

From Category to Scenery: An End-to-End Framework for Multi-Person Human-Object Interaction Recognition in Videos

Tanqiu Qiao[✉], Ruochen Li[✉], Frederick W. B. Li[✉], and Hubert P. H. Shum[✉]

Durham University, Durham, United Kingdom

{tanqiu.qiao, ruochen.li, frederick.li, hubert.shum}@durham.ac.uk

Abstract. Video-based Human-Object Interaction (HOI) recognition explores the intricate dynamics between humans and objects, which are essential for a comprehensive understanding of human behavior and intentions. While previous work has made significant strides, effectively integrating geometric and visual features to model dynamic relationships between humans and objects in a graph framework remains a challenge. In this work, we propose a novel end-to-end category to scenery framework, CATS, starting by generating geometric features for various categories through graphs respectively, then fusing them with corresponding visual features. Subsequently, we construct a scenery interactive graph with these enhanced geometric-visual features as nodes to learn the relationships among human and object categories. This methodological advance facilitates a deeper, more structured comprehension of interactions, bridging category-specific insights with broad scenery dynamics. Our method demonstrates state-of-the-art performance on two pivotal HOI benchmarks, including the MPHUI-72 dataset for multi-person HOIs and the single-person HOI CAD-120 dataset.

Keywords: Human-object interaction · Multi-person interaction · Feature fusion.

1 Introduction

Human-Object Interaction (HOI) recognition delves into the subtle dynamics between humans and objects, aiming to capture the breadth of their interactions from basic actions to complex activities. This field transcends mere identification to explore the depth of their interactions, from elementary actions to intricate sequences, which are essential for a comprehensive understanding of human behavior and intentions [30,35,51]. Accurate HOI recognition is crucial across various domains, serving as a cornerstone for developing sophisticated surveillance [7,37], enhancing video analysis techniques [32,27,25], and facilitating effective human-robot collaboration [39,31].

Prior work in Human-Object Interaction (HOI) detection predominantly examines interactions within static images, offering crucial insights yet constrained by the lack of temporal dynamics [13,28,12]. The emergence of single-person

HOI video datasets marks a significant advancement [18,8,19], enabling the development of models that understand spatio-temporal actions through visual cues [34,15,30]. A notable progression is presented by [35], which leverages geometric features informed networks for HOI recognition in videos, broadening the scope to encompass two-person HOIs with the introduction of a novel dataset.

While fusing geometric and visual features achieves remarkable performance, video-based HOI recognition still faces challenges in effectively fusing these features and learning dynamic relationships between humans and objects in a graph model. 2G-GCN [35] attempts to enrich visual data with geometric information via a graph-based network. However, merging geometric features of all humans and objects with individual visual features in a single graph leads to a critical flaw by neglecting category-specific characteristics. This fusion difficulty hampers accurate and specific HOI learning, especially in complex multi-person scenes.

Categorization simplifies learning and improves behavior discrimination by grouping similar features, enhancing model accuracy in identifying diverse interactions. In this work, we follow natural cognitive processes [26,3] to learn HOIs from category-level feature fusion to scenery-level graph representation, facilitating a structured and comprehensive understanding. This strategy enables a more sophisticated integration of varied feature types, ensuring each level is fully leveraged for enhanced representational efficacy. We propose a novel end-to-end CATegory to Scenery framework (CATS), which initially generates geometric features via a graph for different categories, integrating them with corresponding visual features. Subsequently, a scenery interactive graph is constructed using these enriched geometric-visual features as nodes, to deeply understand the interaction dynamics among all humans and objects.

Our approach surpasses state-of-the-art performance on two HOI benchmarks, including the two-person MPHUI-72 [35] dataset and the single-person HOI CAD-120 [18] dataset. Additionally, we conduct ablation studies to evaluate the core components of our model. Our main contributions are:

- We propose an end-to-end framework CATS ranging from category-level feature fusion to scenery-level graph for multi-person HOI recognition in videos ¹.
- We propose a multi-category multi-modality fusion module that fuses visual features and graph-based geometric features for human and object categories, respectively.
- We propose a scenery interactive graph to learn the relationships among human and object categories via an attention-based graph.

2 Related Work

2.1 HOI Recognition in Videos

There are two setups for video-based HOI recognition, where the more challenging setup focuses on segmenting and recognizing distinct human sub-activities in

¹ The code is open-sourced at <https://github.com/tanqiu98/CATS>.

videos. Deep neural networks (DNNs) and graphical models have been combined in recent works. A paradigm for integrating the effectiveness of spatio-temporal graphs with Recurrent Neural Networks (RNNs) in sequence learning is presented by Jain et al. [15]. Using learnable graph structures for videos, Qi et al. [34] expand previous graphical models in DNNs and pass messages through GPNN. For the intention of acquiring spatial relations, Dabral et al. [6] compare GCNs to Convolutional Networks and Capsule Networks. In attempting to investigate the evolution of spatio-temporal connections and identify objects in a scene [23,43], STIGPN [43] utilizes visual-based multi-modal features and a multi-stream fusion strategy to enhance the reasoning capability of the model. Morais et al. [30] present a visual feature attention model to learn asynchronous and sparse HOI in videos. Xing et al. [44] represent the 2D or 3D spatial relation of human skeletons and object center points from the detection results in video data as a graph. Based on prior visual-only and geometric-only approaches, 2G-GCN [35] incorporates geometric features to complement visual features into the HOI recognition network through a graph network. Nevertheless, the fusion of geometric and visual features introduces certain design complexities that offer opportunities for further refinement.

