

HDR-Net-Fusion: Real-time 3D dynamic scene reconstruction with a hierarchical deep reinforcement network

Hao-Xuan Song¹, Jiahui Huang¹, Yan-Pei Cao², and Tai-Jiang Mu¹ (✉)

© The Author(s) 2021.

Abstract Reconstructing dynamic scenes with commodity depth cameras has many applications in computer graphics, computer vision, and robotics. However, due to the presence of noise and erroneous observations from data capturing devices and the inherently ill-posed nature of non-rigid registration with insufficient information, traditional approaches often produce low-quality geometry with holes, bumps, and misalignments. We propose a novel 3D dynamic reconstruction system, named HDR-Net-Fusion, which learns to simultaneously reconstruct and refine the geometry on the fly with a sparse embedded deformation graph of surfels, using a hierarchical deep reinforcement (HDR) network. The latter comprises two parts: a global HDR-Net which rapidly detects local regions with large geometric errors, and a local HDR-Net serving as a local patch refinement operator to promptly complete and enhance such regions. Training the global HDR-Net is formulated as a novel reinforcement learning problem to implicitly learn the region selection strategy with the goal of improving the overall reconstruction quality. The applicability and efficiency of our approach are demonstrated using a large-scale dynamic reconstruction dataset. Our method can reconstruct geometry with higher quality than traditional methods.

Keywords dynamic 3D scene reconstruction; deep reinforcement learning; point cloud completion; deep neural networks

1 Introduction

3D reconstruction is a key technique in computer graphics with various applications in virtual and augmented reality and animation. In recent years, many advances have been made in both reconstruction quality and speed. Since the early success of KinectFusion [1], scanning with a commodity RGB-D camera and reconstructing the captured geometry in an online fashion have become commonplace. Subsequent work has either improved system scalability to support larger scenes and finer details by introducing new persistent data structures [2–4], or focused on enhancing reconstruction quality through accurate frame-to-model registration [5–7].

While research on reconstructing and modeling static indoor scenes [8] has matured in the past few years, reconstructing dynamic objects (e.g., humans, animals, and other freely moving objects) still remains an open problem in both the graphics and robotics communities (referred to as *dynamic SLAM* [9–13]). Given an input sequence recording a non-rigid deforming object, the goal of dynamic reconstruction is to recover the moving object's underlying shape in a canonical pose as well as the deformation field for each frame so that the geometry at each instant of time can be recovered. The seminal work of DynamicFusion [14] described a general pipeline adopted by many other algorithms: by parameterizing the per-frame deformation as a warp field defined on a sparse set of transformation nodes skinned from the full geometry, the underlying shape can be registered to the depth observations at a particular time by solving a non-rigid iterative closest point (NR-ICP) problem [15], yielding the transformation for every node. To further improve robustness, many industrial and academic solutions use dedicated hardware [16, 17], or exploit

1 BNRist, Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China. E-mail: H.-X. Song, songhx17@mails.tsinghua.edu.cn; J. Huang, huang-jh18@mails.tsinghua.edu.cn; T.-J. Mu, taijiang@tsinghua.edu.cn (✉).

2 Kuaishou Technology Co., Ltd., Beijing 100085, China. E-mail: caoyanpei@gmail.com.

Manuscript received: 2021-03-01; accepted: 2021-03-27

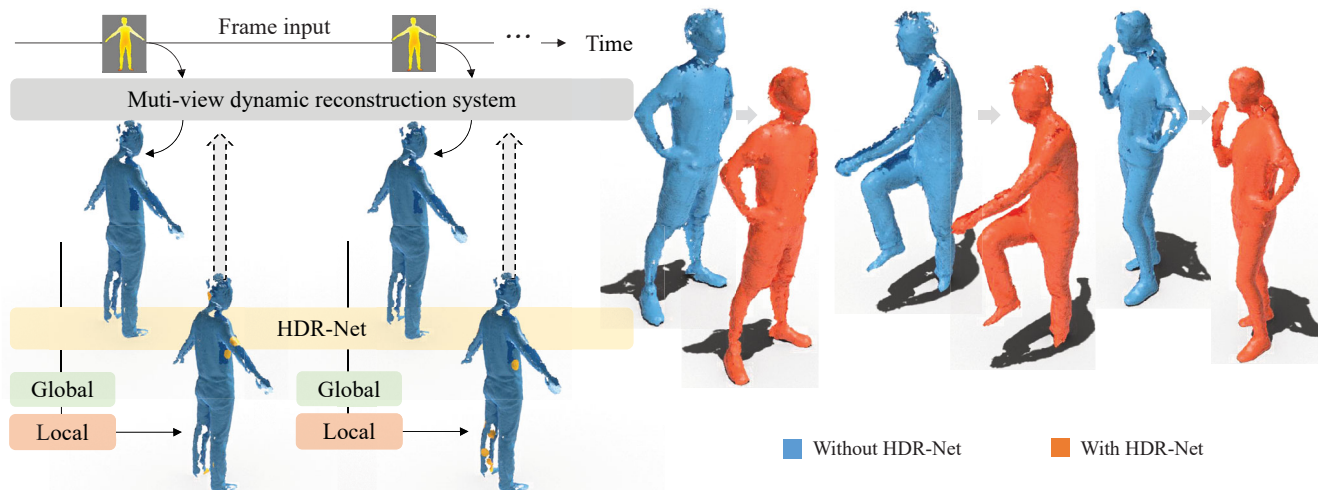


Fig. 1 HDR-Net-Fusion simultaneously reconstructs and performs patch-based geometric refinement in an online fashion. Unlike traditional reconstruction methods, it can automatically complete missing regions and refine noisy patches, improving the overall reconstruction quality. Yellow regions indicate refined patches.

a common deformable template [18, 19] as a prior to regularize the final result. On the other hand, multi-view reconstruction systems like Fusion4D [5] and FusionMLS [20] leverage more complete observations from a large number of cameras to reconstruct the geometry with higher quality.

However, reconstructing scene dynamics is inherently an ill-posed problem because the solution space for occluded regions which are not observed by any camera can be infinitely large [21]. Various regularization terms, such as as-rigid-as-possible constraints, are used to tackle this problem to some extent, but they are not always appropriate to real-world scenarios. Another challenge is the low-quality output provided by the capture device, which tends to contain noise and erroneous depth observations, resulting in artifacts in the final reconstructed model such as holes and bumps. To address the above challenges, we pursue a data-driven framework based on state-of-the-art deep learning techniques, which can be easily integrated into an existing dynamic reconstruction pipeline, to enhance the fusion quality.

Deep neural networks (DNN) have shown their applicability to a wide range of graphics applications such as shape completion [22, 23], geometric registration [24, 25], and flow/correspondence estimation [26, 27]. Recently, 3D deep learning has also gained ever more attention in reconstruction applications [28]. However, there are few attempts to embed deep models directly into reconstruction systems, primarily due to efficiency and generalization

considerations. Furthermore, in online systems where succeeding frames rely directly on previous fusion results, deep models directly operating over the already fused geometry [23, 29] cannot utilize intermediate fusion results, and it is impossible to recover from any catastrophic tracking failure in the reconstruction system.

In order to maintain system efficiency while exploiting the power of deep learning models, we present HDR-Net-Fusion, a highly-efficient dynamic 3D reconstruction system based on surfel representation [30], which can reconstruct and refine dynamic scenes simultaneously with a hierarchical deep reinforcement network, called HDR-Net. The core of HDR-Net consists of two parts: a Global-HDR-Net and a Local-HDR-Net. The global net first considers the overall geometric structure of the current model, and determines those local patches which may have poor quality, potentially leading to bad registration results for future frames. Then, the local net fixes such detected regions, performing patch-based geometric refinement using a data-driven neural network. We formulate the training of the global HDR-Net as a reinforcement learning problem: the optimal region selection strategy is implicitly learned to minimize the overall reconstruction error, considering both short-term and long-term loss during the fusion process. Our system is empirically shown to be accurate, robust, and efficient. As far as we know, this is the first work to integrate deep neural networks into reconstruction using a reinforcement learning approach.

In brief, this paper makes the following contributions:

- the first efficient hierarchical deep reinforcement network integrated with real-time dynamic multi-view 3D reconstruction,
- a reinforcement learning model for efficient and progressive selection of region to be fixed, and
- a deep neural network for high quality local reconstruction refinement.

