CIRP Annals - Manufacturing Technology 72 (2023) 1–4

Contents lists available at ScienceDirect

CIRP Annals - Manufacturing Technology

journal homepage: https://www.editorialmanager.com/CIRP/default.aspx

A collaborative intelligence-based approach for handling human-robot collaboration uncertainties

Pai Zheng[a], Shufei Li[a], Junming Fan[a], Chengxi Li[a], Lihui Wang (1)[b,*]

[a] *Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, HKSAR, China*
[b] *Department of Production Engineering, KTH Royal Institute of Technology, Sweden*

ARTICLE INFO

Article history:
Available online 15 April 2023

Keywords:
Human-robot collaboration
Manufacturing system
Collaborative intelligence

ABSTRACT

Human-Robot Collaboration (HRC) has played a pivotal role in today's human-centric smart manufacturing scenarios. Nevertheless, limited concerns have been given to HRC uncertainties. By integrating both human and artificial intelligence, this paper proposes a Collaborative Intelligence (CI)-based approach for handling three major types of HRC uncertainties (i.e., human, robot and task uncertainties). A fine-grained human digital twin modelling method is introduced to address human uncertainties with better robotic assistance. Meanwhile, a learning from demonstration approach is offered to handle robotic task uncertainties with human intelligence. Lastly, the feasibility of the proposed CI has been demonstrated in an illustrative HRC assembly task.

1. Introduction

Existing automation systems have reached its bottleneck in handling various flexible manufacturing tasks, such as the assembly or disassembly of complicated products with frequent changes [1]. To satisfy human-centric needs in these working conditions, Human-Robot Collaboration (HRC) is becoming an active area that combines the strength, repeatability and accuracy of robots with the cognitive flexibility of humans to achieve a high-level flexible automation [2]. Especially, recent studies on the proactive HRC approaches immerse mutual-cognitive, predictable, and even self-organising elements into a human-robot execution loop [3]. In this context, human and robotic agents can understand bi-directional operation intention, learn task operation sequences and even execute adaptive manipulations in a self-organised manner [4]. Nevertheless, many uncertainties still exist in real-life HRC manufacturing scenarios, which impede their successful and effective implementations. To the authors' best knowledge, those HRC uncertainties can be further classified into three major types, namely *human uncertainty, task uncertainty* and *robot uncertainty*.

Human uncertainty is the most common one, since humans have a much higher degree of autonomy and spontaneity. It mainly consists of abnormal human behaviours [5] (e.g., sudden change of gesture, collisions with the predefined robotic trajectories), which pose a great challenge to the robotic collaborators.

Task uncertainty is also common in many HRC manufacturing scenarios, especially in the disassembly/inspection process of used products, due to their uncertain shapes, sizes and physical conditions, such as rusted, loosened and defective components, stained surfaces, and geometry changes [6].

Robot uncertainty denotes its runaway motions in unstructured workspaces, which may bring confusion in operation arrangement and human safety issues in human-robot physical interaction.

To tackle these problems, this paper introduces a Collaborative Intelligence (CI)-based approach to handling industrial uncertainties in HRC systems with respective tolerance settings. According to Wilson and Daugherty [7], CI looks forward to joining forces between human intelligence and artificial intelligence from both humans-assisting-machines and machines-assisting-humans perspectives. In HRC scenarios, the proposed CI-based approach is mainly embodied in 1) *human-assisting-robot*, humans train robots how to adjust operations for new situations, explain which decision is taken, and sustain what knowledge is updated in task and robot uncertainties; and 2) *robot-assisting-human*, robot can amplify human easy working conditions, let humans interact with work with flexible decisions, and embody human-centric needs when facing human uncertainties. Following this manner, the rest of this paper is organised as follows. First, Section 2 depicts the systematic methodology of CI-based HRC. Then, Section 3 gives a demonstrative case study to validate its feasibility. Section 4 finally highlights the conclusions and future works.

2. CI-based HRC uncertainties handling

To handle various HRC uncertainties, the proposed CI-based system framework with core procedures is shown in Fig. 1.