Another more relaxed setup in HOI recognition aims to generate `<human, predicate, object>` triplets, neglecting a more detailed analysis of specific actions and interactions. For example, in recent years, SERVO-HOI [1] presents a robust end-to-end framework adept at recognizing HOIs within in-the-wild videos, especially effective in high label-skew settings. Zeng et al. [47] introduce the Relation-Pose Transformer (RPT), a novel framework designed to intricately model the spatial and temporal dynamics between relations and poses, adept at encapsulating spatially contextualized information and the temporal evolution of relationships. Furthermore, Zhang et al. [49] explore a new task, Human-Object-Object Interaction (HOOI) detection, focusing on localizing the human and identifying their interactions within untrimmed videos as a quadruple `<human, interaction, object1, object2>`. In this work, our study concentrates on the more challenging aspect of video-based HOI recognition, specifically the segmentation and recognition of distinct human sub-activities along the video timeline.

2.2 Graph-based HOI Analysis

Graphical models facilitate the sharing of contextual information among nodes. Qi et al. [34] introduce this concept in HOI detection, where they propose a fully-connected graph with detected instances as nodes and update node features with a message passing algorithm. Wang et al. [42] suggest that adaptation to two sets of heterogeneous nodes, human and object, is essential for graph-based HOI analysis. This necessitates modelling intra-class messages differently from inter-class messages during message passing. Incorporating the heterogeneity of nodes, Gao et al. [11] create separate human-centric and object-centric graphs for HOI detection by treating human-object pairs as nodes and employing the pairwise spatial relations as node encoding. VSGNet et al. [41] leverages graph convolution

and spatial configuration to refine visual features of human-object pairs and exploits structural connections between them. SCG [48] develops a bipartite graph to model interrelationships between nodes in HOI scene where each human node is connected to each object node. Building upon SCG, Park et al. [33] design a graph with a pose-conditioned self-loop structure to update the encoding of human nodes with local features of skeleton joints. Additionally, Zhang et al. [50] construct an interaction-centric graph by treating selected interaction proposals as graph nodes to examine inter-interaction semantic structure and intra-interaction spatial structure.

Recent advancements in HOI recognition tasks are also inspired by graphical models. LIGHTEN [40] employs a graph structure to model human and object embeddings, which serves them as nodes in the scene. In a similar vein, Dabral et al. [6] investigate the efficacy of GCNs in spatial relation learning compared to Convolutional Networks and Capsule Networks. Wang et al. [43] propose the STIGPN to understand the evolution of spatio-temporal relationships and distinguish the objects involved in the background using parsed graphs. Xing et al. [44] introduce a novel spatial attention mechanism that can enhance action recognition by adaptively generating a spatial-relation graph during HOIs. Inter-Diff [45] utilizes a diffusion model [4] combined with a physics-informed predictor to anticipate 3D HOIs, effectively capturing complex, long-term interactions by modeling dynamic objects and whole-body motion in a spatial-temporal graph neural network. In 2G-GCN [35], linking collective geometric features with individual visual features causes hierarchical misalignment, as high-level spatial information may not align well with detailed, entity-specific visual data. This focuses on less relevant objects and fails to explicitly learn HOIs. In this study, we develop an understanding of HOIs by progressing from category-level feature fusion to scenery-level graph representation, enabling a structured and thorough comprehension of interactions.

3 Methodology

We propose an end-to-end framework CATS (Fig. 1) to learn HOIs from category-level to scenery-level, which first focuses on the inherent characteristics of different categories, capturing their physical properties and contextual visual cues to achieve a rich feature representation. It then adopts a graph attention neural network to learn multi-category features as a scenery graph representation, which represents the true HOI. This approach mirrors natural cognitive processes [26,3] facilitating a structured and comprehensive understanding of interactions within various contexts.

Alternative architecture performs suboptimally, an approach treats each human and object as an entity independently, ignoring the correlation between the same category and compromising the model’s ability to understand complex dynamics. An alternative method [35] groups all human poses and object bounding boxes into a single category for geometric feature learning, and then combines these geometric features with visual features in a single graph learning, which

complicates entity representation and hampers explicit HOI learning. We compare these alternative architectures with our method in Experimental Results 4.

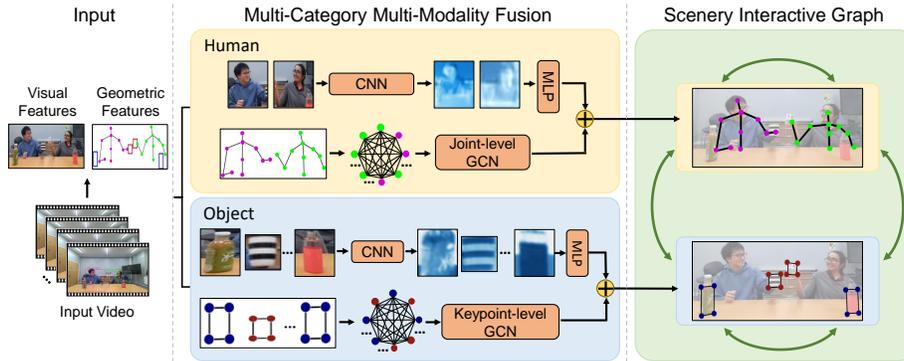


Fig. 1. Overview of our end-to-end framework CATS. We first learn geometric features via a graph for human and object categories, fusing them with corresponding visual features. Subsequently, a scenery interactive graph is constructed to deeply understand the interaction dynamics between multi-categories.