2 Related work

2.1 Dynamic reconstruction

Inferring dynamic scene geometry remains an open research topic. Some work [18, 31–34] adopts strong semantic scene priors (e.g., a human body or hand template) to facilitate accurate correspondence and registration. Other methods [14, 35–37] instead choose to aggregate and denoise geometry in a canonical static space, and only track the per-frame deformation field over time, without knowledge of the reconstructed scene beforehand. To tackle the inherent ambiguity of the deformation field, and achieve better reconstruction fidelity, Refs. [5, 6, 36] introduce sparse image feature tracking, silhouette constraints, and albedo inference into the non-linear optimization to make tracking more robust, while Refs. [38, 39] bypass the correspondence estimation stage by imposing divergence constraints over the entire deformation vector field, and Refs. [17, 40] give dedicated hardware designs for obtaining cleaner and more complete depth and texture information. Readers are referred to Ref. [41] for a comprehensive literature review. Our approach introduces deep neural models to efficiently learn geometric priors from data for higher fusion quality.

2.2 Point set deep networks

For our surfel-based representation of the reconstructed geometry, we apply deep networks which directly consume point clouds. PointNet [43] and its variants [44–46] are a standard choice for encoding point set features while providing a good description of multi-scale details. The work in Ref. [47] is the first point set decoder combining fully-connected and deconvolution layers. In order to enforce a uniform structure onto the generated point set, FoldingNet [48] and AtlasNet [49] use one or more uniform grids to condition the shape descriptor for shape generation. The designs of

various deep point set networks support a variety of applications in both graphics and vision, such as point upsampling [29, 50] and shape completion [23, 51].

2.3 Deep reinforcement learning

Traditional reinforcement learning aims to learn from past experience and make better decisions in a principled way. The successful combination of deep neural networks and reinforcement learning algorithms is capable of dealing with higher-dimensional state and action spaces which were previously intractable [52]. Deep reinforcement learning has various applications in video games [42], generating animation [53], and indoor navigation [54]. A prominent approach is provided by the Deep Q-Network (DQN) [42] and its variants [55, 56], which approximate value functions with off-policy learning. Another line of approaches is based on policy gradients or actor-critics, where the model directly learns a stochastic policy [57, 58]. Our work formulates dynamic reconstruction as a Markov decision process and applies DQN to learn how to achieve minimum reconstruction error. We believe this is the first application of deep reinforcement learning in dynamic reconstruction pipelines.

3 Overview

Our HDR-Net-Fusion takes sequential depth maps captured using several commodity RGBD cameras as input and progressively reconstructs the geometry of the dynamic scene for every frame. As shown in Fig. 2, during the testing phase of our algorithm, for each incoming frame, the warping field which best aligns the current depth observations and the reference geometry is first found, and then a traditional fusion process is applied by our basic reconstruction system (Section 5). After that, the Global-HDR-Net is applied to the embedded deformation nodes to compute an expected reward for each node (Section 6.3). The node with the highest expected reward is selected and the local surfel patch surrounding that deformation node is fed into the Local-HDR-Net, which locally refines the patch geometry and completes missing areas (Section 6.2). The refined patch is then integrated into the reference geometry maintained by the reconstruction system to improve the quality of reconstruction and assist future tracking and registration.

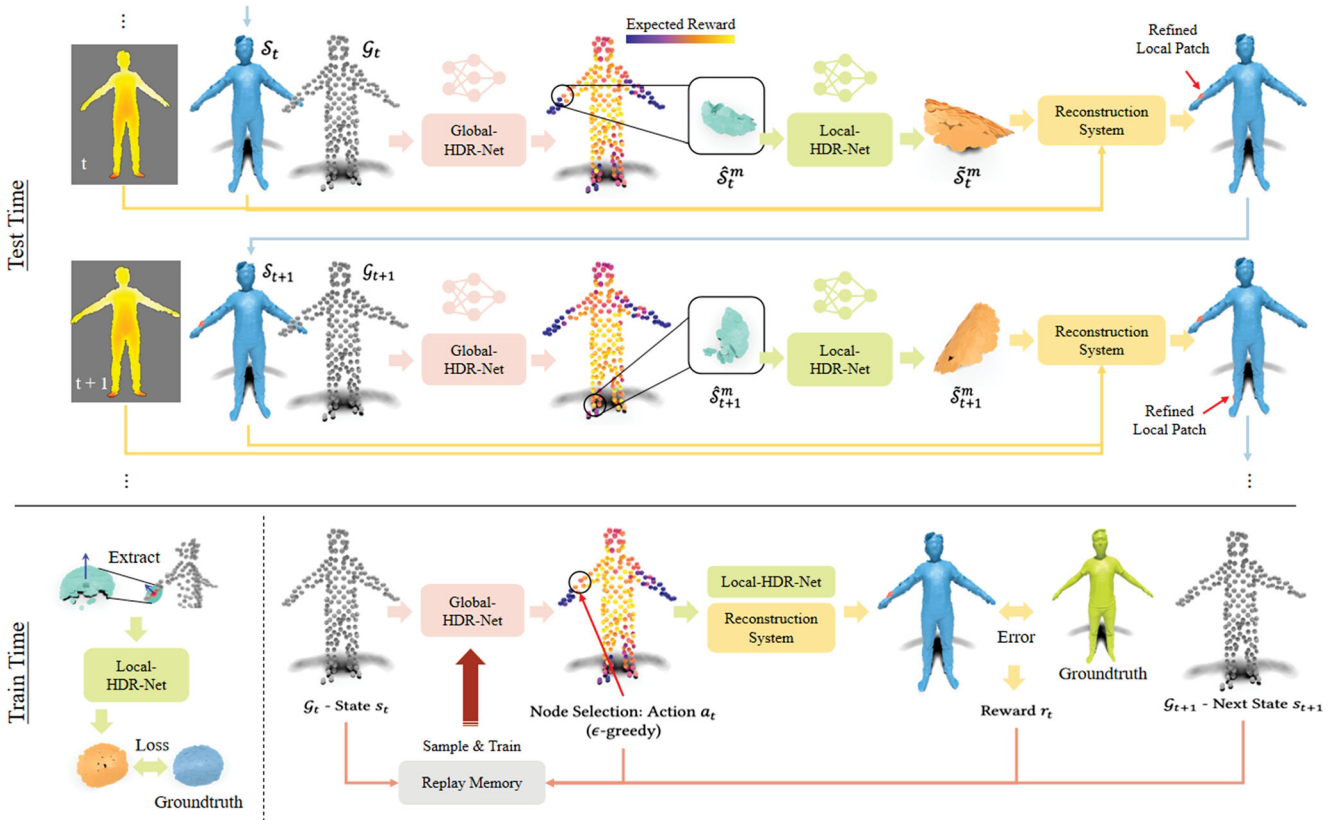


Fig. 2 Overview of our reconstruction system. During testing, the deformation nodes \mathcal{G}_t of the live geometry \mathcal{S}_t are fed into the Global-HDR-Net; the local patch $\hat{\mathcal{S}}_t^m$ with the highest expected reward is fed into the Local-HDR-Net for refinement. The refined patch $\tilde{\mathcal{S}}_t^m$ replaces the original geometry and is fused into the whole model, which is used for registering the next incoming frame. To train the global and local hierarchical networks, we first supervise Local-HDR-Net with groundtruth full patches, and then we fix its weight and train the Global-HDR-Net represented as a point-set-based DQN [42].

4 Notation and scene representation

The dynamically reconstructed scene is represented by a set of deformation nodes $\mathcal{G} = \{\mathbf{g}^m \in \mathbb{R}^3\}$ and a set of surfels with neighborhoods $\mathcal{S} = \{s^i, N^i\}$. \mathcal{S} is a dense reconstruction of the entire scene, while the nodes in \mathcal{G} are scattered sparsely over the surface represented by the surfels. Each surfel $s^i = (\mathbf{p}^i, \mathbf{n}^i, r^i)$ is represented by its center $\mathbf{p}^i \in \mathbb{R}^3$, normal $\mathbf{n}^i \in \mathbb{R}^3$, and radius $r^i \in \mathbb{R}$. A neighbourhood set $N^i \subset \mathcal{G}$ is attached to each surfel, initialized as the nearest K neighbours of s^i among all deformation nodes \mathbf{g}^m in \mathcal{G} . Similarly, a neighbourhood set $N^m \subset \mathcal{G}$ can be built for all the deformation nodes to establish their spatial relationships.

