

Scene-Centric Joint Parsing of Cross-View Videos

Hang Qi,^{1*} Yuanlu Xu,¹ Tao Yuan,¹ Tianfu Wu,² Song-Chun Zhu¹

¹Dept. Computer Science and Statistics, University of California, Los Angeles (UCLA)

²Dept. Electrical and Computer Engineering, NC State University

{hangqi, yuanluxu}@cs.ucla.edu, taoyuan@ucla.edu, tianfu_wu@ncsu.edu, sczhu@stat.ucla.edu

Abstract

Cross-view video understanding is an important yet under-explored area in computer vision. In this paper, we introduce a joint parsing framework that integrates view-centric proposals into scene-centric parse graphs that represent a coherent scene-centric understanding of cross-view scenes. Our key observations are that overlapping fields of views embed rich appearance and geometry correlations and that knowledge fragments corresponding to individual vision tasks are governed by consistency constraints available in commonsense knowledge. The proposed joint parsing framework represents such correlations and constraints explicitly and generates semantic scene-centric parse graphs. Quantitative experiments show that scene-centric predictions in the parse graph outperform view-centric predictions.

1 Introduction

During the past decades, remarkable progress has been made in many vision tasks, e.g., image classification, object detection, pose estimation. Recently, more comprehensive visual tasks probe deeper understanding of visual scenes under interactive and multi-modality settings, such as visual Turing tests (Geman et al. 2015; Qi et al. 2015) and visual question answering (Antol et al. 2015). In addition to discriminative tasks focusing on binary or categorical predictions, emerging research involves representing fine-grained relationships in visual scenes (Krishna et al. 2017; Aditya and Fermuller 2016) and unfolding semantic structures in contexts including caption or description generation (Yao et al. 2010), and question answering (Tu et al. 2014; Zhu et al. 2016).

In this paper, we present a framework for uncovering the semantic structure of scenes in a cross-view camera network. The central requirement is to resolve ambiguity and establish cross-reference among information from multiple cameras. Unlike images and videos shot from single static point of view, cross-view settings embed rich physical and geometry constraints due to the overlap between fields of

*Hang Qi, Yuanlu Xu and Tao Yuan contributed equally to this paper. This work is supported by ONR MURI project N00014-16-1-2007, DARPA XAI Award N66001-17-2-4029, and NSF IIS 1423305.

Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

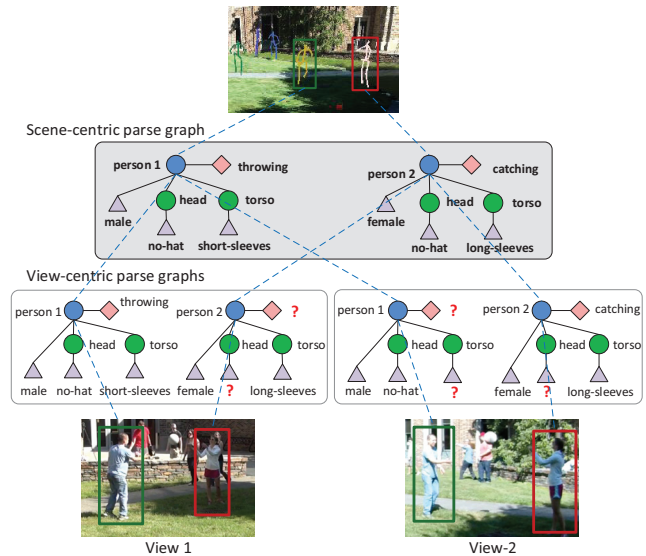


Figure 1: An example of the spatio-temporal semantic parse graph hierarchy in a visual scene captured by two cameras.

views. While multi-camera setups are common in real-world surveillance systems, large-scale cross-view activity dataset are not available due to privacy and security reasons. This makes data-demanding deep learning approaches infeasible.

Our joint parsing framework computes a hierarchy of spatio-temporal parse graphs by establishing cross-reference of entities among different views and inferring their semantic attributes from a scene-centric perspective. For example, Fig. 1 shows a parse graph hierarchy that describes a scene where two people are playing a ball. In the first view, person 2's action is not grounded because of the cluttered background, while it is detected in the second view. Each view-centric parse graph contains local recognition decisions in an individual view, and the scene centric parse graph summarizes a comprehensive understanding of the scene with coherent knowledge.