3.1 Multi-Category Multi-Modality Fusion

Previous CNN-based methods for HOI recognition in videos have predominantly focused on visual features [29,21,30], which may not be sufficient in cases of occlusion. While more advanced approaches like 2G-GCN [35] have attempted to incorporate geometric features to complement visual features, they categorize all human skeletons and object bounding boxes under a single category for geometric feature learning, thereby neglecting the distinct characteristics unique to each category and potentially generating skewed geometric features.

To this end, we propose a multi-category multi-modality fusion module that first learns geometric features via a graph for human and object two categories and then fuses them with corresponding visual features (Fig. 1). These category-specific features establish a rich multimodal context, providing a solid foundation for subsequent accurate interaction recognition.

Geometric Features For feature representation in human category and other related tasks, following previous successes [35,24], we concatenate the position and velocity of all humans into keypoint channels, forming human geometric features $\mathcal{HG} = \{hg_{t,h,j}\}_{t=1,h=1,j=1}^{T,H,J} \in \mathbb{R}^4$, where $hg_{t,h,j}$ denotes the body joint of type j in human h at time t , T denotes the total number of frames in the video, H and J denote the total number of humans and keypoints of a human body in a frame, respectively. Similar to humans, object geometric features

$\mathcal{OG} = \{og_{t,o,u}\}_{t=1, o=1, u=1}^{T,O,2} \in \mathbb{R}^4$, where $og_{t,o,u}$ denotes the bounding box diagonal points u in object o at time t and O denotes the total number of objects.

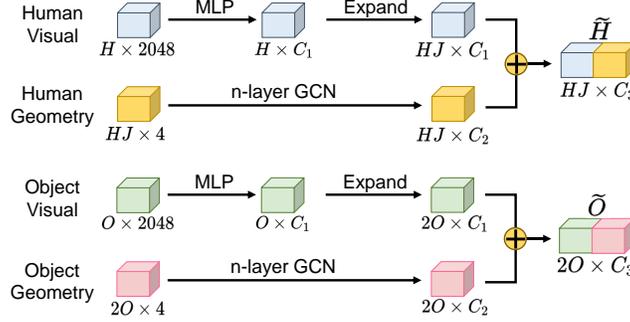


Fig. 2. The process of learning and fusing geometric and visual features for human and object categories.

As shown in Fig. 2, human and object geometric features are adopted n-layer GCNs to capture spatial dynamics and interactions in each category. This enables deeper analysis through successive transformations, allowing the graph-based network to learn intricate patterns of spatial dynamic interactions at multiple levels of abstraction [46,9]. Here, taking human geometric features as an example, the operation of each GCN layer is formalized as:

$$H^{(l+1)} = \sigma \left(AH^{(l)}W^{(l)} \right), \quad (1)$$

where $H^{(l)}$ represents the activation matrix at the l th layer ($H^{(0)} = \mathcal{HG}$ for the initial layer), A is the adjacency matrix defining the graph structure, $W^{(l)}$ is the weight matrix for the l th layer, and σ is the Tanh activation function.

For an n-layer GCN, this transformation is applied iteratively to obtain the final embedded human geometric features:

$$HG' = H^{(n)} = \sigma \left(AH^{(n-1)}W^{(n-1)} \right) \quad (2)$$

where n is the total number of GCN layers, iterating the process from $l = 0$ to $n - 1$. We choose $n = 4$ based on empirical experimental results. Through this operation, we can obtain the embedded human and object geometric features: $HG' \in \mathbb{R}^{T \times HJ \times C_2}$ and $OG' \in \mathbb{R}^{T \times 2O \times C_2}$.

Visual Features In contrast to geometric features, visual features in videos offer a wealth of contextual information and essential feature representations. Following [30,35], we derive 2048-dimensional visual features of entities from Region of Interest (ROI) pooled 2D bounding boxes around humans and objects in video frames. As shown in Fig. 2, they are subsequently reduced dimensionally

to C_1 through an MLP with learnable embeddings and aligned dimensionally with geometric features. This process results in the embedded human and object visual features: $HV' \in \mathbb{R}^{T \times HJ \times C_1}$ and $OV' \in \mathbb{R}^{T \times 2O \times C_1}$.

Multi-Modality Fusion Finally, we fuse embedded geometric and visual features in the human and object keypoint channel, producing new enriched human and object feature representations, respectively:

$$\tilde{H} = HG' \oplus HV' \in \mathbb{R}^{T \times HJ \times C_3}; \quad (3)$$

$$\tilde{O} = OG' \oplus OV' \in \mathbb{R}^{T \times 2O \times C_3}, \quad (4)$$

where \oplus represents concatenate operation and $C_3 = C_1 + C_2$. This refined fusion of geometric and visual cues creates a richly contextualized blend, laying a solid foundation for enhanced scenery graph learning of HOIs.