For each frame t we compute a warp field $\mathcal{W}_t = \{q_t^m \in \mathbb{SE}(3)\}$ defined at each node in \mathcal{G} , where q_t^m is the transformation applied to \mathbf{g}^m in frame t ; it is represented using dual quaternions [59]. Let $\mathcal{G}^t = \{\mathbf{g}_t^m \in \mathbb{R}^3\}$ be the transformed version of \mathcal{G} where $\mathbf{g}_t^m = q_t^m \cdot \mathbf{g}^m$. The surfels skin the deformation nodes

and the transformation for each surfel is found by interpolating nearby node transformations as $\hat{q}_t^i = \sum_{m \in N_i} w_{im} q_t^m$, where $w_{im} = \exp(-\|\mathbf{p}^i - \mathbf{g}^m\|_2^2 / \sigma^2)$, σ representing the node sampling distance [37]. We denote the transformed version of the surfels at frame t after applying \hat{q}_t^i by $\mathcal{S}_t = \{s_t^i, N^i\}$, $s_t^i = (\mathbf{p}_t^i, \mathbf{n}_t^i, r^i)$, where $\mathbf{p}_t^i = \hat{q}_t^i \cdot \mathbf{p}^i$, and \mathbf{n}_t^i is the normal \mathbf{n}^i transformed by the rotation part of \hat{q}_t^i only.

5 Basic reconstruction system

5.1 Design

The design of our basic reconstruction system is inspired by Ref. [37] and illustrated in Fig. 3. The initial surfels \mathcal{S} are called the *reference geometry* and the up-to-date surfels \mathcal{S}_t are called the *live geometry*. For each frame, \mathcal{W}_t is determined and the new surfels introduced in the current frame are appended to the live geometry; matched surfels are updated according to the running mean integration protocol [30]. The live geometry is then warped back




Fig. 3 Pipeline of the basic reconstruction system. For each incoming frame, we first forward warp the reference geometry into live geometry according to the current warp field, and then a new deformation field is found to align the reference geometry and depth observation. After that, the depth observation is fused with the new live geometry. Finally, we warp the live geometry back to the reference geometry.

to the reference geometry which provides a canonical shape representation.

5.2 Energy function

A key step in dynamic reconstruction is to find the per-frame warp field \mathcal{W}_t , which is solved by minimizing the following energy, consisting of a data term, a correspondence term and a regularization term:

$$E(\mathcal{W}_t) := E_{\text{data}}(\mathcal{W}_t) + \lambda_c E_{\text{cor}}(\mathcal{W}_t) + \lambda_r E_{\text{reg}}(\mathcal{W}_t) \quad (1)$$

where λ_c and λ_r are balancing weights. The data term:

$$E_{\text{data}}(\mathcal{W}_t) := \sum_{v=1}^V \sum_{i \in \mathcal{V}_t^v} \|\mathbf{n}_{d_t^{i,v}}^T (\mathbf{p}_t^i - \mathbf{p}_{d_t^{i,v}}^i)\|^2 \quad (2)$$

is a depth-to-plane ICP error summed across all V input views per frame, where \mathcal{V}_t^v is the visible surfel index set for the current \mathcal{W}_t from the v -th view, and $d_t^{i,v}$ is the corresponding depth observation of the i -th surfel found by re-projecting s_t^i into the v -th camera view, transformed by the camera extrinsic. The correspondence term is

$$E_{\text{cor}}(\mathcal{W}_t) := \sum_{v=1}^V \sum_{(s_t^i, u_t^i) \in \mathcal{C}_t^v} \|\mathbf{p}_{u_t^i}^v - \mathbf{p}_t^i\|^2 \quad (3)$$

which is a distance between two sparsely related points found by global patch collider [60] and \mathcal{C}_t^v is the correspondence set containing tuples of matched surfels s_t^i and pixels u_t^i for the v -th view. $\mathbf{p}_{u_t^i}^v$ is the 3D point re-projected from pixel u_t^i in view v . The regularization term:

$$E_{\text{reg}}(\mathcal{W}_t) := \sum_{\mathbf{g}^m \in \mathcal{G}} \sum_{n \in \mathcal{N}^m} \|\mathbf{g}_t^m - \mathbf{q}_t^n \cdot \mathbf{g}^m\|_2^2 \quad (4)$$

is an as-rigid-as-possible constraint encouraging nearby nodes to share the same transformation.

5.3 Challenges

Generally, our basic reconstruction system works well on simple datasets; special cases like topology changes and tracking failure can be fixed by re-initialization [37]. However, without any prior knowledge about the dynamic scene structure, it is still very challenging to track fast motions and much information is lost during re-initialization. The situation becomes worse when there are erroneous depth observations or when parts of the dynamic structure are occluded in any views.

6 HDR-Net: Hierarchical deep reinforcement network for dynamic reconstruction

6.1 Concept

To address the challenges faced by the traditional reconstruction system discussed in Section 5, repair of the erroneous and occluded parts is necessary, while the efficiency of the reconstruction system should also be guaranteed. We thus propose a hierarchical reinforcement network (HDR-Net) that first finds the to-be-fixed regions (Global-HDR-Net) efficiently using reinforcement learning algorithms and then fixes these regions by exploiting the power of deep neural network (Local-HDR-Net).

6.2 Local-HDR-Net

6.2.1 Network architecture

For each frame, given a selected deformation node \mathbf{g}_t^m from Global-HDR-Net, we gather all surfels influenced by that node as $\hat{\mathcal{S}}_t^m := \{s_t^i | \mathbf{g}_t^m \in \mathcal{N}^i\}$ and feed that local patch into the Local-HDR-Net. Its job is to generate $\hat{\mathcal{S}}_t^m$, a completed and de-noised version of $\hat{\mathcal{S}}_t^m$. Two design requirements exist for our model: (i) as the network is applied to the reconstructed geometry on a per-frame basis, the model should be lightweight, requiring minor additional computation, and (ii) in order to resolve the inherent ambiguity of point set completion, knowledge of the entire scene geometry should be taken into consideration.

We therefore propose a hybrid encoder-decoder structure using the order-agnostic PointNet [43] as the backbone, as shown in Fig. 4. To integrate global geometric knowledge, we use \mathcal{G}_t as a summary of the current coarse shape: it gives a good global shape approximation which can effectively summarize the overall scene structure to the network.

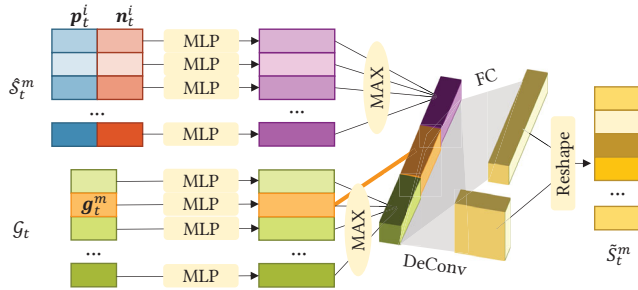


Fig. 4 Local-HDR-Net model. The network takes the local surfel patch and all deformation nodes as input. The concatenated latent feature vector is decoded using fully-connected and deconvolution layers.

In the encoder part of our model, G_t and \hat{S}_t^m are first encoded separately, extracting features with respective point-shared MLPs. As the encoded feature of each point in G_t , we take its globally aggregated feature vector as well as the point feature vector for g_t^m . These two feature vectors are then concatenated with the aggregated per-point feature of \hat{S}_t^m . The overall aggregated latent representation of the local region now contains information summarizing the patch geometry in its global context.

In the decoder, we find that using a classic fully connected and deconvolution combination [47] generates the best results while still allowing real-time processing. Deformation-based decoders [48] easily lead to over-smoothed surfaces lacking detail, while implicit-function-based decoders [22] involve heavy sampling computations during inferencing. The direct output of our model is simply the 3D surfel center position's offset to the selected node \tilde{o}_t^i , which is easier to learn than the surfel normal, given its spatial continuity.