The structure of each individual parse graph fragment is induced by an ontology graph that regulates the domain of interests. A parse graph hierarchy is used to represent the correspondence of entities between the multiple views and

the scene. We use a probabilistic model to incorporate various constraints on the parse graph hierarchy and formulate the joint parsing as a MAP inference problem. A MCMC sampling algorithm and a dynamic programming algorithm are used to explore the joint space of scene-centric and view-centric interpretations and to optimize for the optimal solutions. Quantitative experiments show that scene-centric parse graphs outperforms the initial view-centric proposals.

Contributions. The contributions of this work are three-fold: (i) a unified hierarchical parse graph representation for cross-view person, action, and attributes recognition; (ii) a stochastic inference algorithm that explores the joint space of scene-centric and view-centric interpretations efficiently starting with initial proposals; (iii) a joint parse graph hierarchy that is an interpretable representation for scene and events.

2 Related Work

Our work is closely related to three research areas in computer vision and artificial intelligence.

Multi-view video analytics. Typical multi-view visual analytics tasks include object detection (Liebelt and Schmid 2010; Utasi and Benedek 2011), cross-view tracking (Berclaz et al. 2011; Leal-Taixe, Pons-Moll, and Rosenhahn 2012; Xu et al. 2016; 2017), action recognition (Wang et al. 2014), person re-identification (Xu et al. 2013; 2014) and 3D reconstruction (Hofmann, Wolf, and Rigoll 2013). While heuristics such as appearances and motion consistency constraints have been used to regularize the solution space, these methods focus on a specific multi-view vision task whereas we aim to propose a general framework to jointly resolve a wide variety of tasks.

Semantic representations. Semantic and expressive representations have been developed for various vision tasks, e.g., image parsing (Han and Zhu 2009), 3D scene reconstruction (Liu, Zhao, and Zhu 2014; Pero et al. 2013), human-object interaction (Koppula and Saxena 2016), pose and attribute estimation (Wei et al. 2016). In this paper, our representation also falls into this category. The difference is that our model is defined upon cross-view spatio-temporal domain and is able to incorporate a variety of tasks.

Interpretability. Automated generation of explanations regarding predictions has a long and rich history in artificial intelligence. Explanation systems have been developed for a wide range of applications, including simulator actions (Van Lent, Fisher, and Mancuso 2004; Lane et al. 2005; Core et al. 2006), robot movements (Lomas et al. 2012), and object recognition in images (Biran and McKeown 2014; Hendricks et al. 2016). Most of these approaches are rule-based and suffer from generalization across different domains. Recent methods including (Ribeiro, Singh, and Guestrin 2016) use proxy models or data to interpret black box models, while our scene-centric parse graphs are explicit representations of the knowledge by definition.

3 Representation

A scene-centric spatio-temporal parse graph represents humans, their actions and attributes, interaction with other ob-

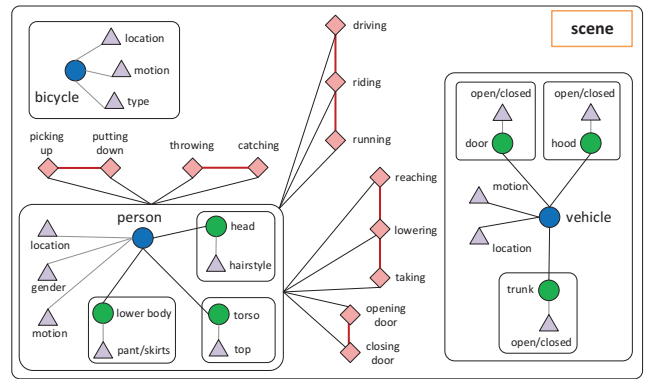


Figure 2: An illustration of the proposed ontology graph describing objects, parts, actions and attributes.

jects captured by a network of cameras. We will first introduce the concept of ontology graph as domain definitions, then we will describe parse graphs and parse graph hierarchy as view-centric and scene-centric representations respectively.