Firstly, adaptive HRC task planning can be generated from a holistic scene perception, including object detection, human action recognition [8], and point cloud segmentation of the surrounding environment. The task planning strategy is denoted by the HRC Knowledge Graph (KG), which leverages a dynamic graph embedding-based method to infer human-robot relations and task structures in different stages, achieving a mutual understanding of teamwork. *In certain situations*, the adaptive task planning module assigns operations from three aspects, (1) robot-centric mode, (2) human-centric mode, and (3) adaptive mode. The former two modes grant authorities to robots and humans respectively, so that they can perform different types of tasks based on their own capabilities. Lastly, the adaptive mode allows humans and robots to

* Corresponding author.
E-mail address: lihui.wang@iip.kth.se (L. Wang).

https://doi.org/10.1016/j.cirp.2023.04.057

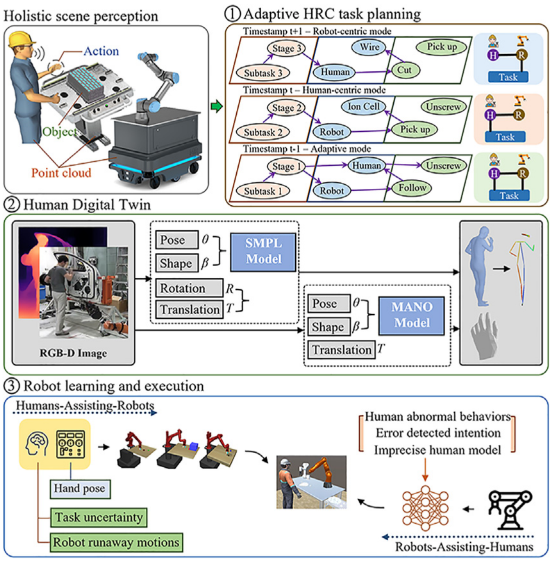

Fig. 1. The proposed CI-based approach for handling HRC uncertainties.

coordinately take operations. *In uncertain situations*, the Human Digital Twin (HDT) model and robot learning methods are triggered. The HDT is developed to reconstruct the fine-grained human hand and full body pose during execution. When a human operator discovers task uncertainties or robot runaway motions, one can teach a robot and update its manipulation skills from a flexible demonstration via the proposed Learning from Demonstration (LfD) method. When the HDT captures human uncertain activities in the task process, the robot can self-optimise its trajectories via a Deep Reinforcement Learning (DRL) based approach to adapt to humans' abnormal behaviours for safe and efficient collaboration. Thus, the CI-based HRC maximises human-robot complementarity and resists uncertain situations in tasks.

2.1. Dynamic graph embedding-based adaptive HRC task planning

Adaptive task planning algorithms free existing HRC systems from predefined instructions, for enhanced adaptability in various manufacturing scenarios. Our previous works have introduced the scene graph (SG) and KG approaches [4] to generate task planning strategies for HRC systems. Meanwhile, Raatz et al. [9] utilised a genetic algorithm to optimise task scheduling based on capabilities and time assumptions in HRC. However, these methods focus on distilling temporal knowledge representation of HRC configurations, while ignoring updating the information to next-stage task arrangement.

With perceptual results, a DynGraphGAN model [10] is leveraged to temporally construct and update HRC KG for on-demand task allocation, as shown in Fig. 2. Considering input of human behaviours, detected objects and task structures, there are numerous possible configurations of their relations. Firstly, a generator generates adjacency matrices to represent the relation evolution between Human-Robot-Task-Workpiece-Environment (HRTWE) nodes over time. The updated HRTWE connections may introduce fake edges. Thus, a continuous discriminator is leveraged to distinguish the authentic and fake links between HRTWE nodes via Gated Recurrent Unit (GRU) algorithm. Optimised operation