6.2.2 Loss function

We use the earth-mover distance (EMD) as the loss function to train the network:

$$\mathcal{L} := \frac{1}{|\tilde{S}_t^m|} \min_{\phi: \tilde{\mathbf{p}} \rightarrow \tilde{\mathbf{p}}} \sum_{s_t^i \in \tilde{S}_t^m} \|\tilde{\mathbf{p}}_t^i - \phi(\tilde{\mathbf{p}}_t^i)\|_2^2 \quad (5)$$

where $\tilde{\mathbf{p}}_t^i = \tilde{o}_t^i + \mathbf{g}_t^m$ is the center position of each output surfel in world coordinates. ϕ is a bijection; the best linear assignment expressed by the min operator can be computed efficiently using the approach in Ref. [47]. $\tilde{\mathbf{p}}$ is the groundtruth surfel center.

6.3 Global-HDR-Net

6.3.1 Problem formulation

For each frame t , Global-HDR-Net aims to select the node \mathbf{g}_t^m for the subsequent local patch refinement

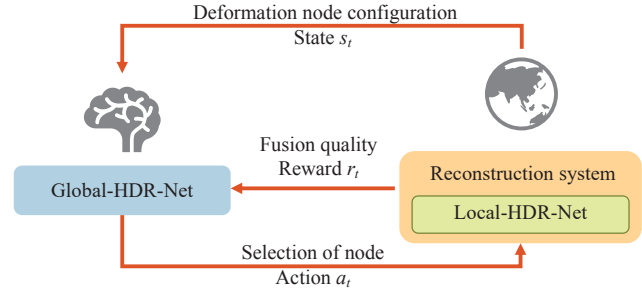


Fig. 5 Reinforcement learning for training Global-HDR-Net. The reconstruction system as well as Local-HDR-Net serve as the environment while Global-HDR-Net is the agent whose task is to select a deformation node in G_t to be refined by the Local-HDR-Net for each frame.

operation described in Section 6.2. One could choose to refine more than one node, or in the extreme case, all nodes in one single frame. However, too many passes of network inferencing will drastically affect the system's real-time performance. In fact, as the output reward is not expected to vary much given subtle changes in the input node positions, all nodes with high reward will be eventually picked up in time based on the high capturing frame rate. By instead performing inferencing on one node at a time, we distribute the computation across the entire session so that speed is guaranteed while still not preventing any nodes from being chosen.

One simple strategy for this module could be to always select S_t^m with the worst geometric quality. However, a greedy algorithm will not necessarily lead to a globally optimal result as it does not consider possible future registration error and the empirical performance of Local-HDR-Net. We instead pursue an algorithm that is aware of both short-term and long-term reconstruction quality and takes the properties of both the underlying dynamic reconstruction system and Local-HDR-Net into account.

We solve this problem using ideas from reinforcement learning (RL) which implicitly model the environment using existing experience gained through trial and error. A natural analogy can be made between Global-HDR-Net and a reinforcement learning agent. The dynamic reconstruction system and the local net serve as the environment, which receives an *action* (a deformation node \mathbf{g}_t^m) from the network, performs internal fusion and local patch refinement, and emits the reconstructed result as the new observation. The rewards for

the action performed can be modeled by the score of reconstruction quality. By choosing different actions at each timestamp, the Global-HDR-Net agent influences internal state of the system and all succeeding reconstruction steps.

The target of RL is to learn an optimal policy which can be later executed during inferencing. The optimal policy maximizes the expected return along the state transition path, which, in our case, effectively minimizes reconstruction error over all time steps.

From a theoretical point of view, two key propositions have to be met for the above formulation to be meaningful. Firstly, the reconstruction system should obey the Markov property, where the state of the current step is solely dependent on the previous step's state. Secondly, an appropriate choice of the deformation node can be made solely from the configuration of \mathcal{G}_t . The first proposition is naturally satisfied because for each frame, the depth observation is integrated only with the fused geometry from the previous frame. Also, we have found that regions with poor reconstruction quality often have highly complicated or mostly occluded parts, which to a certain degree justifies us in assuming the second proposition to be true.

6.3.2 Learning algorithm and network architecture

We employ DQN [42], which uses an efficient off-policy value-function based approach, as our reinforcement learning algorithm. DQN aims to learn the Q-function (expected reward given state and action) through past experience, and approximates $Q(s_t, \cdot)$ using a deep neural network (i.e., the Q-Network) to model the high-dimensional state and action space. Here we use s_t and a_t to denote the state and action for frame t . Specifically, s_t represents the positions of global nodes \mathcal{G}_t up to frame t and a_t is the integer index m of the selected deformation node in \mathcal{G}_t . By enforcing Bellman equality and minimizing temporal difference error δ_t , the Markov process of the environment can be precisely modeled by the Q-function and our final policy can be greedily selected as $\pi^*(s_t) := \arg\max_a Q(s_t, a)$ so that in each frame t we maximize $\sum_{t'=t}^T \gamma^{t'-t} r_{t'}$, where $\gamma > 0$ is the discount factor, r_t is the reward for frame t and T is the number of total frames.

Following Ref. [42] we define the temporal difference error as

$$\delta_t(\Theta) := Q(s_t, a_t; \Theta) - (r_t + \gamma \max_a Q(s_{t+1}, a; \Phi)) \quad (6)$$

where Θ is the parameter of the policy deep network and Φ is the parameter of the target deep network. During training, we execute the reconstruction system with Local-HDR-Net several times and gather (s_t, a_t, r_t, s_{t+1}) tuples. Here, the actions a_t are chosen using an ϵ -greedy policy which interpolates between the currently best found policy π^* and a completely random policy with factor ϵ . It can be proved that this strategy converges to an optimal policy, balancing exploration and exploitation in state space. We store multiple state-action-reward tuples across different episodes in a common replay memory. Mini-batches are then sampled from the replay memory to train the policy network parameter Θ using back propagation, so that $\delta_t^2(\Theta)$ is minimized. Φ is usually fixed and updated to Θ only every few episodes to guarantee stable training.

We choose PointNet++ [44] as our Q-network. It takes in point set \mathcal{G}_t at frame t and the numbers of local surfel patches $\mathcal{S}_t^m (m \in \mathcal{G}_t)$ as input and predicts the Q-value, i.e., the expected reward for each point (possible action).

The reward r_t is taken as the negative chamfer distance \mathcal{D} , defined as

$$\begin{aligned} \mathcal{D} := & \frac{1}{|\mathcal{S}_p|} \sum_{\mathbf{p}^i \in \mathcal{S}_p} \min_{\mathbf{p}^j \in \mathcal{S}_g} \|\mathbf{p}^i - \mathbf{p}^j\|_2 \\ & + \frac{1}{|\mathcal{S}_g|} \sum_{\mathbf{p}^j \in \mathcal{S}_g} \min_{\mathbf{p}^i \in \mathcal{S}_p} \|\mathbf{p}^i - \mathbf{p}^j\|_2 \end{aligned} \quad (7)$$

where \mathcal{S}_p are the surfel positions of the current geometry using the reconstruction system and \mathcal{S}_g is the groundtruth reference geometry.

7 Results and discussions

In this section, we introduce the experimental setup for implementing our system, give results and comparisons, and validate the design of our method.

7.1 Experimental setup

7.1.1 Dataset

Our experiments used sequences from the Human10 dataset [4] to test our algorithm. This dataset contains 10 long sequences of several human actors performing various actions, of which 9 are publicly usable. In Human10, each sequence was recorded using 4 fixed-position 512×512 resolution RGB-D cameras distributed uniformly around a 360° viewing circle. The frustum of each depth camera covers

a partial view of the entire human body. The limited sensor quality, leading to severe depth error and loss, and many fast large motions as well as topology changes, present very challenging data to the reconstruction system, resulting in very frequent tracking loss and re-initialization. To measure reconstruction quality, the dataset provides a groundtruth 3D mesh reconstructed using a free-viewpoint-video [61] capture system.