Ontology graph. To define the scope of our representation on scenes and events, an ontology is used to describe a set of plausible objects, actions and attributes. We define an ontology as a graph that contains nodes representing objects, parts, actions, attributes respectively and edges representing the relationships between nodes. Specifically, every object and part node is a concrete type of object that can be detected in videos. Edges between object and part nodes encodes “part-of” relationships. Action and attribute nodes connected to an object or part node represent plausible actions and appearance attributes the object can take. For example, Fig. 2 shows an ontology graph that describes a domain including people, vehicles, bicycles. An object can be decomposed into parts (i.e., green nodes), and enriched with actions (i.e., pink nodes) and attributes (i.e., purple diamonds). The red edges among action nodes denote their incompatibility. The ontology graph can be considered a compact AOG (Liu, Zhao, and Zhu 2014; Wei et al. 2016) without the compositional relationships and event hierarchy. In this paper, we focus on a restricted domain inspired by (Qi et al. 2015), while larger ontology graphs can be easily derived from large-scale visual relationship datasets such as (Krishna et al. 2017) and open-domain knowledge bases such as (Liu and Singh 2004).

Parse graphs. While an ontology describes plausible elements, only a subset of these concepts can be true for a given instance at a given time. For example, a person cannot be both “standing” and “sitting” at the same time, while both are plausible actions that a person can take. To distinguish plausible facts and satisfied facts, we say a node is *grounded* when it is associated with data. Therefore, a subgraph of the ontology graph that only contains grounded nodes can be used to represent a specific *instance* (e.g. a specific person) at a specific time. In this paper, we refer to such subgraphs as *parse graphs*.

Parse graph hierarchy. In cross-view setups, since each

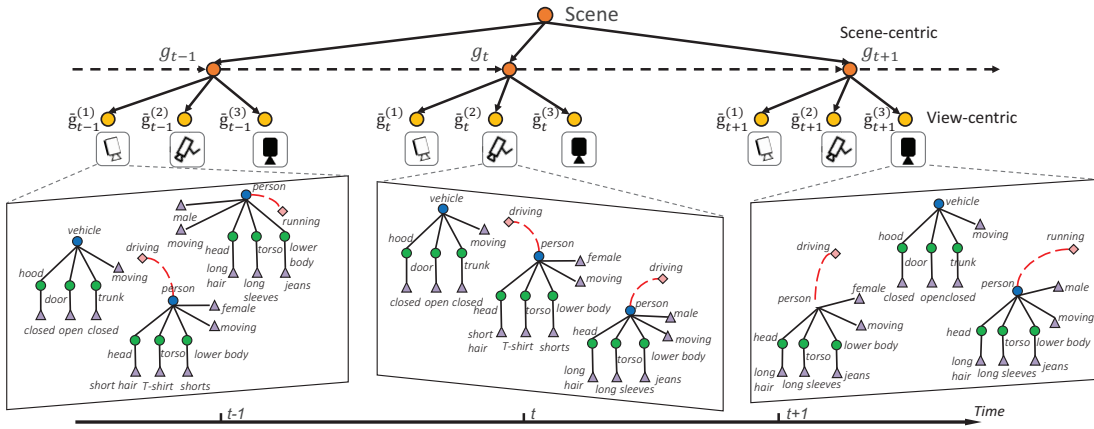


Figure 3: The proposed spatio-temporal parse graph hierarchy. (Better viewed electronically and zoomed).

view only captures an incomplete set of facts in a scene, we use a spatio-temporal hierarchy of parse graphs to represent the collective knowledge of the scene and all the individual views. To be concrete, a view-centric parse graph \tilde{g} contains nodes grounded to a video sequence captured by an individual camera, whereas a scene-centric parse graph g is an aggregation of view-centric parse graphs and therefore reflects a global understanding of the scene. As illustrated in Fig. 3, for each time step t , the scene-centric parse graph g_t is connected with the corresponding view-centric parse graphs $\tilde{g}_t^{(i)}$ indexed by the views, and the scene-centric graphs are regarded as a Markov chain in the temporal sequence. In terms of notations, in this paper we use a tilde notation to represent the view-centric concepts \tilde{x} corresponding to scene-centric concepts x .

4 Probabilistic Formulation

The task of joint parsing is to infer the spatio-temporal parse graph hierarchy $G = \langle \Phi, g, \tilde{g}^{(1)}, \tilde{g}^{(2)}, \dots, \tilde{g}^{(M)} \rangle$ from the input frames from video sequences $I = \{I_t^{(i)}\}$ captured by a network of M cameras, where Φ is an object identity mapping between scene-centric parse graph g and view-centric parse graphs $\tilde{g}^{(i)}$ from camera i . Φ defines the structure of parse graph hierarchy. In this section, we discuss the formulation assuming a fixed structure, while defer the discussion of how to traverse the solution space to section 5.