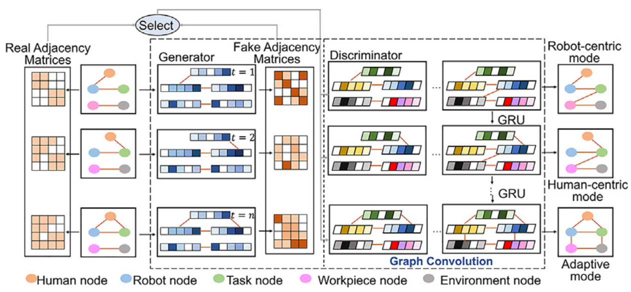

Fig. 2. Dynamic graph embedding-based adaptive HRC task planning.

arrangement and sequential orders are obtained by connecting HRTWE nodes with updated embeddings, achieving HRC KG update. The connection from a robot node to a task node indicates robot action types, while the following edge from the task node to a workpiece node denotes the robot endpoints and operation goals. Various task-planning strategies are represented in an explainable graphical manner.

2.2. Vision-based fine-grained HDT modelling

Based on the adaptive HRC task planning result, the ability to constantly perceive and model the human body is also essential, which can provide necessary information for the collaborative robot to further cope with human related uncertainties during HRC. Previous endeavours in the HRC literature have been devoted to perceiving human body skeleton for active collision avoidance [11] or recognising human action intention [12] for robotic decision-makings. However, they can only model the human operator in a relatively coarse manner with insufficient representation fidelity or limited recognition accuracy. To this end, a vision-based fine-grained HDT modelling scheme of a human operator is proposed and depicted in Fig. 3, which mainly consists of two parts: 1) fine-grained human pose reconstruction, and 2) spatial-temporal human behaviour intention recognition.

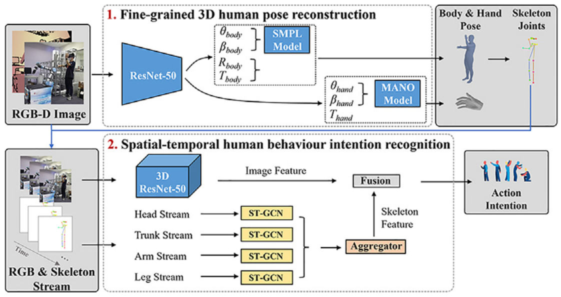

Fig. 3. Vision-based fine-grained HDT modelling scheme.

Fine-grained human pose reconstruction. In the first part of the HDT, a deep learning model that can simultaneously reconstruct the fine-grained 3D dense mesh and skeleton joints of the human body is proposed. Concretely, the RGB-D images of the human operator will be processed by a ResNet-50 backbone network to extract the geometric features, which is then utilised to regress the pose parameters $\theta_{body} \in \mathbb{R}^{3\times K}$, shape parameters $\beta_{body} \in \mathbb{R}^{M}$, 3D rotation $R_{body} \in \mathbb{R}^{3\times 3}$, and 3D translation $T_{body} \in \mathbb{R}^{3}$. The pose and shape parameters are subsequently sent to the SMPL (Skinned Multi-Person Linear model) human body model – a differentiable function that outputs a triangulated mesh $M(\theta, \beta) \in \mathbb{R}^{3\times N}$. The adoption of the SMPL model can largely simplify the reconstruction process to achieve real-time performance by relying on a template human body mesh as a priori which will be bended and stretched to the target human pose according to the estimated pose and shape parameters. The predicted global 3D rotation and translation will then be applied to obtain the correctly posed human body mesh. The 3D skeleton points $X(\theta, \beta) \in \mathbb{R}^{3\times K}$ are further obtained by linear regression from the final mesh vertices. To further refine hand pose reconstruction of human operators, the ResNet-50 backbone will also regress hand parameters including pose parameters $\theta_{hand} \in \mathbb{R}^{3\times K}$, shape parameters $\beta_{hand} \in \mathbb{R}^{M}$, and 3D translation $T_{hand} \in \mathbb{R}^{3}$, which will be processed by the MANO (hand Model with Articulated and Non-rigid defOrmations) to enhance the SMPL model and to facilitate the human-assisted robot LfD. Once the fine-grained human mesh is reconstructed, it can represent the precise geometric occupancy of the human body, which can substantially reduce the perception error of human body during HRC.