To verify the generalization of the network, we split the 9 sequences into training sequences (human1/2/4/6/9) for HDR-Net and testing sequences (human0/3/7/8). Weights for both Local-HDR-Net and Global-HDR-Net were learned and cross-validated solely from the training frames. In order to train Local-HDR-Net, patch-level surfel and deformation node data were generated. We first generated surfels and nodes from depth observations every single frame without warping, to simulate the artifacts caused by re-initialization and camera quality. On the other side, we used Poisson disk sampling to sample equally-spaced surfels over the groundtruth mesh. We then gathered surfels from both the reconstruction results and the groundtruth surrounding each node to form a complete patch, using a ball query, forming a local patch training pair for the supervised learning of patch completion. Additionally, we balanced the distribution of training pairs by their *completeness* score, defined as the portion of groundtruth surfels closer than a certain threshold to its nearest neighbour in the partial surfel patch (extracted from the input depth map). Empirically we found better overall system performance can be achieved with this balanced dataset, most of whose training pairs would otherwise be almost complete.

7.1.2 Training protocol

We use a common supervised training strategy to optimize the Local-HDR-Net using the dataset described above: an AdaGrad optimizer is used with a learning rate of 10^{-3} . The training of Global-HDR-Net is based on 200 randomly selected consecutive frames from each episode. For each frame, we randomly sample a state-action-reward tuple batch from the replay memory and optimize Θ with the RMSprop optimizer. The network weight Φ is updated to Θ every 3 episodes. During execution of the ϵ -greedy policy, we start with 90% probability of

selecting random nodes and decrease the probability exponentially to 5% with a 200 frame decay rate. The discount factor is set to $\gamma = 0.999$. In total, we train for around 160 episodes to get a fairly convergent result. The loss curve is shown in Fig. 6.

To ensure fair evaluation of system performance, all input frames are never seen by either network. Specifically, among the only four sequences (human0/3/8/9) which contain RGB information and can achieve a good result in multi-view sequences' tracking, we choose human9 for training Global-HDR-Net. Considering the dependency of Global-HDR-Net's training on Local-HDR-Net's performance, Global-HDR-Net would learn non-generalizable policies if Local-HDR-Net is too familiar with the sequence we train Global-HDR-Net on. Therefore, we only take human1/2/4/6 for Local-HDR-Net training and evaluation, excluding human9.

7.1.3 Implementation details

We implemented our multi-view reconstruction system in C++/CUDA. The training code for both Local-HDR-Net and Global-HDR-Net is written using PyTorch. We interfaced the reconstruction system and the deep network so that both training and inferencing are tightly coupled and trained effectively end-to-end. The algorithm was tested on a workstation with an Nvidia Titan RTX graphics card, running the reconstruction pipeline, Local-HDR-Net and Global-HDR-Net simultaneously.

In practice, we find that Local-HDR-Net does not guarantee perfectly smoothed output, which may degrade overall system performance. Hence we separately adopt a post-processing step to directly reject bad local net output $\tilde{\mathcal{S}}_m^t$. This post-rejection step finds all nearest neighbour pairs in $\tilde{\mathcal{S}}_m^t$ and computes the dot product of normal vectors of the pair. This generated surfel patch is rejected if the




Fig. 6 Variation of the average reward per epoch during the training process of Global-HDR-Net.

mean value of all dot products is less than ε , which can guarantee the smoothness of the surfels we add.

In addition, the number of fixed nodes in each frame can be adjusted in practice to achieve better quality. Since the Global-HDR-Net can output the expected reward of all nodes, we can select several nodes within an acceptable range instead of the best one and fix them one by one so that we can refine a necessary number of nodes every frame while still enabling real-time performance.

In the presence of fast motion, the tracking module may fail. We detect this abnormality by checking the residuals of the registration solver and perform re-initialization [5, 37].

7.1.4 Parameter selection

Parameter choice for the reconstruction system is application dependent. To effectively track and recover human geometry, we empirically set the reconstruction parameters in Section 5 to $\lambda_r = 2.3$

and $\lambda_c = 1.3$ while the node sampling distance [37] $\sigma = 0.04$ cm. The post-processing rejection threshold ε is set to 0.9.

In terms of network structure, our Local-HDR-Net encodes surfels in $\hat{\mathcal{S}}_t^m$ using a shared MLP with sequentially 32, 64, and 256 channels, encodes each point in \mathcal{G}_t with 32, 64, 256 channels, and transforms the concatenated latent feature into two patches of 256 3D points with FC and DeConv layers separately. The Global-HDR-Net downsamples input deformation nodes into 256, 64, and 16 points sequentially with set abstraction layers and the local feature is interpolated using a feature propagation layer. The input nodes \mathbf{G}_t are padded to a minimum size of 400.

7.2 Overall performance

We present some qualitative results for our entire HDR-Net-Fusion framework in Fig. 7. Compared to results without HDR-Net, combining our hierarchical



Fig. 7 Selected frames demonstrating the overall reconstruction quality of HDR-Net-Fusion. Top to bottom: human0, human3, human8 from Human10. Each pair shows the reconstructed result without (blue) and with (red) HDR-Net. Our algorithm successfully identifies missing or noisy regions and refines them reasonably.

network can effectively complete and refine missing or inaccurate regions of the fused model. Leveraging the geometric prior of the underlying scenes using our carefully-designed deep network, a plausible completion can be generated, filling in holes in occluded regions (e.g., body parts partially hidden by moving arms) or wrongly-observed regions (e.g., regions with dark hair whose depth cannot be accurately measured by the sensor). During re-initialization caused by large registration errors, most of the fused model is deleted, which can be quickly fixed by our model and subsequent registration artifacts can be minimized. Table 1 compares the number of re-initializations required by systems with and without HDR-Net. When there are frequent re-initializations caused by large motions, our repair can prevent further subsequent registration artifacts, leading to a significant reduction in the number of re-initializations.

A close-up of the reconstructed geometry is shown in Fig. 8. The regions of the actor's head and shoulders contain holes as a result of the erroneous observation depth, while the region behind his left arm is empty because of a recent re-initialization. Both of the artifacts can be effectively fixed by our method.

The behaviour of our algorithm can be further analyzed by visualizing the expected reward computed

Table 1 Number of re-initializations during reconstruction by our framework without or with HDR-Net. Our method can reduce re-initialization especially when re-initialization is frequent

Sequence	Without HDR-Net	With HDR-Net
Human0	20	13
Human3	15	11
Human8	7	6




Fig. 8 Close-up comparison of geometry reconstructed by the basic method and our method. HDR-Net can resolve both artifacts caused by erroneous depth observations and incompleteness caused by re-initialization.

by the global net, answering the questions of what the Global-HDR-Net has learned and why it is useful in our setting. As shown in Fig. 9, Global-HDR-Net can find places with holes and bumps efficiently and accurately. In addition, it tends to repair regions of the model's boundaries such as the shoulders and feet. This is valid in the sense that these parts are more likely to move rapidly later and need to be refined to make the tracking more reliable. Otherwise, it will be harder to track a broken arm, as we can see in Fig. 7, leading to frequent re-initialization as shown in Table 1. Presumably, as the input to our global net only contains the deformation nodes for the current frame, the model implicitly learns to predict the potential node motions given the static pose and jointly considers both spatial and temporal cues when making decisions. Again, the policy is implicitly learned for lower reconstruction error, which is hard for a hand-crafted heuristic to imitate as demonstrated in Section 7.5.

7.3 Speed

Our reconstruction system takes about 25 ms per frame for a single-view sequence of Human10, and more time for the image pre-processing as it uses multi-view sequences, which can be parallelized if there are multiple processors. The average inferencing time is about 2 and 4.5 ms for Local-HDR-Net and Global-HDR-Net, respectively, adding little overhead to the underlying reconstruction. This is due to the lightweight design of our deep models and the scalable surfel representation. In conclusion, our system can reach 25 Hz with 5 nodes fixed per frame. Taking the parallel running of Local-HDR-Net into consideration, the process could be made even faster.




Fig. 9 Expected reward for the deformation nodes computed by Global-HDR-Net (superimposed on the surfels). Since the Q-value's interval for each sequence is not uniform, we use relative coloring, nodes in blue and yellow having relatively lower and higher values.

7.4 Comparison with the traditional method

To demonstrate the advance of our framework over traditional reconstruction methods, we also reconstructed the test data with SurfelWarp [37] for each single perspective to make a comparison with our method. Figure 10 presents several selected frames from the sequences reconstructed by both SurfelWarp and our system. The result shows our method can refine holes and bumps effectively, which is impossible for traditional methods without correct input. Therefore, our system can give a reconstruction result with higher quality than



Fig. 10 Selected frames from the test set demonstrating the single-view performance of our framework (red) compared with Ref. [37] (blue). Our method provides results with richer features and fewer artifacts.

traditional methods when there are heavy noise and erroneous depth observations. In addition, broken or disconnected legs and arms are quickly completed, confirming our conclusion in Section 7.2.