We formulate the inference of parse graph hierarchy as a MAP inference problem in a posterior distribution $p(G|I)$ as follows

$$G^* = \arg \max_G p(I|G) \cdot p(G). \quad (1)$$

Likelihood. The likelihood term models the grounding of nodes in view-centric parse graphs to the input video sequences. Specifically,

$$\begin{aligned} p(I|G) &= \prod_{i=1}^M \prod_{t=1}^T p(I_t^{(i)} | \tilde{g}_t^{(i)}) \\ &= \prod_{i=1}^M \prod_{t=1}^T \prod_{v \in V(\tilde{g}_t^{(i)})} p(I(v)|v), \end{aligned} \quad (2)$$

where $\tilde{g}_t^{(i)}$ is the view-centric parse graph of camera i at time t and $V(\tilde{g}_t^{(i)})$ is the set of nodes in the parse graph. $p(I(v)|v)$ is the node likelihood for the concept represented by node v being grounded on the data fragment $I(v)$. In practice, this probability can be approximated by normalized detection and classifications scores (Pirsiavash, Ramanan, and Fowlkes 2011).

Prior. The prior term models the compatibility of scene-centric and view-centric parse graphs across time. We factorize the prior as

$$p(G) = p(g_1) \prod_{t=1}^{T-1} p(g_{t+1}|g_t) \prod_{i=1}^M \prod_{t=1}^T p(\tilde{g}_t^{(i)}|g_t), \quad (3)$$

where $p(g_1)$ is a prior distribution on parse graphs that regulates the combination of nodes, and $p(g_t|g_{t-1})$ is a transitions probability of scene-centric parse graphs across time. Both probability distributions are estimated from training sequences. $p(\tilde{g}_t^{(i)}|g_t)$ is defined as a Gibbs distribution that models the compatibility of scene-centric and view-centric parse graphs in the hierarchy (we drop subscripts t and camera index i for brevity).

$$\begin{aligned} p(\tilde{g}|g) &= \frac{1}{Z} \exp\{-\mathcal{E}(g, \tilde{g})\} \\ &= \frac{1}{Z} \exp\{-w_1 \mathcal{E}_S(g, \tilde{g}) - w_2 \mathcal{E}_A(g, \tilde{g}) \\ &\quad - w_3 \mathcal{E}_{Act}(g, \tilde{g}) - w_4 \mathcal{E}_{Attr}(g, \tilde{g})\}, \end{aligned} \quad (4)$$

where energy $\mathcal{E}(g, \tilde{g})$ is decomposed into four different terms described in detail in the subsection below. The weights are tuning parameters that can be learned via cross-validation. We consider view-centric parse graphs for videos from different cameras are independent conditioned on scene-centric parse graph under the assumption that all cameras have fixed and known locations.

4.1 Cross-view Compatibility

In this subsection, we describe the energy function $\mathcal{E}(g, \tilde{g})$ for regulating the compatibility between the occurrence of

objects in the scene and an individual view from various aspects. Note that we use a tilde notation to represent the node correspondence in scene-centric and view-centric parse graphs (i.e., for a node $v \in g$ in a scene-centric parse graph, we refer to the corresponding node in a view-centric parse graph as \tilde{v}).

Appearance similarity. For each object node in the parse graph, we keep an appearance descriptor. The appearance energy regulates the appearance similarity of object o in the scene-centric parse graph and \tilde{o} in the view-centric parse graphs.

$$\mathcal{E}_A(g, \tilde{g}) = \sum_{o \in g} \|(\phi(o) - \phi(\tilde{o}))\|_2, \quad (5)$$

where $\phi(\cdot)$ is the appearance feature vector of the object. At the view-level, this feature vector can be extracted by pre-trained convolutional neural networks; at the scene level, we use a mean pooling of view-centric features.

