using the above semantic reasoner.

Sensor Events Segmentation

We generalize the concept of sensory events to contain status and usage information gathered in the IoT networks. Inspired by Mallick *et al.*,¹² we adopt a transaction-based segmentation as activities may be multiplexed. Given a set of sensory events $E = \{e_1, e_2, \dots, e_n \mid e_i \in E^s\}$, we segment it into multiple transactions, denoted by $t r_i = \{e_i, e_{i+1}, \dots, e_{i+j}\} \in Tr$. Suppose ProperCut is the transaction tr_i that matches exactly with activity a_i , OverCut includes the transactions that do not contain all transactions, and UnderCut covers those transactions that involve multiple activities. Then, our goal becomes to minimize the



FIGURE 3. (a) Overall testing success rate to training episodes. (b) Overall learning rewards obtained to training episodes.

number of overcutting and undercutting transactions. To this end, we identify the minimal transactions of sensory events that contain complete activities for our downstream processing, the activity prediction, and use the *MinMax* algorithm to embed and cluster contextual information of sensors and IoT devices, where we use distances and temporal sequences for segmentation.

Activity Prediction

We infer a user's possible next moves for our downstream processing, such as giving commands to robotic assistants. Then, we employ two stages of the method proposed by Altulyan *et al.*¹³ to suggest the next items to be used by the human:

- 1) *Complex activity recognition*: At the current stage, the system learns to recognize simple activities, such as the posture and movement of human. The next step is to recognize complex activities based on predefined ontology and rules.
- 2) *Recommendation*: At the current stage, the system prompts activities and items to be used by human or robot in the next step. We learn from past trajectories of activities with Q-learning.

Semantic Searching for Interactive Internet of Things

The system uses semantic search to identify the suitable objects and generate the procedure for a robot to act toward the inferred human intent. We use Mask R-CNN,¹⁴ a widely used pixel-level object segmentation model that outputs 1 or 0 to indicate

whether a pixel belongs to an object, for object detection. We also combine objective detection with depth information to get 3D positions of objects of interest. Such position information is then used to adapt the learning goals and training/running processes of robots.

Context-Aware Hypermedia Interactions

We introduce an RL-based approach named Policy Continuation with Hindsight Inverse Dynamics (PCHID) to generate hardware commands for instructing robotic arms' movements.

Policy Continuation With Hindsight Inverse Dynamics

PCHID¹⁵ supports self-supervised learning for goal-conditional tasks and can extrapolate the learned policy to complex and nonlinear Hindsight Inverse Dynamics (HID) applications. Specifically, PCHID introduces the goal into inverse dynamics as Hindsight Inverse Dynamics and outputs the minimal k -step actions to achieve the goal under the optimal policy. This enables us to complete a complex task in multiple steps, each achieving a different subgoal.

Multistep Goals With Policy Continuation

While PCHID¹⁵ makes the system aware of the hindsight goal, more sophisticated settings are introduced and it is considerable challenging. We extend the original Hindsight Experience Replay (HER) model with the ideal proposed by Sun *et al.*¹⁵ to let HID embrace k -step solubility for better optimization. As in the




FIGURE 4. Learning results for the task “Remove cup” on some selected intermediate episodes, where (a) task failed with a gripper of robotic arm hitting the desk; (b) the cup was gripped yet other items were knocked off; (c) the cup was gripped and removed successfully.

Universal Value Function Approximators (UVFA),¹⁶ we have possible goals $g \in \mathcal{G}$ and a corresponding reward $r_g : \mathcal{S} \times \mathcal{A} \rightarrow R$, where \mathcal{S}, \mathcal{A} are the state space and action space, respectively. At timestamp t , we have $r_t = r_g(s_t, g_t)$ and policy $\pi : \mathcal{S} \times \mathcal{G} \rightarrow \mathcal{A}$. The reward is extended from the $(0, 1)$ binary problem set to the following:

$$r(s_t, a_t, g) = \lambda |g - s_t^{object}|^p - |g - s_{t+1}^{object}|^p.$$

The hyperparameter $p \in \{1, 2\}$, and λ may not be limited to 0 or 1 (to be discussed later in our evaluation and experiments).