3.2 Scenery Interactive Graph

To effectively model the interactions between humans and objects, the existing method [30] focuses exclusively on their visual features to construct an interaction graph. This approach taps into the visual aspect of interactions, which is essential but insufficient for grasping the dynamic spatial relationships critical to understanding the complexities of HOI. Furthermore, 2G-GCN [35] offers a more comprehensive view but fuse geometric features representing all entities with visual features representing individuals, which results in hierarchical misalignment and fails to explicitly learn HOIs.

To overcome the constraints of prior approaches, we propose a scenery interactive graph that adopts a graph attention neural network to learn interactions between different categories with enriched feature representation (Fig. 1), to deeply understand the interaction dynamics among all humans and objects. This structured approach facilitates a comprehensive understanding of interactions within various contexts.

GAT for Learning Scenery Graph Specifically, we adopt Graph Attention Networks (GAT) [14] in learning scenery graph interactions is particularly advantageous due to their ability to dynamically adjust to rapid changes in human and object interactions within scenery graphs, thanks to their adaptive edge weighting and handling of non-static features. This ensures a precise focus on relevant entities and their evolving relationships, optimizing the model’s responsiveness to the complex dynamics of interactions.

We construct the HOI scenery graph $\mathcal{G}_{s-t} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} \in \mathbb{R}^{T \times (HJ+2O) \times C_3}$ represents the node features, which is obtained by concatenating the local human feature representation \tilde{H} and object feature representation \tilde{O} , and $\mathcal{E} \in \mathbb{R}^{T \times (HJ+2O) \times (HJ+2O)}$ denotes the initialized fully-connected adjacency matrix. For each node \mathcal{V}_i at time step $t \in [1, \dots, T]$, the feature representation

is:

$$\mathcal{V}_i^t = \sigma \left(\sum_{j \in \mathcal{N}_{(i) \cup i}} \alpha_{i,j}^t \Theta \mathcal{V}_j^t \right), \quad (5)$$

and the attention coefficients $\alpha_{i,j}$ are computed as:

$$\alpha_{i,j}^t = \frac{\exp(\text{LeakyReLU}(\mathbf{W}^\top [\Theta \mathcal{V}_i^t, \|\cdot\|, \Theta \mathcal{V}_j^t]))}{\sum_{n \in \mathcal{N}_{(i) \cup i}} \exp(\text{LeakyReLU}(\mathbf{W}^\top [\Theta \mathcal{V}_i^t, \|\cdot\|, \Theta \mathcal{V}_n^t]))}, \quad (6)$$

where $\Theta(\cdot)$ is the transformation function, $\mathcal{N}(\cdot)$ is the neighbor set of node i and \mathbf{W} represents learnable parameters. This dynamic weighting is crucial as it allows the model to adaptively focus on the most relevant nodes and edges, reflecting the changing nature of interactions and relationships within the scene.

RNN-based Network for Learning Temporal Dependency After obtaining the learned HOI scenery graph representations at each time step t , we employ an RNN-based network to learn the temporal dependencies across all the time steps. Specifically, we utilize a Bi-direction Gated Recurrent Unit (Bi-GRU) [5] that enables our model to integrate both past and future contexts, enhancing its understanding of the sequential dynamics in human-object interactions. The GRU’s gating mechanisms effectively manage long-term dependencies, ensuring robust temporal modeling. For the learned step-wise feature representations, we utilize a Gumbel-Softmax module [16], enabling precise and adaptable delineation of sub-event lengths in video sequences. This module is instrumental in enabling gradient-based optimization while maintaining probabilistic integrity in segmenting actions, a crucial aspect when dealing with the inherently fluctuating characteristics of video content. Subsequently, we employ another Bi-GRU to discern the temporal relations among segmented sub-actions. The processed features are then leveraged to identify specific sub-activities associated with humans, with the granularity of recognition tailored to suit the requirements of the specific dataset.

4 Experiments

4.1 Datasets

We evaluate CATS on two datasets: MPHUI-72 [35] and CAD-120 [18], showcasing the superior results on multi-person and single-person HOI recognition.

The MPHUI-72 dataset is valuable for two-person HOI tasks. It contains 72 videos of 8 pairs of people performing 3 distinct activities (*Cheering*, *Hair cutting* and *Co-working*) with 13 human sub-activities (e.g., *Sit*, *Pour*). Each video showcases two participants interacting with 2-4 objects from 3 unique angles. Geometric features and human sub-activities labels are frame-wise annotated.

CAD-120 is a prominent dataset for single-person HOI recognition. It contains 120 RGB-D videos, capturing 10 distinct activities executed by 4 participants,

each repeated three times. In each video, a participant interacts with 1-5 objects. The dataset provides frame-wise annotations for 10 human sub-activities (e.g., *opening, placing*).

4.2 Evaluation Protocol

Following the evaluation protocol of [30,35], we assess CATS across two specific tasks: joint segmentation and label recognition for pre-segmented entities. The initial task involves both segmenting and classifying the timeline of each entity in a video, while the second extends this by assigning labels to pre-segmented sections with known ground truth. We adopt the $F1@k$ metric [22] for evaluation, using standard thresholds of $k = 10\%$, 25% , and 50% . This metric, prevalent in segmentation research [22,10,30], determines the correctness of a predicted action segment based on its minimum Intersection over Union (IoU) overlap with the ground truth and is particularly effective for assessing brief actions and detailed segmentation. For dataset evaluation, we implement a leave-two-subjects-out strategy for the MPHUI-72 dataset and a leave-one-subject-out cross-validation approach for CAD-120.