Spatial consistency. At each time point, every object in a scene has a fixed physical location in the world coordinate system while appears on the image plane of each camera according to the camera projection. For each object node in the parse graph hierarchy, we keep a scene-centric location $s(o)$ for each object o in scene-centric parse graphs and a view-centric location $s(\tilde{o})$ on the image plane in view-centric parse graphs. The following energy is defined to enforce the spatial consistency:

$$\mathcal{E}_S(g, \tilde{g}) = \sum_{o \in g} \|s(o) - h(s(\tilde{o}))\|_2, \quad (6)$$

where $h(\cdot)$ is a perspective transform that maps a person’s view-centric foot point coordinates to the world coordinates on the ground plane of the scene with the camera homography, which can be obtained via the intrinsic and extrinsic camera parameters.

Action compatibility. Among action and object part nodes, scene-centric human action predictions shall agree with the human pose observed in individual views from different viewing angles:

$$\mathcal{E}_{Act}(g, \tilde{g}) = \sum_{l \in g} -\log p(l|\tilde{p}), \quad (7)$$

where l is an action node in scene-centric parse graphs and \tilde{p} are positions of all human parts in the view-centric parse graph. In practice, we separately train a action classifier that predicts action classes with joint positions of human parts and uses the classification score to approximate this probability.

Attribute consistency. In cross-view sequences, entities observed from multiple cameras shall have a consistent set of attributes. This energy term models the commonsense constraint that scene-centric human attributes shall agree with the observation in individual views:

$$\mathcal{E}_{Attr}(g, \tilde{g}) = \sum_{a \in g} \mathbf{1}(a \neq \tilde{a}) \cdot \xi, \quad (8)$$

where $\mathbf{1}(\cdot)$ is an indicator function and ξ is a constant energy penalty introduced when the two predictions mismatch.

5 Inference

The inference process consists of two sub-steps: (i) matching object nodes Φ in scene-centric and view-centric parse graphs (i.e. the structure of parse graph hierarchy) and (ii) estimating optimal values of parse graphs $\{g, \tilde{g}^{(1)}, \dots, \tilde{g}^{(M)}\}$.

The overall procedure is as follows: we first obtain view-centric objects, actions, and attributes proposals from pre-trained detectors on all video frames. This forms the initial view-centric predictions $\{\tilde{g}^{(1)}, \dots, \tilde{g}^{(M)}\}$. Next we use a Markov Chain Monte Carlo (MCMC) sampling algorithm to optimize the parse graph structure Φ . Given a fixed parse graph hierarchy, variables within the scene-centric and view-centric parse graphs $\{g, \tilde{g}^{(1)}, \dots, \tilde{g}^{(M)}\}$ can be efficiently estimated by a dynamic programming algorithm. These two steps are performed iteratively until convergence.

5.1 Inferring Parse Graph Hierarchy

We use a stochastic algorithm to traverse the solution space of the parse graph hierarchy Φ . To satisfy the detailed balance condition, we define three reversible operators $\Theta = \{\Theta_1, \Theta_2, \Theta_3\}$ as follows.

Merging. The merging operator Θ_1 groups a view-centric parse graph with an other view-centric parse graph by creating a scene-centric parse graph that connects the two. The operator requires the two operands to describe two objects of the same type either from different views or in the same view but with non-overlapping time intervals.

Splitting. The splitting operator Θ_2 splits a scene-centric parse graph into two parse graphs such that each resulting parse graph only connects to a subset of view-centric parse graphs.

Swapping. The swapping operator Θ_3 swaps two view-centric parse graphs. One can view the swapping operator as a shortcut of merging and splitting combined.

We define the proposal distribution $q(G \rightarrow G')$ as an uniform distribution. At each iteration, we generate a new structure proposal Φ' by applying one of the three operators Θ_i with respect to probability 0.4, 0.4, and 0.2, respectively. The generated proposal is then accepted with respect to an acceptance rate $\alpha(\cdot)$ as in the Metropolis-Hastings algorithm (Metropolis et al. 1953):

$$\alpha(G \rightarrow G') = \min \left(1, \frac{q(G' \rightarrow G) \cdot p(G'|x)}{q(G \rightarrow G') \cdot p(G|x)} \right), \quad (9)$$

where $p(G|x)$ the posterior is defined in Eqn. (1).