Based on the UVFA model of HID, we apply the k -step solvable extension, thus decomposing the state-goal $\mathcal{S} \times \mathcal{G}$ into $\mathcal{S} \times \mathcal{G} = (\mathcal{S} \times \mathcal{G})_0 \cup (\mathcal{S} \times \mathcal{G})_1 \cup$

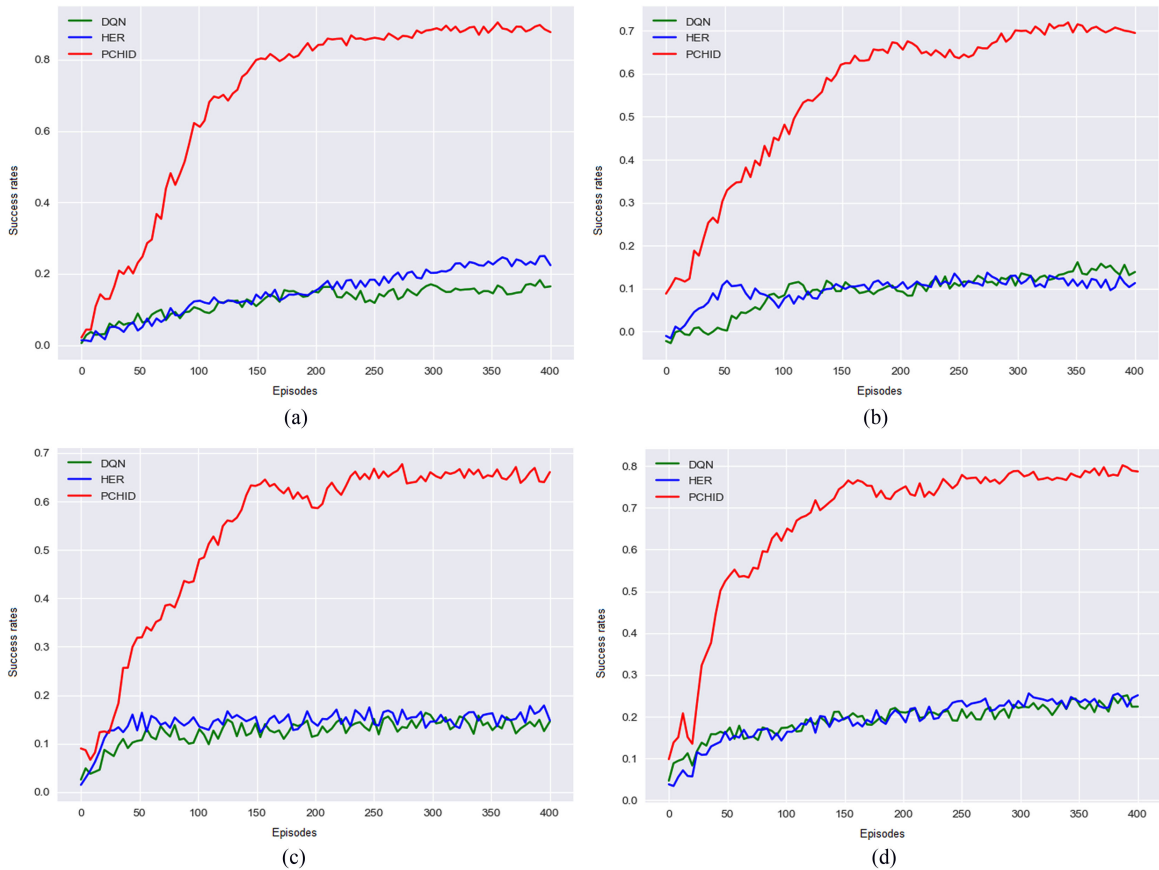


FIGURE 5. Testing success rates to training episodes for certain tasks. (a) Pick and lift. (b) Place cup. (c) Remove cup. (d) Press switch.

$\dots \cup (\mathcal{S} \times \mathcal{G})_T \cup (\mathcal{S} \times \mathcal{G})_U$ (see Sun *et al.*¹⁵ for more technical details). Our objective here is to find the optimal policy $\pi^* = \pi_0^* \cup \pi_1^* \cup \dots \cup \pi_T^* \cup \pi_U^*$, where each policy π_t^* at timestamp t leads to the next subgoal. This objective is solvable by stochastic gradient descent (SGD) models, as shown below

$$\theta_k = \arg \min_{\theta_{k-1}} \sum_{s_t, s_{t+1}^i, a_t^i, i \in 1 \dots k} \left\| f_{\theta_k} \left((s_t, g_{t+1}'), (s_{t+1}, g_{t+1}') \right) - a_t \right\|^2.$$

EVALUATION

We have implemented live demonstrations of some selected activities^d and evaluate the feasibility of our system in real-life applications based on simulation.