Spatial-temporal human behaviour intention recognition. The spatial-temporal human behaviour intention estimation amounts to a higher semantic level, which is essential to complete a holistic HDT in HRC scenarios. In this module, the RGB video stream and associated skeleton stream are regarded as the input data source. For the RGB stream, a 3D ResNet-50 will be used to process the spatial-temporal features in a unified convolution structure and extract the image feature. Meanwhile, the skeleton stream will be split into 4 branches including head, trunk, arm, and leg branches, each of which will be processed by an ST-GCN (Spatial Temporal Graph Convolutional Network) model. After the spatial-temporal feature extraction of the local body parts, an aggregator network will be utilised to fuse them into the global skeleton feature, which will be fused with the extracted image feature to discriminate the type of current human behaviour and if it is abnormal. Since the model is exclusively trained on normal behaviour data, it will only be able to provide a random guess with extremely low confidence score for an unseen abnormal behaviour sequence, of which setting up a tolerance on the confidence scores can eliminate any wrong detection. For normal human behaviours, the recognised action intention will be passed to robots to make action and motion planning ahead of time, while for abnormal behaviours that cannot be properly parsed by the system, a warning will be signalled to notify human for behaviour correction or optionally trigger the LfD protocol.

2.3. Human-assisting-robot via LfD

Considering robot and task uncertainties, a human operator can potentially transfer one's experience to robot manipulation skills for flexible and adaptive task execution via the LfD approach.

Human-in-the-loop robot control. To better transfer human experts' engineering practice, the hand pose extracted from HDT is introduced to implement a seamless hand gesture-enabled robot control system. As shown in Fig. 4, the method extracts the worker's hand poses and corresponds them to the robot end-effector poses, which could intuitively enable the robot to mimic the hand movement. The worker's left-hand pose can be accurately extracted via his palm location/orientation and transformed into the end-effector pose of the robot accordingly.

Fig. 4. Examples of mapping between human hand and robot end-effector.

Robot learn from demonstration & dataset aggregation. In addition to explicit imitation, the robot movement trajectories demonstrated by the experts are recorded as datasets and fed into the LfD algorithm, a promising learning approach for empowering the skills of robots. With the dataset, the workers' patterns could be implicitly extracted and drive robots to learn the uncertainties-oriented control policy for subsequent adaptive robot programs.

In this work, to approximate the control policy function $\pi_\theta(a|s)$, an LfD approach adopted is Behavioural Cloning (BC) algorithm. To fit into the LfD algorithm settings, the data trajectories i.e., $\tau_1, \tau_2, \ldots, \tau_m$ consist of the environment observation s_1^i and robot motion action a_1^i. The element of s includes the task information, working environment, task conditions, human information etc., and a is the robot movement demonstrated by the human-expert regarding the cases. Each demonstration of the dataset is represented as $\tau_i = \langle s_1^i, a_1^i, s_2^i, a_2^i, \ldots, s_{n+1}^i \rangle$ and the whole dataset $\mathscr{D}$ is denoted as $\mathscr{D} = \{(S_1, A_1), (S_2, A_2), (S_3, A_3)\ldots\}$. Essentially, the BC learning is to approximate the maximum likelihood estimation of policy function, which could minimise the difference between the model-generated state-action trajectory probability distribution (robot control policy) and the input trajectory probability distribution (human-expert policy):

$$\max_\theta \mathbb{E}_{(s,a)\sim\mathscr{D}}[\pi_\theta(a|s)] \tag{1}$$

$$s.t. \sum_{a\in A} \widehat{\pi}_\theta(a|s) = 1, \ \forall s \in S$$

During the parameter optimisation process of the maximum likelihood estimation, a policy $\pi_\theta(s)$ is trained to minimise the difference between robot behaviour patterns and human demonstration with the objective function $\mathbb{E}_{(s,a)\sim\mathscr{D}} \| \pi_\theta(s) - a \|^2$. In practice, $\pi_\theta(s)$ approximates the expert policy by using the deep neural network, which is optimised via the gradient descent with an aim to gain the optimal robot control policy function.