5.2 Inferring Parse Graph Variables

Given a fixed parse graph hierarchy, we need to estimate the optimal value for each node within each parse graph. As illustrated in Fig. 3, for each frame, the scene-centric node g_t and the corresponding view-centric nodes $\tilde{g}_t^{(i)}$ form a star model, and the whole scene-centric nodes are regarded as a Markov chain in the temporal order. Therefore the proposed model is essentially a Directed Acyclic Graph (DAG). To infer the optimal node values, we can simply apply the standard factor graph belief propagation (sum-product) algorithm.

6 Experiments

6.1 Setup and Datasets

We evaluate our scene-centric joint-parsing framework in tasks including object detection, multi-object tracking, action recognition, and human attributes recognition. In object detection and multi-object tracking tasks, we compare with published results. In action recognition and human attributes tasks, we compare the performance of view-centric proposals without joint parsing and scene-centric predictions after joint parsing as well as additional baselines. The following datasets are used to cover a variety of tasks.

The **CAMPUS** dataset (Xu et al. 2016)¹ contains video sequences from four scenes each captured by four cameras. Different from other multi-view video datasets focusing solely on multi-object tracking task, videos in the CAMPUS dataset contains richer human poses and activities with moderate overlap in the fields of views between cameras. In addition to the tracking annotation in the CAMPUS dataset, we collect new annotation that includes 5 action categories and 9 attribute categories for evaluating action and attribute recognition.

The **TUM Kitchen** dataset (Tenorth, Bandouch, and Beetz 2009)² is an action recognition dataset that contains 20 video sequences captured by 4 cameras with overlapping views. As we only focusing on the RGB imagery inputs in our framework, other modalities such as motion capturing, RFID tag reader signals, magnetic sensor signals are not used as inputs in our experiments. To evaluate detection and tracking task, we compute human bounding boxes from motion capturing data by projecting 3D human poses to the image planes of all cameras using the intrinsic and extrinsic parameters provided in the dataset. To evaluate human attribute tasks, we annotate 9 human attribute categories for every subject.

In our experiments, both the CAMPUS and the TUM Kitchen datasets are used in all tasks. In the following subsection, we present isolated evaluations.

6.2 Evaluation

Object detection & tracking. We use FasterRCNN (Ren et al. 2015) to create initial object proposals on all video frames. The detection scores are used in the likelihood term in Eqn. (2). During joint parsing, objects which are not initially detected on certain views are projected from object’s scene-centric positions with the camera matrices. After joint parsing, we extract all bounding boxes that are grounded by object nodes from each view-centric parse graph to compute multi-object detection accuracy (DA) and precision (DP). Concretely, the accuracy measures the fraction of correctly detected objects among all ground-truth objects and the precision is computed as fraction of true-positive predictions among all output predictions. A predicted bounding box is considered a match with a ground-truth box only if the intersection over union (IoU) score is greater than 0.5. When more than one prediction overlaps with a ground-truth box,

¹bitbucket.org/merayxu/multiview-object-tracking-dataset

²ias.in.tum.de/software/kitchen-activity-data

CAMPUS-S1	DA (%)	DP (%)	TA (%)	TP (%)	IDSW	FRAG
Fleuret et al.	24.52	64.28	22.43	64.17	2269	2233
Berclaz et al.	30.47	62.13	28.10	62.01	2577	2553
Xu et al.	49.30	72.02	56.15	72.97	320	141
Ours	56.00	72.98	55.95	72.77	310	138
CAMPUS-S2	DA (%)	DP (%)	TA (%)	TP (%)	IDSW	FRAG
Fleuret et al.	16.51	63.92	13.95	63.81	241	214
Berclaz et al.	24.35	61.79	21.87	61.64	268	249
Xu et al.	27.81	71.74	28.74	71.59	1563	443
Ours	28.24	71.49	27.91	71.16	1615	418
CAMPUS-S3	DA (%)	DP (%)	TA (%)	TP (%)	IDSW	FRAG
Fleuret et al.	17.90	61.19	16.15	61.02	249	235
Berclaz et al.	19.46	59.45	17.63	59.29	264	257
Xu et al.	49.71	67.02	49.68	66.98	219	117
Ours	50.60	67.00	50.55	66.96	212	113
CAMPUS-S4	DA (%)	DP (%)	TA (%)	TP (%)	IDSW	FRAG
Fleuret et al.	11.68	60.10	11.00	59.98	828	812
Berclaz et al.	14.73	58.51	13.99	58.36	893	880
Xu et al.	24.46	66.41	24.08	68.44	962	200
Ours	24.81	66.59	24.63	68.28	938	194
TUM Kitchen	DA (%)	DP (%)	TA (%)	TP (%)	IDSW	FRAG
Fleuret et al.	69.88	64.54	69.67	64.76	61	57
Berclaz et al.	72.39	63.27	72.20	63.51	48	44
Xu et al.	86.53	72.12	86.18	72.37	9	5
Ours	89.13	72.21	88.77	72.42	12	8

Table 1: Quantitative comparisons of multi-object tracking on CAMPUS and TUM Kitchen datasets.