7.5 Global-HDR-Net comparison

To test the performance of Global-HDR-Net and to show that our network actually learns effective information during its training, we set up two competing agents executing different policies.

- *Random*: nodes are uniformly sampled from \mathcal{G}_t .
- *Heuristic*: we first remove all candidate nodes not satisfying the following criteria:
 - the number of surfels related to the node should be greater than 20;
 - the mean Euclidean distance from each surfel to the node should be smaller than 0.08;
 - the surfel confidence maintained within the reconstruction system, representing the point's stability and reliability, should be greater than 2.0.

The above criteria ensure Local-HDR-Net acquires sufficient information for inferencing.

Then the eligible node with fewest surfels is selected by this policy since patches with fewer surfels should be given higher priority for refinement.

The performance of the policies are compared using the quality of the reconstructed model computed using two-way chamfer distance as defined in Eq. (7). Figure 11 shows both qualitative and quantitative




Fig. 11 Comparisons of different policies of node selection on the human0/3 from Human10. The leftmost figures plot sorted errors across the entire test sequence (i.e., the smaller the better). The error map based on the groundtruth model is placed on the front of each geometry reconstructed by different policies.

comparisons over two of the Human10 sequences. Results show that the manually designed heuristic policy leads to better reconstruction quality than the random policy most of the time, but the effect is not strong or particularly stable. Interestingly, we find that the random policy can sometimes lead to worse results than the simple reconstruction system without Local-HDR-Net. This is because that some randomly selected nodes may contain too much noise or have a low completeness score, so their corresponding complete geometry is too challenging for Local-HDR-Net to recover, generating many noisy outliers. In the contrast, our policy provides an effective refinement to the geometry. Clearly, choosing the correct node is as important as the geometry refinement process, which needs careful handling.

Compared to our policy, which outperforms all baselines and is learned with the direct goal of minimizing reconstruction error, the heuristic policy is the closest competitor but is unaware of the behaviour of Local-HDR-Net and the underlying reconstruction system. It is non-trivial to manually build a spatiotemporally aware criterion as analyzed in Section 7.2.

7.6 Local-HDR-Net comparison

Local-HDR-Net mainly focuses on completing and refining local patch geometries. We compare it with three baselines:

- FoldingNet [48], whose decoder is designed by concatenating sampled points on a uniform grid with the global feature vector. The network learns how to deform such a uniform grid to the desired shape; surface smoothness is guaranteed.
- The Point Completion Network [23], employing a coarse-to-fine completion strategy and aggregating multiple deformable grids to assemble the final completed shape.
- A variant of Local-HDR-Net, lacking the branch taking in the scene deformation nodes \mathcal{G}_t . This variant is used to test the utility of the scene structural guidance.

All baseline models and Local-HDR-Net were trained for 200 epochs. Network hyper-parameters including model architecture and learning strategies were separately tuned for each model to give best cross-validation performance.

For evaluation we use the earth mover distance (Eq. (5)) between prediction and groundtruth surfel

positions as our metric. Distance error is plotted for different ranges of patch completeness (from 0.0 to 1.0) in Fig. 12. As there will be much noise and error in the actual reconstruction, we also tested the performance of the networks when Gaussian noise with a standard deviation of 5 mm is added in the normal direction of each surfel of the input. For patches with low completeness (< 0.1), most methods find it challenging to infer missing fine details due to loss of information. In addition, patches with too much noise will also lead to a bad result. This also justifies our strategy of rejecting patches with too few surfels and returning a negative reward: this explicitly discourages the global net from choosing overly challenging patches for the local net.

Compared to the baselines, our Local-HDR-Net yields much smaller distance errors in most cases and exhibits better stability in the case of noisy input. Figure 13 renders results qualitatively in the form of surfels. FoldingNet [48] can generate smooth surfaces, but its uniform grid parameterization leads to distorted boundaries and the deformation in complex areas is unnatural. The Point Completion Network [23] assembles the surface from many smaller patch grids, resulting in overlaps and uneven distribution of surfels. It is unsuitable for small-scale geometry refinement. Our baseline without nodes completely discards the scene structural guidance, i.e., the global context, which is very important when the patches' information is very limited or contains much noise.




Fig. 12 Comparison with different Local-HDR-Net baselines for patches with different completeness scores. The last two columns show the average EMD loss of all patches without and with artificial noise, respectively.

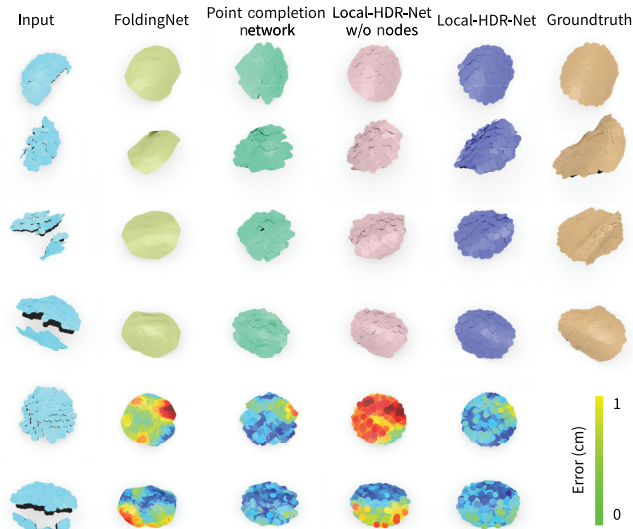


Fig. 13 Qualitative comparison to Local-HDR-Net baselines using randomly selected patches from our dataset. Error maps on the last two rows are overlaid on the predicted patch \hat{S}_t^m .

7.7 Generalization

To demonstrate the ability of our system to generalize on real world scenes, we also selected some sequences from the DeepDeform dataset [7] and sequence human7 from the Human10 dataset [4] to test our method. Results are shown in Fig. 14. Our method can fix artifacts and achieve better results than a traditional method [37]. Furthermore, in addition to the human body, our system can fix the artifacts in other objects as well.

7.8 Limitations

There are two typical limitations of our method: Firstly, as shown in Fig. 15(a), although our Local-



Fig. 14 Reconstruction results of sequences from DeepDeform dataset [7] and Human10 dataset [4] (bottom-left two). Compared to a traditional method [37] (blue), our method (red) can complete objects other than the human body and real scenes including interaction of people and things like a bag or ball.




Fig. 15 Limitations of our method. In each pair we show the original reconstructed model in blue and ours in yellow.

HDR-Net can provide good refinement for most model parts, it still remains challenging to learn features of complicated and subtle structures like hands. A possible reason is that it is hard for the network to represent all features of such a large batch. Narrowing the range of a single patch when the structure is complicated may provide a better result. Furthermore, a more powerful deep neural network with self-attention mechanisms [62] could be adopted to learn more discriminative features for point cloud completion.

Secondly, as shown in Fig. 15(b), selecting a constant number of nodes per frame can lead to problems since the model's completeness is continuously changing during reconstruction. For some models already sufficiently complete, Global-HDR-Net will still select some completed patches to refine, resulting in computational inefficiency and even leading to worse results. A more adaptive node selection strategy could be applied by rejecting previously chosen nodes or selecting nodes by considering their predicted completeness.

8 Conclusions

This paper has presented HDR-Net-Fusion, a novel dynamic reconstruction system using a hierarchical deep reinforcement network to improve reconstruction quality. Its applicability and effectiveness have been experimentally demonstrated using a large-scale dynamic fusion dataset. Our approach formulates the global selection of a local geometric patch for refinement in terms of reinforcement learning and uses a point-based neural network to complete and improve the local geometry. We hope this work can inspire future work pursuing better dynamic

reconstruction quality using powerful deep learning and reinforcement learning algorithms.

Acknowledgements

We thank the anonymous reviewers for their helpful comments on this paper. This work was supported by the National Natural Science Foundation of China (Grant Nos. 61902210 and 61521002).