Algorithm 1
Pseudo code of Dagger algorithm.

Input: Original Dataset D
Output: Optimal updated policy $\tilde{\pi}_{update}$, Aggregated dataset D_{agg}
Initialisation: Original policy π_o
$\pi_{i=1} \leftarrow \pi_o$
For *episode* $i = 1, 2, \ldots, T$ **do**
Sample $T-step$ trajectories with trained policy π_i
Generate/Demonstrate dataset $D_i = \{(s, \pi^*(s)\}$ via expert π^* for unsolved case trial (determined by expert) by π_i
Aggregate the Dataset $D \leftarrow D \cup D_i$
Re-train the control policy function π_{i+1} via behavioural cloning
End For
Return *optimal policy* π_i *determined by evaluating the success rate of task*

However, owing to the diversity of samples, the effectiveness of LfD is limited by the number and variance amongst expert demonstrations. Therefore, the robot policy that trained by BC algorithm still lacks the flexibility and adaptability towards newly occurred task or robot uncertainties. To address that, an online learning approach with dataset aggregation (Dagger) mechanism is introduced in the LfD process, whose workflow is shown in Algorithm 1. With the Dagger mechanism, not only existing uncertainties but also similar but new ones can be addressed. Thereby, the robot can better resolve the dynamic manufacturing tasks with the help of expert intelligence more efficiently.

2.4. Robot-assisting-human via DRL

Considering abnormal behaviours caused by human uncertainties, a robot may dynamically re-plan its motion to complete assigned tasks from HRC KG, and also to ensure human safety. A DRL-based approach is introduced in this work to achieve human uncertainty-aware robot control for safe and adaptive HRC.

The human uncertainty factors detected from HDT are embedded in the full-body skeleton position, abnormal behaviours warnings, and intentions. In the implementation, the DRL-approach leverages uncertainty factors, motion planning success rate, and safe constraints in HRC scenarios as the optimisation indexes. Following that, the robot motion planning process can be formulated as a Markov Decision Process to optimise the control policy π^* and reinforce the robot to choose an action $a_t \in A$ in the state $s_t \in S$ with the largest cumulated reward. The DRL settings are shown as follows:

Observation state (*O*) is a state representation of the human-robot working scene, which consists of the human data extracted from the abovementioned HDT, including whole-body skeleton position *(P)*, abnormal behaviours warnings (*B*), human intentions (*I*), and also robot own state *(M).* The information can be concatenated in a state vector, $\boldsymbol{O} = (P, B, I, M)$.

Action space (*A*) refers to the reachability of the robot. In the experiment, the inverse kinematics is combined to transform the robot's joint space into the three-dimensional spatial coordinates of the end effector $\boldsymbol{A} = (X, Y, Z) \in \mathscr{R}^3$ and enable DRL algorithm to explore the feasible trajectory.

Reward (*R*) combines multiple safe motion planning tolerance settings, including safety (e.g., human-robot distance ≥10 mm), and task completion index evaluation (e.g., execution time ≤ 30 s), task completion progress and success rate (e.g., target reaching deviation ≤1 mm), which is denoted as ***R***= (*Rs,Rt*).

Meanwhile, the actor-critic is employed to learn and control the corresponding actions, which maximises the expected return $J(\theta)$ and optimises the moving path in terms of safety:

$$J(\theta) = \mathbb{E}_{\tau\sim p_\theta(\tau)}\left[\sum_{t=0}^{T} \gamma^t r_t\right] \tag{2}$$

where $p_\theta(\tau) = p(s_0)\prod_{t=0}^{T-1}[(s_{t+1}|s_t, a_t)\pi_\theta(a_t|s_t)]$ is the probability distribution over all possible state-action trajectories $\tau = (s_0, a_0, s_1, \ldots, a_{T-1}, s_T)$, $\gamma^t \in [0, 1]$ is the discount factor at time t, and $d_\theta(s_t)$ is the state distribution under the policy θ_π.