4.3 Network Setting

The visual features of humans and objects are extracted from 2D bounding boxes within the video using a Faster R-CNN module [36] that has been pre-trained [2] on the Visual Genome dataset [20]. For multi-modality fusion, we set $C_1 = 512$ and $C_2 = 256$, resulting in a fused dimension of $C_3 = 768$, which supports varied feature dimensions as shown in Fig. 2.

4.4 Quantitative Comparison

Multi-person HOIs In the MPHUI-72 dataset, results in Table 1 demonstrate CATS not only surpasses the previous state-of-the-art models, ASSIGN [30] and 2G-GCN [35], showcasing significant performance improvements, but also exhibits unparalleled stability. This is highlighted by CATS’s superior performance across all F1 configurations coupled with substantially lower standard deviations. Specifically, in the $F1@10$ score, CATS achieves 71.3%, which is approximately 3% and 12% higher than 2G-GCN and ASSIGN, respectively, marking a clear advancement in both predictive accuracy and consistency in the domain of human-object interaction recognition. These experimental outcomes further underscore the significance of geometric features in the multi-person Human-Object Interaction (MPHUI) domain. Models based solely on visual features, such as ASSIGN, are noticeably outperformed by those that incorporate both visual and geometric information. Although 2G-GCN integrates both visual and geometric features, its sub-optimal performance can be attributed to a lack of specificity in representing individual entities. Consequently, our model’s superior performance and stability are not just a result of integrating multiple types of features but also our model’s ability to specifically and effectively capture the nuanced dynamics of each entity involved in the interaction.

Table 1. Joined segmentation and label recognition on MPHUI-72.

Model	Sub-activity		
	F ₁ @10	F ₁ @25	F ₁ @50
ASSIGN [30]	59.1 ± 12.1	51.0 ± 16.7	33.2 ± 14.0
2G-GCN [35]	68.6 ± 10.4	60.8 ± 10.3	45.2 ± 6.5
CATS	71.3 ± 5.0	65.8 ± 3.9	48.8 ± 5.3

Single-person HOIs In the CAD-120 dataset, as presented in Table 2, CATS demonstrates strong competitiveness in the single-person HOI scenarios. For both human sub-activity and object affordance labelling tasks, CATS surpasses various prior methods, including those reliant on visual features like ATCRF[17] and [30], as well as the more sophisticated visual-geometric approach offered by 2G-GCN [35]. Notably, CATS secures SOTA performance in both F1@10 and F1@25 metrics, registering improvements of 1.6% and 0.1% over ASSIGN and 2G-GCN, respectively. This achievement underscores CATS’s exceptional capability to accurately model and predict the dynamics of interactions, highlighting its adaptability and efficiency across different HOI challenges.

Table 2. Joined segmentation and label recognition on CAD-120.

Model	Sub-activity		
	F ₁ @10	F ₁ @25	F ₁ @50
rCRF [38]	65.6 ± 3.2	61.5 ± 4.1	47.1 ± 4.3
Independent BiRNN	70.2 ± 5.5	64.1 ± 5.3	48.9 ± 6.8
ATCRF [17]	72.0 ± 2.8	68.9 ± 3.6	53.5 ± 4.3
Relational BiRNN	79.2 ± 2.5	75.2 ± 3.5	62.5 ± 5.5
ASSIGN [30]	88.0 ± 1.8	84.8 ± 3.0	73.8 ± 5.8
2G-GCN [35]	89.5 ± 1.6	87.1 ± 1.8	76.2 ± 2.8
CATS	89.6 ± 2.1	87.3 ± 1.5	76.0 ± 3.5

4.5 Qualitative Comparison

In this section, we present a qualitative comparison of CATS with the state-of-the-art method across the MPHUI-72 and CAD-120 datasets.

Fig. 3 and Fig. 4 illustrate *Cheering* and *Hair Cutting* activities within the MPHUI-72 dataset, comparing the segmentation and labeling tasks performed by CATS and 2G-GCN [35] against the ground truth. Significant segmentation errors are marked with red dashed boxes. Although both methods exhibit some discrepancies in their predictions, CATS more closely aligns with the ground truth, offering a more precise and stable visualization across a variety of actions. Conversely, 2G-GCN is prone to generating inappropriate sub-activities such as *cheers* and *lift* in the *Cheering* activity. Moreover, in the *Hair Cutting* activity, 2G-GCN inaccurately presents the *cut* sub-activity into *place* sub-activity, further deviating from the expected interaction dynamics. This comparison underscores

the superior accuracy and reliability of CATS in capturing and visualizing complex human-object interactions within diverse scenarios.

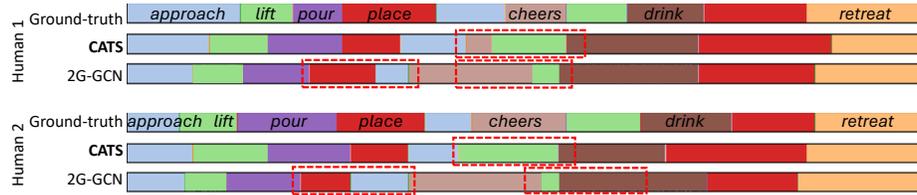


Fig. 3. Visualization of segmentation on MPHUI-72 for *Cheering* activity. Red dashed boxes highlight major segmentation errors.