We test the task learning process of robotic assistants based on RLbench,¹⁷ an open source robotics-related manipulation benchmark that contains 100 daily activities. While RLbench uses the PyRep simulator, we implement our system with the Deepbot framework¹⁸ in the OpenAI Gym environment,¹⁹ which enables us to synchronize the simulation with Robot Operating System (ROS). Our setting involves a Lynxmotion AL5D robotic arm, a TurtleBot waffle Pi equipped with a 3D stereo camera, and some interactive items, such as kitchen utensils.

As aforementioned, the robotic assistant actuates based on the UVFA model of HID, an extension to DQN.²⁰ We compare PCHID with DQN and its successor HER to demonstrate the effectiveness of our framework in RLbench tasks. Specifically, we set the reward to 0 when the final state is within the tolerable range of the subgoal for each step and -1 otherwise. Figure 3 shows the overall success rates and rewards with respective to episodes of training. The PCHID method is effective and learns significantly faster than HER and the original DQN methods. Figure 4 shows some intermediate learning outcomes: a) and b) failed to complete the task of "Pick and lift" while c) succeeded at episode 20, 50, and 100. Figure 5 further demonstrates the effectiveness of PCHID in some selected tasks: a) Pick and lift; b) Place cups; c) Remove cups; d) Press switch. In our experiments, we set the steps k to 5—we observed a larger step tended to improve the results.

CONCLUSION

In this article, we proposed an interactive ambient intelligence system that enables robots to actively infer human intents and assist humans in accomplishing daily tasks. Our system relies on human activity

prediction and reinforcement learning on IoT sensory data in a controlled smart-home environment. In the future, we will develop an interactive item-discovery method to dynamically expand the system's knowledge base about the environment and prompt robotic actions.

ACKNOWLEDGMENTS

The work of L. Yao and Q. Z. Sheng was supported in part by Australian Research Council under Grant DP21010723 and Grant LP190100140. The work of Z. Wang and X. Xu was supported by the National Key Research and Development Program of China under Grant 2018YFB1402500 and the National Science Foundation of China under Grant 61832004 and Grant 61772155.

REFERENCES

1. C. Mahieu, F. Ongenaes, F. De Backere, P. Bonte, F. De Turck, and P. Simoens, "Semantics-based platform for context-aware and personalized robot interaction in the internet of robotic things," *J. Syst. Softw.*, vol. 149, pp. 138–157, 2019.
2. L. Yao *et al.*, "WITS: An IoT-endowed computational framework for activity recognition in personalized smart homes," *Comput.*, vol. 100, no. 4, pp. 369–385, 2018.
3. D. J. Cook, A. S. Crandall, B. L. Thomas, and N. C. Krishnan, "CASAS: A smart home in a box," *IEEE Comput.*, vol. 46, no. 7, pp. 62–69, Jul. 2013.
4. W. Ruan, Q. Z. Sheng, L. Yao, T. Gu, M. Ruta, and L. Shangguan, "Device-free indoor localization and tracking through human-object interactions," in *Proc. IEEE 17th Int. Symp. World Wireless, Mobile Multimedia Netw.*, 2016, pp. 1–9.
5. L. Yao, Q. Z. Sheng, and S. Dustdar, "Web-based management of the Internet of Things," *IEEE Internet Comput.*, vol. 19, no. 4, pp. 60–67, Jul./Aug. 2015.
6. L. Yao, Q. Z. Sheng, N. J. G. Falkner, and A. H. H. Ngu, "ThingsNavi: Finding most-related things via multi-dimensional modeling of human-thing interactions," in *Proc. 11th Int. Conf. Mobile Ubiquitous Syst., Comput., Netw. Serv.*, 2014, pp. 20–29.
7. A. Shemshadi, Q. Z. Sheng, Y. Qin, A. Sun, W. E. Zhang, and L. Yao, "Searching for the internet of things: Where it is and what it looks like," *Pers. Ubiquitous Comput.*, vol. 21, no. 6, pp. 1097–1112, 2017.
8. O. Vermesan *et al.*, "Internet of robotic things: Converging sensing/actuating, hypoconnectivity, artificial intelligence and IoT platforms," in *Cognitive-Hyperconnected Digital Transformation – Internet of Things Intelligence*, Singapore: River Publishers, 2017.