Electronic Supplementary Material Supplementary material is available at the online version of this article at <https://doi.org/10.1007/s41095-021-0230-z>.

References

- [1] Newcombe, R. A.; Davison, A. J.; Izadi, S.; Kohli, P.; Hilliges, O.; Shotton, J.; Hodges, S.; Fitzgibbon, A. W. KinectFusion: Real-time dense surface mapping and tracking. In: Proceedings of the 10th IEEE International Symposium on Mixed and Augmented Reality, 127–136, 2011.
- [2] Whelan, T.; McDonald, J. B.; M. Kaess, M.; M. F. Fallon, M. F.; Johannsson, H.; Leonard, J. J. Kintinuous: Spatially extended KinectFusion. In: Proceedings of the RSS Workshop on RGB-D: Advanced Reasoning with Depth Cameras, 2012.
- [3] Nießner, M.; Zollhöfer, M.; Izadi, S.; Stamminger, M. Real-time 3D reconstruction at scale using voxel hashing. *ACM Transactions on Graphics* Vol. 32, No. 6, Article No. 169, 2013.
- [4] Liu, Z. N.; Cao, Y. P.; Kuang, Z. F.; Kobbelt, L.; Hu, S. M. High-quality textured 3D shape reconstruction with cascaded fully convolutional networks. *IEEE Transactions on Visualization and Computer Graphics* Vol. 27, No. 1, 83–97, 2021.
- [5] Dou, M. S.; Khamis, S.; Degtyarev, Y.; Davidson, P.; Fanello, S. R.; Kowdle, A.; Orts-Escolano, S.; Rhemann, C.; Kim, D.; Taylor, J. et al. Fusion4D: Real-time performance capture of challenging scenes. *ACM Transactions on Graphics* Vol. 35, No. 4, Article No. 114, 2016.
- [6] Dou, M. S.; Davidson, P.; Fanello, S. R.; Khamis, S.; Kowdle, A.; Rhemann, C.; Tankovich, V.; Izadi, S. Motion2fusion. *ACM Transactions on Graphics* Vol. 36, No. 6, Article No. 246, 2017.
- [7] Božič, A.; Zollhöfer, M.; Theobalt, C.; Nießner, M. DeepDeform: Learning non-rigid RGB-D reconstruction with semi-supervised data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 7000–7010, 2020.
- [8] Chen, K.; Lai, Y. K.; Hu, S. M. 3D indoor scene modeling from RGB-D data: A survey. *Computational Visual Media* Vol. 1, No. 4, 267–278, 2015.
- [9] Rünz, M.; Agapito, L. Co-fusion: Real-time segmentation, tracking and fusion of multiple objects. In: Proceedings of the IEEE International Conference on Robotics and Automation, 4471–4478, 2017.
- [10] Runz, M.; Buffier, M.; Agapito, L. MaskFusion: Real-time recognition, tracking and reconstruction of multiple moving objects. In: Proceedings of the IEEE International Symposium on Mixed and Augmented Reality, 10–20, 2018.
- [11] Huang, J. H.; Yang, S.; Zhao, Z. S.; Lai, Y. K.; Hu, S. M. ClusterSLAM: A SLAM backend for simultaneous rigid body clustering and motion estimation. *Computational Visual Media* Vol. 7, No. 1, 87–101, 2021.
- [12] Huang, J. H.; Yang, S.; Mu, T. J.; Hu, S. M. ClusterVO: Clustering moving instances and estimating visual odometry for self and surroundings. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2165–2174, 2020.
- [13] Du, Z. J.; Huang, S. S.; Mu, T. J.; Zhao, Q. H.; Martin, R.; Xu, K. Accurate dynamic SLAM using CRF-based long-term consistency. *IEEE Transactions on Visualization and Computer Graphics* doi: 10.1109/TVCG.2020.3028218, 2020.
- [14] Newcombe, R. A.; Fox, D.; Seitz, S. M. DynamicFusion: Reconstruction and tracking of non-rigid scenes in real-time. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 343–352, 2015.
- [15] Brown, B. J.; Rusinkiewicz, S. Global non-rigid alignment of 3-D scans. *ACM Transactions on Graphics* Vol. 26, No. 3, 21–es, 2007.
- [16] Orts-Escolano, S.; Rhemann, C.; Fanello, S. R.; Chang, W.; Kowdle, A.; Degtyarev, Y.; Kim, D.; Davidson, P. L.; Khamis, S.; Dou, M. V. et al. Holoportation: Virtual 3D teleportation in real-time. In: Proceedings of the 29th Annual Symposium on User Interface Software and Technology, 741–754, 2016.
- [17] Guo, K. W.; Lincoln, P.; Davidson, P.; Busch, J.; Yu, X. M.; Whalen, M.; Harvey, G.; Orts-Escolano, S.; Pandey, R.; Dourgarian, J. et al. The relightables: Volumetric performance capture of humans with realistic relighting. *ACM Transactions on Graphics* Vol. 38, No. 6, Article No. 217, 2019.
- [18] Yu, T.; Guo, K. W.; Xu, F.; Dong, Y.; Su, Z. Q.; Zhao, J. H.; Li, J.; Dai, Q.; Liu, Y. BodyFusion: Real-time capture of human motion and surface geometry using a single depth camera. In: Proceedings of the IEEE

- International Conference on Computer Vision, 910–919, 2017.
- [19] Yu, T.; Zhao, J. H.; Zheng, Z. R.; Guo, K. W.; Dai, Q. H.; Li, H.; Pons-Moll, G.; Liu, Y. DoubleFusion: Real-time capture of human performances with inner body shapes from a single depth sensor. *IEEE Transactions on Pattern Analysis and Machine Intelligence* Vol. 42, No. 10, 2523–2539, 2020.
- [20] Meerits, S.; Thomas, D.; Nozick, V.; Saito, H. FusionMLS: Highly dynamic 3D reconstruction with consumer-grade RGB-D cameras. *Computational Visual Media* Vol. 4, No. 4, 287–303, 2018.
- [21] Fujiwara, K.; Nishino, K.; Takamatsu, J.; Zheng, B.; Ikeuchi, K. Locally rigid globally non-rigid surface registration. In: Proceedings of the International Conference on Computer Vision, 1527–1534, 2011.
- [22] Park, J. J.; Florence, P.; Straub, J.; Newcombe, R.; Lovegrove, S. DeepSDF: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 165–174, 2019.
- [23] Yuan, W. T.; Khot, T.; Held, D.; Mertz, C.; Hebert, M. PCN: Point completion network. In: Proceedings of the International Conference on 3D Vision, 728–737, 2018.
- [24] Wang, Y.; Solomon, J. Deep closest point: Learning representations for point cloud registration. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 3522–3531, 2019.
- [25] Gojcic, Z.; Zhou, C. F.; Wegner, J. D.; Guibas, L. J.; Birdal, T. Learning multiview 3D point cloud registration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 1756–1766, 2020.
- [26] Gu, X. Y.; Wang, Y. J.; Wu, C. R.; Lee, Y. J.; Wang, P. Q. HPLFlowNet: Hierarchical permutohedral lattice FlowNet for scene flow estimation on large-scale point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 3249–3258, 2019.
- [27] Liu, X. Y.; Qi, C. R.; Guibas, L. J. FlowNet3D: Learning scene flow in 3D point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 529–537, 2019.
- [28] Xiao, Y. P.; Lai, Y. K.; Zhang, F. L.; Li, C. P.; Gao, L. A survey on deep geometry learning: From a representation perspective. *Computational Visual Media* Vol. 6, No. 2, 113–133, 2020.
- [29] Li, R. H.; Li, X. Z.; Fu, C. W.; Cohen-Or, D.; Heng, P. A. PU-GAN: A point cloud upsampling adversarial network. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 7202–7211, 2019.
- [30] Keller, M.; Lefloch, D.; Lambers, M.; Izadi, S.; Weyrich, T.; Kolb, A. Real-time 3D reconstruction in dynamic scenes using point-based fusion. In: Proceedings of the International Conference on 3D Vision, 1–8, 2013.
- [31] Garcia Cifuentes, C.; Issac, J.; Wüthrich, M.; Schaal, S.; Bohg, J. Probabilistic articulated real-time tracking for robot manipulation. *IEEE Robotics and Automation Letters* Vol. 2, No. 2, 577–584, 2017.
- [32] Tzionas, D.; Gall, J. Reconstructing articulated rigged models from RGB-D videos. In: *Computer Vision – ECCV 2016 Workshops. Lecture Notes in Computer Science*, Vol. 9915. Hua, G.; Jégou, H. Eds. Springer Cham, 620–633, 2016.
- [33] Taylor, J.; Bordeaux, L.; Cashman, T.; Corish, B.; Keskin, C.; Sharp, T.; Soto, E.; Sweeney, D.; Valentin, J. P. C.; Luff, B. et al. Efficient and precise interactive hand tracking through joint, continuous optimization of pose and correspondences. *ACM Transactions on Graphics* Vol. 35, No. 4, Article No. 143, 2016.
- [34] Schmidt, T.; Newcombe, R.; Fox, D. DART: Dense articulated real-time tracking with consumer depth cameras. *Autonomous Robots* Vol. 39, No. 3, 239–258, 2015.
- [35] Innmann, M.; Zollhöfer, M.; Nießner, M.; Theobalt, C.; Stamminger, M. VolumeDeform: Real-time volumetric non-rigid reconstruction. In: *Computer Vision – ECCV 2016. Lecture Notes in Computer Science*, Vol. 9912. Leibe, B.; Matas, J.; Sebe, N.; Welling, M. Eds. Springer Cham, 362–379, 2016.
- [36] Guo, K. W.; Xu, F.; Yu, T.; Liu, X. Y.; Dai, Q. H.; Liu, Y. B. Real-time geometry, albedo and motion reconstruction using a single RGBD camera. *ACM Transactions on Graphics* Vol. 36, No. 4, Article No. 32, 2017.
- [37] Gao, W.; Tedrake, R. SurfelWarp: Efficient non-volumetric single view dynamic reconstruction. In: Proceedings of Robotics: Science and Systems, 2018.
- [38] Slavcheva, M.; Baust, M.; Ilic, S. SobolevFusion: 3D reconstruction of scenes undergoing free non-rigid motion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2646–2655, 2018.
- [39] Slavcheva, M.; Baust, M.; Cremers, D.; Ilic, S. KillingFusion: Non-rigid 3D reconstruction without correspondences. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 5474–5483, 2017.
- [40] Guo, K. W.; Taylor, J.; Fanello, S.; Tagliasacchi, A.; Dou, M. S.; Davidson, P.; Kowdle, A.; Izadi, S. TwinFusion: High framerate non-rigid fusion through fast correspondence tracking. In: Proceedings of the International Conference on 3D Vision, 596–605, 2018.