Fig. 5 and Fig. 6 illustrate the *Cleaning Objects* and *Making Cereal* activities from the single-person CAD-120 dataset, with abnormal segmentation instances accentuated by red dashed boxes. For the *Cleaning Objects* activity, both methods effectively match the overall ground truth. However, CATS provides a visualization that more closely approximates the ground truth. In the *Making Cereal* activity, CATS significantly outperforms 2G-GCN, particularly in sub-activities such as *pouring*, *moving*, and *reaching*, while 2G-GCN yields some inaccurate segmentations. The enhanced precision of CATS in capturing the intricacies of each activity highlights its superior performance, excelling in the identification and precise representation of detailed actions and interactions within the scenes, thus delivering a more accurate and reliable analysis of the activities performed.

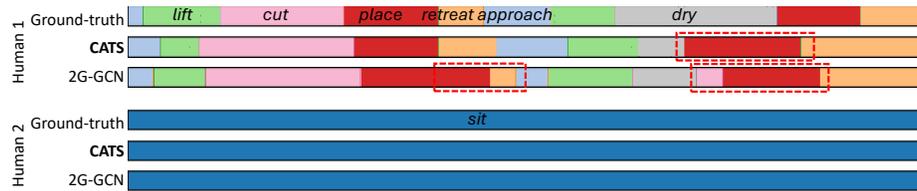


Fig. 4. Visualization of segmentation on MPHUI-72 for *Hair cutting* activity. Red dashed boxes highlight major segmentation errors.

4.6 Alternative Architectures and Ablation Studies

Architecture Alternatives Comparison We evaluate the HOI recognition performance on the MPHUI-72 and CAD-120 datasets by conducting tests on various alternative model structures. The experimental outcomes, as detailed in Tables 3 and 4, reveal that our model consistently delivers superior results compared to these alternatives. This superior performance is likely attributable



Fig. 5. Visualization of segmentation on CAD-120 for *Cleaning objects* activity. Red dashed boxes highlight major segmentation errors.

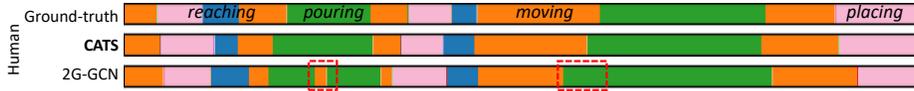


Fig. 6. Visualization of segmentation on CAD-120 for *Making Cereal* activity. Red dashed boxes highlight major segmentation errors.

to the unique consideration our model gives to category-level interactions, specifically the distinct analysis of human-human and object-object interactions. Unlike other approaches that might treat interactions generically or overlook the nuanced distinctions between different types of interactions, our model maintains a comprehensive view.

Table 3. Comparison between architecture alternatives and CATS on MPH01-72.

Model	Sub-activity		
	F ₁ @10	F ₁ @25	F ₁ @50
Independent-entity architecture	65.1 ± 3.3	58.7 ± 1.7	40.4 ± 3.9
2G-GCN [35]	68.6 ± 10.4	60.8 ± 10.3	45.2 ± 6.5
CATS	71.3 ± 5.0	65.8 ± 3.9	48.8 ± 5.3

Table 4. Comparison between architecture alternatives and CATS on CAD-120.

Model	Sub-activity		
	F ₁ @10	F ₁ @25	F ₁ @50
Independent-entity architecture	85.9 ± 4.0	84.1 ± 4.9	72.8 ± 5.2
2G-GCN [35]	89.5 ± 1.6	87.1 ± 1.8	76.2 ± 2.8
CATS	89.6 ± 2.1	87.3 ± 1.5	76.0 ± 3.5

GCN Layers for Geometric Feature Learning In this section, we conduct ablation studies to elucidate the impact of the depth of GCN layers on the geometric learning of human joints and object keypoints within our network, results are shown in Table 5. To assess the influence of GCN layer depth on model performance, we explore configurations with 1, 2, 3, 4, and 5 GCN layers. Through this comparative analysis, we aim to identify the most effective layer depth that balances computational efficiency with the nuanced understanding of spatial relationships essential for interpreting complex interactions between humans and objects. The results indicate that a configuration of 4-layer GCN offers the optimal balance, providing the best performance in terms of both accuracy and computational efficiency. This depth allows for sufficient complexity to understand and model the geometric relationships critical for accurate interaction recognition, without incurring the diminishing returns or increased computational demand associated with additional layers.

Table 5. Results of different GCN layers in multi-category multi-modality fusion on MPHOI-72.