3. Case study

To illustrate the performance of the HRC system when tackling various uncertainties, comparative experiments are conducted on some HRC assembly tasks in our lab environment, containing visual sensors (Azure Kinect and Intel D435), a GPU server (RTX 3080), one human operator and collaborative robot (UR5). Firstly, the effectiveness of the proposed HDT model is evaluated with other baseline methods for finding human abnormal behaviours and assisting robots. Then, the LfD and DRL control policies are demonstrated and assessed in several *human-assisting-robot* and *robot-assisting-human* uncertain tasks, while robot uncertainties are handled manually by emergency stop and human inspection.

3.1. Fine-grained HDT modelling for accurate robot assistance

The performance evaluation of the proposed HDT model mainly consists of two parts: (1) human behaviour recognition accuracy and (2) 3D human pose reconstruction error. For the prior one, RGB-D data was collected via an Azure Kinect camera capturing the HRC activities including 5 types of human-centric HRC subtasks: (1) dismantling, (2) part picking, (3) robot handover, (4) robot guiding, and (5) robot stopping. After trimming and cleaning the raw data, a total of 939 valid action clips remained for evaluation purpose, of which 751 were used for training and 188 for testing. As for the evaluation of the fine-grained human pose reconstruction, we utilise MPJPE (Mean Per Joint Position Error) metric to evaluate body and hand posture reconstruction errors, respectively. The comparative results are listed in Table 1, which clearly shows that our proposed HDT modelling scheme performs better on both behaviour recognition and human pose reconstruction compared with previous approaches which only targeted at a sole recognition task. The evaluation results demonstrated that the proposed HDT model is capable of conducting subsequent robot assistance action planning based on human activity recognition.

Table 1
Comparative experimental results on the collected HRC data.

Method	Body pose error (mm)	Hand pose error (mm)	Behaviour recognition accuracy	Example
Kanazawa et al. [13] **Ours**	67.19 **52.14**	– –	97.89% **98.94%**	
Hasson et al. [14] **Ours**	– –	50.95 **38.41**	– –	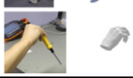

3.2. LfD and DRL experiments for handling HRC uncertainties

Human-assisting-robot, the LfD-based robot control experiments are conducted based on RoboMimic [15] dataset to evaluate the effectiveness of the Dagger integrated BC to handle HRC task uncertainties, including picking, sorting, and assembly of bearing and base. Given 20 validated task trials, different numbers of expert demonstrations (i.e., 50, 100, 150, 180) are fed to Dagger, and their corresponding success rates are shown in Table 2. It is found that BC can effectively extract human manipulation patterns to drive the robot for completing tasks. Meanwhile, the Dagger mechanism with BC could significantly enhance the robustness of robot control policy with growing human demonstrations.

Robot-assisting-human, the success rate of safe robotic motion planning is used to identify the effectiveness of tackling human uncertainties. More specifically, it is evaluated based on the overlapped distance between human operator and the work envelope of the robot. In this study, the Actor-Critic algorithm is designed and executed for robotic safe motion planning in the high fidelity Unity3D environment via 2 million steps. The resulting success rates of various overlapped distances are shown in Table 3, which can largely improve safe motion planning for HRC works.

Table 2
Comparative experimental results on the Dagger-based LfD.

Demonstrated tasks/trials		50	100	150	180
Success rate	Picking	90%	96%	100%	100%
	Sorting	60%	82%	86%	88%
	Assembly	50%	72%	82%	88%

Table 3
Experimental results of the DRL-based safe motion planning.