- [41] Zollhöfer, M.; Stotko, P.; Görlitz, A.; Theobalt, C.; Nießner, M.; Klein, R.; Kolb, A. State of the art on 3D reconstruction with RGB-D cameras. *Computer Graphics Forum* Vol. 37, No. 2, 625–652, 2018.
- [42] Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G. et al. Human-level control through deep reinforcement learning. *Nature* Vol. 518, No. 7540, 529–533, 2015.
- [43] Charles, R. Q.; Hao, S.; Mo, K. C.; Guibas, L. J. PointNet: Deep learning on point sets for 3D classification and segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 77–85, 2017.
- [44] Qi, C. R.; Yi, L.; Su, H.; Guibas, L. J. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In: *Proceedings of the 31st Conference on Neural Information Processing Systems*, 5099–5108, 2017.
- [45] Wang, Y.; Sun, Y. B.; Liu, Z. W.; Sarma, S. E.; Bronstein, M. M.; Solomon, J. M. Dynamic graph CNN for learning on point clouds. *ACM Transactions on Graphics* Vol. 38, No. 5, Article No. 146, 2019.
- [46] Wu, W. X.; Qi, Z.; Fuxin, L. PointConv: Deep convolutional networks on 3D point clouds. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9613–9622, 2019.
- [47] Fan, H. Q.; Su, H.; Guibas, L. A point set generation network for 3D object reconstruction from a single image. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2463–2471, 2017.
- [48] Yang, Y. Q.; Feng, C.; Shen, Y. R.; Tian, D. FoldingNet: Point cloud auto-encoder via deep grid deformation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 206–215, 2018.
- [49] Groueix, T.; Fisher, M.; Kim, V. G.; Russell, B. C.; Aubry, M. A Papier-Mache approach to learning 3D surface generation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 216–224, 2018.
- [50] Yifan, W.; Wu, S. H.; Huang, H.; Cohen-Or, D.; Sorkine-Hornung, O. Patch-based progressive 3D point set upsampling. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5951–5960, 2019.
- [51] Tchapmi, L. P.; Kosaraju, V.; Rezatofighi, H.; Reid, I.; Savarese, S. TopNet: Structural point cloud decoder. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 383–392, 2019.
- [52] Arulkumaran, K.; Deisenroth, M. P.; Brundage, M.; Bharath, A. A. A brief survey of deep reinforcement learning. *arXiv preprint arXiv: 1708.05866*, 2017.
- [53] Peng, X. B.; Abbeel, P.; Levine, S.; van de Panne, M. DeepMimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics* Vol. 37, No. 4, Article No. 143, 2018.
- [54] Zhu, Y. K.; Mottaghi, R.; Kolve, E.; Lim, J. J.; Gupta, A.; Fei-Fei, L.; Farhadi, A. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In: *Proceedings of the IEEE International Conference on Robotics and Automation*, 3357–3364, 2017.
- [55] Van Hasselt, H.; Guez, A.; Silver, D. Deep reinforcement learning with double Q-learning. In: *Proceedings of the 13th AAAI Conference on Artificial Intelligence*, 2094–2100, 2016.
- [56] Wang, Z.; Schaul, T.; Hessel, M.; van Hasselt, H.; Lanctot, M.; de Freitas, N. Dueling network architectures for deep reinforcement learning. In: *Proceedings of the 33rd International Conference on Machine Learning*, 1995–2003, 2016.
- [57] Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; Wierstra, D. Continuous control with deep reinforcement learning. In: *Proceedings of the International Conference on Learning Representations*, 2016.
- [58] Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T. P.; Harley, T.; Silver, D.; Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In: *Proceedings of the 33rd International Conference on Machine Learning*, Vol. 48, 1928–1937, 2016.
- [59] Kavan, L.; Collins, S.; Žára, J.; O’Sullivan, C. Geometric skinning with approximate dual quaternion blending. *ACM Transactions on Graphics* Vol. 27, No. 4, Article No. 105, 2008.
- [60] Wang, S. L.; Fanello, S. R.; Rhemann, C.; Izadi, S.; Kohli, P. The global patch collider. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 127–135, 2016.
- [61] Collet, A.; Chuang, M.; Sweeney, P.; Gillett, D.; Evseev, D.; Calabrese, D.; Hoppe, H.; Kirk, A. G.; Sullivan, S. High-quality streamable free-viewpoint video. *ACM Transactions on Graphics* Vol. 34, No. 4, Article No. 69, 2015.
- [62] Guo, M. H.; Cai, J. X.; Liu, Z. N.; Mu, T. J.; Martin, R. R.; Hu, S. M. PCT: Point cloud transformer. *Computational Visual Media* Vol. 7, No. 2, 187–199, 2021.



Hao-Xuan Song is currently an undergraduate in the Department of Computer Science and Technology, Tsinghua University. His research interests include computer vision and computer graphics.



Tai-Jiang Mu is currently an assistant researcher at Tsinghua University, where he received his B.S. and Ph.D. degrees in computer science and technology in 2011 and 2016, respectively. His research interests include computer vision, robotics, and computer graphics.



Jiahui Huang received his B.S. degree in computer science and technology from Tsinghua University in 2018. He is currently a Ph.D. candidate in the Department of Computer Science and Technology, Tsinghua University. His research interests include computer vision, robotics, and computer graphics.



computer vision.

Yan-Pei Cao received his B.S. and Ph.D. degrees in computer science and technology from Tsinghua University in 2013 and 2018, respectively. He is currently a research engineer at Y-tech, Kuaishou Technology. His research interests include geometric modeling and processing, 3D reconstruction, and 3D

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made.

The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Other papers from this open access journal are available free of charge from <http://www.springer.com/journal/41095>. To submit a manuscript, please go to <https://www.editorialmanager.com/cvmj>.