Model	Sub-activity		
	F ₁ @10	F ₁ @25	F ₁ @50
1-layer GCN	70.4 ± 1.7	62.0 ± 2.5	43.9 ± 3.8
2-layer GCN	68.8 ± 4.3	62.1 ± 4.3	44.0 ± 3.3
3-layer GCN	67.4 ± 4.2	63.3 ± 3.4	44.2 ± 1.3
5-layer GCN	70.4 ± 5.7	60.0 ± 2.3	43.7 ± 2.2
4-layer GCN (Ours)	71.3 ± 5.0	65.8 ± 3.9	48.8 ± 5.3

5 Conclusion

In conclusion, we propose CATS, an advanced end-to-end framework that enhances video-based HOI recognition through sophisticated integration of category and scenery level analyses. It first fuses multi-modal features of different categories, and then construct a scenery interactive graph to learn the relationships between these categories. CATS demonstrates superior performance on key benchmarks such as MPHOI-72 and CAD-120 datasets, showcasing the effectiveness of multi-person and single-person HOI recognition.

Acknowledgement

This research is supported in part by the EPSRC NorthHFutures project (ref: EP/X031012/1).

References

1. Agarwal, A., Dabral, R., Jain, A., Ramakrishnan, G.: Skew-robust human-object interactions in videos. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 5098–5107 (2023)
2. Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., Zhang, L.: Bottom-up and top-down attention for image captioning and visual question answering. In: CVPR. pp. 6077–6086 (2018)
3. Baldassano, C., Beck, D.M., Fei-Fei, L.: Human–object interactions are more than the sum of their parts. *Cerebral Cortex* **27**(3), 2276–2288 (2017)
4. Chang, Z., Koulieris, G.A., Shum, H.P.H.: On the design fundamentals of diffusion models: A survey. *arXiv* (2023). <https://doi.org/10.48550/arXiv.2306.04542>
5. Chung, J., Gulcehre, C., Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555* (2014)
6. Dabral, R., Sarkar, S., Reddy, S.P., Ramakrishnan, G.: Exploration of spatial and temporal modeling alternatives for hoi. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2281–2290 (2021)
7. Dogariu, M., Stefan, L.D., Constantin, M.G., Ionescu, B.: Human-object interaction: Application to abandoned luggage detection in video surveillance scenarios. In: 2020 13th International Conference on Communications (COMM). pp. 157–160. IEEE (2020)

8. Dreher, C.R., Wächter, M., Asfour, T.: Learning object-action relations from bimanual human demonstration using graph networks. *IEEE Robotics and Automation Letters* **5**(1), 187–194 (2020)
9. Du, S.S., Hou, K., Salakhutdinov, R.R., Póczos, B., Wang, R., Xu, K.: Graph neural tangent kernel: Fusing graph neural networks with graph kernels. *NeurIPS* **32** (2019)
10. Farha, Y.A., Gall, J.: Ms-tcn: Multi-stage temporal convolutional network for action segmentation. In: *CVPR*. pp. 3575–3584 (2019)
11. Gao, C., Xu, J., Zou, Y., Huang, J.B.: Drg: Dual relation graph for human-object interaction detection. In: *ECCV*. pp. 696–712 (2020)
12. Gao, C., Zou, Y., Huang, J.B.: ican: Instance-centric attention network for human-object interaction detection. *arXiv preprint arXiv:1808.10437* (2018)
13. Gkioxari, G., Girshick, R., Malik, J.: Actions and attributes from wholes and parts. In: *ICCV*. pp. 2470–2478 (2015)
14. Huang, Y., Bi, H., Li, Z., Mao, T., Wang, Z.: Stgat: Modeling spatial-temporal interactions for human trajectory prediction. In: *ICCV*. pp. 6272–6281 (2019)
15. Jain, A., Zamir, A.R., Savarese, S., Saxena, A.: Structural-rnn: Deep learning on spatio-temporal graphs. In: *CVPR*. pp. 5308–5317 (2016)
16. Jang, E., Gu, S., Poole, B.: Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144* (2016)
17. Koppula, H.S., Saxena, A.: Anticipating human activities using object affordances for reactive robotic response. *IEEE TPAMI* **38**(1), 14–29 (2016)
18. Koppula, H.S., Gupta, R., Saxena, A.: Learning human activities and object affordances from rgb-d videos. *The International Journal of Robotics Research* **32**(8), 951–970 (2013)
19. Krebs, F., Meixner, A., Patzer, I., Asfour, T.: The kit bimanual manipulation dataset. In: *IEEE/RAS International Conference on Humanoid Robots (Humanoids)*. pp. 0–0 (2021)
20. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV* **123**(1), 32–73 (2017)
21. Le, H., Sahoo, D., Chen, N.F., Hoi, S.C.: Bist: Bi-directional spatio-temporal reasoning for video-grounded dialogues. *arXiv preprint arXiv:2010.10095* (2020)
22. Lea, C., Flynn, M.D., Vidal, R., Reiter, A., Hager, G.D.: Temporal convolutional networks for action segmentation and detection. In: *CVPR*. pp. 156–165 (2017)
23. Li, L., Shum, H.P.H., Breckon, T.P.: Less is More: Reducing Task and Model Complexity for 3D Point Cloud Semantic Segmentation. In: *CVPR* (2023)
24. Li, L., Shum, H.P.H., Breckon, T.P.: RAPiD-Seg: Range-Aware Pointwise Distance Distribution Networks for 3D LiDAR Segmentation. In: *ECCV* (2024)
25. Li, R., Katsigiannis, S., Shum, H.P.: Multiclass-sgen: Sparse graph-based trajectory prediction with agent class embedding. In: *ICIP*. pp. 2346–2350. *IEEE* (2022)
26. Li, Y.L., Liu, X., Wu, X., Li, Y., Lu, C.: Hoi analysis: Integrating and decomposing human-object interaction. *NeurIPS* **33**, 5011–5022 (2020)
27. Liu, M., Tang, S., Li, Y., Rehg, J.M.: Forecasting human-object interaction: joint prediction of motor attention and actions in first person video. In: *ECCV*. pp. 704–721. *Springer* (2020)
28. Mallya, A., Lazebnik, S.: Learning models for actions and person-object interactions with transfer to question answering. In: *ECCV*. pp. 414–428 (2016)
29. Maraghi, V.O., Faez, K.: Zero-shot learning on human-object interaction recognition in video. In: *2019 5th Iranian conference on signal processing and intelligent systems (ICSPIS)*. pp. 1–7 (2019)