^d[Online]. Available: <https://drive.google.com/drive/folders/18ZsyTtt9GlFzOr5yYG-INMpzCb8VoYSC?usp=sharing>

9. H. Lu, D. Cao, D. Tao, S. Dustdar, and P. Ho, "Guest Editorial for special issue on cognitive computing for collaborative robotics," *IEEE/CAA J. Automatica Sinica*, vol. 8, no. 7, Jul. 2021, Art. no. 1221.
10. A. Ricci, M. Piunti, M. Viroli, and A. Omicini, "Environment programming in CARtAgO," in *Multi-Agent Programming*. Berlin, Germany: Springer, 2009, pp. 259–288.
11. D. Triboan, L. Chen, F. Chen, and Z. Wang, "A semantics-based approach to sensor data segmentation in real-time activity recognition," *Future Gener. Comput. Syst.*, vol. 93, pp. 224–236, 2019.
12. M. Mallick, P. Kodeswaran, S. Sen, R. Kokku, and N. Ganguly, "TSFS: An integrated approach for event segmentation and ADL detection in IoT enabled smarthomes," *IEEE Trans. Mobile Comput.*, vol. 18, no. 11, pp. 2686–2700, Nov. 2019.
13. M. S. Altulayan, C. Huang, L. Yao, X. Wang, and S. Kanhere, "Contextual bandit learning for activity-aware things-of-interest recommendation in an assisted living environment," in *Proc. 32th Australas. Database Conf.*, 2021, pp. 37–49.
14. K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2961–2969.
15. H. Sun, Z. Li, X. Liu, B. Zhou, and D. Lin, "Policy continuation with hindsight inverse dynamics," in *Proc. Conf. Adv. Neural Inf. Process. Syst.*, 2019, pp. 10265–10275.
16. T. Schaul, D. Horgan, K. Gregor, and D. Silver, "Universal value function approximators," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 1312–1320.
17. S. James, Z. Ma, D. R. Arrojo, and A. J. Davison, "RLBench: The robot learning benchmark & learning environment," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 3019–3026, Apr. 2020.
18. M. Kirtas, K. Tsampazis, N. Passalis, and A. Tefas, "Deepbots: A webots-based deep reinforcement learning framework for robotics," in *Proc. IFIP Int. Conf. Artif. Intell. Appl. Innov.*, 2020, pp. 64–75.
19. G. Brockman *et al.*, "OpenAI Gym," 2016, *arXiv:1606.01540*.
20. V. Mnih *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, pp. 529–533, 2015.
21. C. Huang and L. Yao, "Active object estimation for human-robot collaborative tasks," in *Proc. Int. Conf. Neural Inf.*, 2020, pp. 750–761.

CHAORAN HUANG is a postdoctoral fellow with the School of Computer Science and Engineering and School of Minerals and Energy Resources Engineering, UNSW Australia, Sydney, NSW, Australia. Huang received his Ph.D. degree in computer science from UNSW Australia, Sydney, NSW, Australia. Contact him at chaoran.huang@unsw.edu.au

LINA YAO is an associate professor with the School of Computer Science and Engineering, UNSW Australia, Sydney, NSW, Australia. Yao received her Ph.D. degree in computer science from the University of Adelaide, Adelaide, Australia. Contact her at lina.yao@unsw.edu.au.

XIANZHI WANG is a lecturer with the School of Computer Science, University of Technology Sydney, Ultimo, NSW, Australia. Wang received his Ph.D. degree in computer science from Harbin Institute of Technology, Harbin, China. Contact him at xianzhi.wang@uts.edu.au.

QUAN Z. SHENG is a full professor and the Head with the Department of Computing, Macquarie University, Sydney, NSW, Australia. Sheng received his Ph.D. degree in computer science from the University of New South Wales, Sydney. Contact him at michael.sheng@mq.edu.au.

SCHAHRAM DUSTDAR is a full professor of computer science and the Head with the Distributed Systems Research Division, TU Wien, Vienna, Austria. Contact him at dustdar@dsg.tuwien.ac.at.

ZHONGJIE WANG is a full professor with the School of Computer Science and Technology, Harbin Institute of Technology (HIT), Harbin, China. Wang received his Ph.D. degree in computer science from HIT. Contact him at rainy@hit.edu.cn.

XIAOFEI XU is a full professor with the School of Computer Science and Technology, and the Vice President of Harbin Institute of Technology (HIT), Harbin, China. Xu received his Ph.D. degree in computer science from HIT. Contact him at xiaofei@hit.edu.cn.