Figure 4: Confusion matrices of action recognition on view-centric proposals (left) and scene-centric predictions (right).

only the one with the maximum overlap is counted as true positive.

When extracting all bounding boxes on which the view-centric parse graphs are grounded and grouping them according to the identity correspondence between different views, we obtain object trajectories with identity matches across multiple videos. In the evaluation, we compute four major tracking metrics: multi-object tracking accuracy (TA), multi-object track precision (TP), the number of identity switches (IDSW), and the number of fragments (FRAG). A higher value of TA and TP and a lower value of IDSW and FRAG indicate the tracking method works better. We report quantitative comparisons with several published methods (Xu et al. 2016; Berclaz et al. 2011; Fleuret et al. 2008) in Table 1. From the results, the performance measured by tracking metrics are comparable to published results. We conjecture that the appearance similarity is the main drive for establish cross-view correspondence while additional semantic attributes proved limited gain to the tracking task.

Action recognition. View-centric action proposals are

Methods	CAMPUS						TUM Kitchen								
	Run	PickUp	PutDown	Throw	Catch	Overall	Reach	Taking	Lower	Release	OpenDoor	CloseDoor	OpenDrawer	CloseDrawer	Overall
view-centric	0.83	0.76	0.91	0.86	0.80	0.82	0.78	0.66	0.75	0.67	0.48	0.50	0.50	0.42	0.59
baseline-vote	0.85	0.80	0.71	0.88	0.82	0.73	0.80	0.63	0.77	0.71	0.72	0.73	0.70	0.47	0.69
baseline-mean	0.86	0.82	1.00	0.90	0.87	0.88	0.79	0.61	0.75	0.69	0.67	0.67	0.66	0.45	0.66
scene-centric	0.87	0.83	1.00	0.91	0.88	0.90	0.81	0.67	0.79	0.71	0.71	0.73	0.70	0.50	0.70

Table 2: Quantitative comparisons of human action recognition on CAMPUS and TUM Kitchen datasets.

CAMPUS	Methods	Gender	Long hair	Glasses	Hat	T-shirt	Long sleeve	Shorts	Jeans	Long pants	mAP
	view-centric	0.59	0.77	0.56	0.76	0.36	0.59	0.70	0.63	0.35	0.59
baseline-mean	0.63	0.82	0.55	0.75	0.34	0.64	0.69	0.63	0.34	0.60	
baseline-vote	0.61	0.82	0.55	0.75	0.34	0.65	0.69	0.63	0.35	0.60	
scene-centric	0.76	0.82	0.62	0.80	0.40	0.62	0.76	0.62	0.24	0.63	
TUM Kitchen	Methods	Gender	Long hair	Glasses	Hat	T-shirt	Long sleeve	Shorts	Jeans	Long pants	mAP
	view-centric	0.69	0.93	0.32	1.00	0.50	0.89	0.91	0.83	0.73	0.76
baseline-mean	0.86	1.00	0.32	1.00	0.54	0.96	1.00	0.83	0.81	0.81	
baseline-vote	0.64	1.00	0.32	1.00	0.32	0.93	1.00	0.83	0.76	0.76	
scene-centric	0.96	0.98	0.32	1.00	0.77	0.96	0.94	0.83	0.83	0.84	

Table 3: Quantitative comparisons of human attribute recognition on CAMPUS and TUM Kitchen datasets.