30. Morais, R., Le, V., Venkatesh, S., Tran, T.: Learning asynchronous and sparse human-object interaction in videos. In: CVPR. pp. 16041–16050 (2021)
31. Mukherjee, D., Gupta, K., Chang, L.H., Najjaran, H.: A survey of robot learning strategies for human-robot collaboration in industrial settings. *Robotics and Computer-Integrated Manufacturing* **73**, 102231 (2022)
32. Nagarajan, T., Feichtenhofer, C., Grauman, K.: Grounded human-object interaction hotspots from video. In: ICCV. pp. 8688–8697 (2019)
33. Park, J., Park, J.W., Lee, J.S.: Viplo: Vision transformer based pose-conditioned self-loop graph for human-object interaction detection. In: CVPR. pp. 17152–17162 (2023)
34. Qi, S., Wang, W., Jia, B., Shen, J., Zhu, S.C.: Learning human-object interactions by graph parsing neural networks. In: ECCV. pp. 401–417 (2018)
35. Qiao, T., Men, Q., Li, F.W.B., Kubotani, Y., Morishima, S., Shum, H.P.H.: Geometric features informed multi-person human-object interaction recognition in videos. In: ECCV (2022)
36. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: towards real-time object detection with region proposal networks. *IEEE TPAMI* **39**(6), 1137–1149 (2016)
37. Rezaee, K., Rezakhani, S.M., Khosravi, M.R., Moghimi, M.K.: A survey on deep learning-based real-time crowd anomaly detection for secure distributed video surveillance. *Personal and Ubiquitous Computing* **28**(1), 135–151 (2024)
38. Sener, O., Saxena, A.: rcrf: Recursive belief estimation over crfs in rgb-d activity videos. In: *Robotics: Science and systems*. Citeseer (2015)
39. Smith, B.A., Yin, Q., Feiner, S.K., Nayar, S.K.: Gaze locking: passive eye contact detection for human-object interaction. In: *Proceedings of the 26th annual ACM symposium on User interface software and technology*. pp. 271–280 (2013)
40. Sunkesula, S.P.R., Dabral, R., Ramakrishnan, G.: Lighten: Learning interactions with graph and hierarchical temporal networks for hoi in videos. In: *ACM MM*. pp. 691–699 (2020)
41. Ulutan, O., Iftekhhar, A.S.M., Manjunath, B.S.: Vsgnet: Spatial attention network for detecting human object interactions using graph convolutions. In: CVPR. pp. 13617–13626 (2020)
42. Wang, H., Zheng, W.s., Yingbiao, L.: Contextual heterogeneous graph network for human-object interaction detection. In: ECCV. pp. 248–264 (2020)
43. Wang, N., Zhu, G., Zhang, L., Shen, P., Li, H., Hua, C.: Spatio-temporal interaction graph parsing networks for human-object interaction recognition. In: *ACM MM*. pp. 4985–4993 (2021)
44. Xing, H., Burschka, D.: Understanding spatio-temporal relations in human-object interaction using pyramid graph convolutional network. In: *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 5195–5201 (2022)
45. Xu, S., Li, Z., Wang, Y.X., Gui, L.Y.: Interdiff: Generating 3d human-object interactions with physics-informed diffusion. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 14928–14940 (2023)
46. You, Y., Chen, T., Wang, Z., Shen, Y.: L2-gcn: Layer-wise and learned efficient training of graph convolutional networks. In: CVPR. pp. 2127–2135 (2020)
47. Zeng, Z., Dai, P., Zhang, X., Zhang, L., Cao, X.: Cognition guided human-object relationship detection. *IEEE Transactions on Image Processing* (2023)
48. Zhang, F.Z., Campbell, D., Gould, S.: Spatially conditioned graphs for detecting human-object interactions. In: ICCV. pp. 13319–13327 (2021)

49. Zhang, M., Wu, X., Yuan, Z., He, Q., Huang, X.: Human-object-object interaction: Towards human-centric complex interaction detection. In: ACM MM. pp. 2233–2242 (2023)
50. Zhang, Y., Pan, Y., Yao, T., Huang, R., Mei, T., Chen, C.W.: Exploring structure-aware transformer over interaction proposals for human-object interaction detection. In: CVPR. pp. 19548–19557 (2022)
51. Zhuo, T., Cheng, Z., Zhang, P., Wong, Y., Kankanhalli, M.: Explainable video action reasoning via prior knowledge and state transitions. In: ACM MM. pp. 521–529 (2019)