Overlapped distance	Success rate	
No overlapping	98.7%	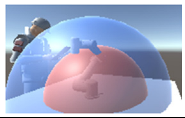
0–15 cm (Blue)	91.7%	
15–30 cm (Red)	85.8%	

4. Conclusions and future work

To ensure the successful implementation of HRC activities, this research proposed a systematic CI-based approach to handling the human, task, and robot uncertainties integrally. The main scientific contributions of this work include: 1) the proposed dynamic graph embedding-based adaptive HRC task planning approach, 2) a novel vision-based HDT modelling method for handling human uncertainties, and 3) the introduced LfD and DRL approaches for handling task/robot uncertainties and human ones, respectively. The performance of the proposed CI has been reported in handling several HRC assembly task uncertainties with preliminary experimental results. In the future, both multimodal intelligence- based HDT and advanced robot learning mechanisms can be explored to tackle multiple HRC scenarios with uncertainties.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Supplementary materials

Supplementary material associated with this article can be found in the online version at doi:10.1016/j.cirp.2023.04.057.

References

[1] Wang L (2022) A Futuristic Perspective on Human-Centric Assembly. *Journal of Manufacturing Systems* 62:199–201.
[2] Wang L, Gao R, Váncza J, Krüger J, Wang XV, Makris S, Chryssolouris G (2019) Symbiotic Human-Robot Collaborative Assembly. *CIRP Annals* 68:701–726.
[3] Li S, Zheng P, Liu S, Wang Z, Wang XV, Zheng L, Wang L (2023) Proactive Human–Robot Collaboration: Mutual-Cognitive, Predictable, and Self-Organising Perspectives. *Robotics and Computer-Integrated Manufacturing* 81:1–30.
[4] Zheng P, Li S, Xia L, Wang L, Nassehi A (2022) A Visual Reasoning-Based Approach for Mutual-Cognitive Human-Robot Collaboration. *CIRP Annals* 71:377–380.
[5] Fan J, Zheng P, Li S (2021) Vision-Based Holistic Scene Understanding Towards Proactive Human–Robot Collaboration. *Robotics and Computer-Integrated Manufacturing* 75:102304.
[6] Meng K, Xu G, Peng X, Youcef-Toumi K, Li J (2022) Intelligent Disassembly of Electric-Vehicle Batteries: A Forward-Looking Overview. *Resources, Conservation and Recycling* 182:106207.
[7] James Wilson H, Daugherty PR (2018) Collaborative Intelligence: Humans and AI are Joining Forces. *Harvard Business Review* 96(4):114–123.
[8] Li S, Zheng P, Fan J, Wang L (2021) Towards Proactive Human Robot Collaborative Assembly: A Multimodal Transfer Learning-Enabled Action Prediction Approach. *IEEE Transactions on Industrial Electronics* 69:8579–8588.
[9] Raatz A, Blankemeyer S, Recker T, Pischke D, Nyhuis P (2020) Task Scheduling Method for HRC Workplaces Based on Capabilities and Execution Time Assumptions for Robots. *CIRP Annals* 69:13–16.
[10] Xiong Y, Zhang Y, Fu H, Wang W, Zhu Y, Yu PS (2019) DynGraphGAN: Dynamic Graph Embedding via Generative Adversarial Networks. *Database Systems for Advanced Applications* 24:536–552.
[11] Bilberg A, Malik AA (2019) Digital Twin Driven Human–Robot Collaborative Assembly. *CIRP Annals* 68(1):499–502.
[12] Liu S, Wang XV, Wang L (2022) Digital Twin-Enabled Advance Execution for Human-Robot Collaborative Assembly. *CIRP Annals* 71(1):25–28.
[13] Kanazawa A, Black MJ, Jacobs DW, Malik J (2018) End-to-End Recovery of Human Shape and Pose. *IEEE Conference on Computer Vision and Pattern Recognition*, 7122–7131.
[14] Hasson Y, Tekin B, Bogo F, Laptev I, Pollefeys M, Schmid C (2020) Leveraging Photometric Consistency Over Time for Sparsely Supervised Hand-Object Reconstruction. *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 571–580.
[15] Mandlekar A., Xu D., Wong J., Nasiriany S., Wang C., Kulkarni R., Fei-Fei L., Savarese S., Zhu Y., Martín-Martín R. (2021) What Matters in Learning From Offline Human Demonstrations for Robot Manipulation. *arXiv preprint* arXiv:2108.03298.