obtained from a fully-connected neural network with 5 hidden layers and 576 neurons which predicts action labels using human pose. For the CAMPUS dataset, we collect additional annotations for 5 human action classes: Run, PickUp, PutDown, Throw, and Catch in total of 8,801 examples. For the TUM Kitchen dataset, we evaluate on the 8 action categories: Reaching, TakingSomething, Lowering, Releasing, OpenDoor, CloseDoor, OpenDrawer, and CloseDrawer. We measure both individual accuracies for each category as well as the overall accuracies across all categories. Table 2 shows the performance of scene-centric predictions with view-centric proposals, and two additional fusing strategies as baselines. Concretely, the *baseline-vote* strategy takes action predictions from multiple views and outputs the label with majority voting, while the *baseline-mean* strategy assumes equal priors on all cameras and outputs the label with the highest averaged probability. When evaluating scene-centric predictions, we project scene-centric labels back to individual bounding boxes and calculate accuracies following the same procedure as evaluating view-centric proposals. Our joint parsing framework demonstrates improved results as it aggregates marginalized decisions made on individual views while also encourages solutions that comply with other tasks. Fig. 4 compares the confusion matrix of view-centric proposals and scene-centric predictions after joint parsing for CAMPUS dataset. To further understand the effect of multiple views, we break down classification accuracies by the number of cameras where persons are observed (Fig. 5). Observing an entity from more cameras generally leads to better performance, while too many conflicting observations may also cause degraded performance. Fig. 6 shows some success and failure examples.

Human attribute recognition. We follow the similar procedure as in the action recognition case above. Additional annotations for 9 different types of human attributes are collected for both CAMPUS and TUM Kitchen dataset. View-centric proposals and score are obtained from an attribute grammar model as in (Park, Nie, and Zhu 2016). We mea-




Figure 5: The breakdown of action recognition accuracy according to the number of camera views in which each entity is observed.

sure performance with average precisions for each attribute categories as well as mean average precision (mAP) as in human attribute literatures. Scene-centric predictions are projected to bounding boxes in each views when calculating precisions. Table 3 shows quantitative comparisons between view-centric and scene-centric predictions. The same baseline fusing strategies as in the action recognition task are used. The scene-centric prediction outperforms the original proposals in 7 out of 9 categories while remains comparable in others. Notably, the CAMPUS dataset is harder than standard human attribute datasets because of occlusions, limited scales of humans, and irregular illumination conditions.

6.3 Runtime

With initial view-centric proposals precomputed, for a 3-minute scene shot by 4 cameras containing round 15 entities, our algorithm performs at 5 frames per second on average. With further optimization, our proposed method can run in real-time. Note that although the proposed framework uses a sampling-based method, using view-based proposals as initialization warm-starts the sampling procedure. Therefore,



Figure 6: Success (1st row) and failure examples (2nd row) of view-centric (labels overlaid on the images) and scene-centric predictions (labels beneath the images) of action and attribute recognition tasks. For failure examples, true labels are in the bracket. “Occluded” means that the locations of objects or parts are projected from scene locations and therefore no view-centric proposals are generated. Better viewed in color.

the overall runtime is significantly less than searching the entire solution space from scratch. For problems of a larger size, more efficient MCMC algorithms may be adopted. For example, the mini-batch acceptance testing technique (Chen et al. 2016) has demonstrated several order-of-magnitude speedups.

7 Conclusion

We represent a joint parsing framework that computes a hierarchy of parse graphs which represents a comprehensive understanding of cross-view videos. We explicitly specify various constraints that reflect the appearance and geometry correlations among objects across multiple views and the correlations among different semantic properties of objects. Experiments show that the joint parsing framework improves view-centric proposals and produces more accurate scene-centric predictions in various computer vision tasks.

We briefly discuss advantages of our joint parsing framework and potential future directions from two perspectives.

Explicit Parsing While the end-to-end training paradigm is appealing in many *data-rich* supervised learning scenarios, as an extension, leveraging loosely-coupled pre-trained modules and exploring commonsense constraints can be helpful when large-scale training data is not available or too expensive to collect in practice. For example, many applications in robotics and human-robot interaction domains share the same set of underlying perception units such as scene understanding, object recognition, etc. Training for every new scenarios entirely could end up with exponential number of possibilities. Leveraging pre-trained modules and explore correlation and constraints among them can be treated as a factorization of the problem space. Therefore, the explicit joint parsing scheme allows practitioners to leverage

pre-trained modules and to build systems with an expanded skill set in a scalable manner.

Interpretable Interface. Our joint parsing framework not only provides a comprehensive scene-centric understanding of the scene, moreover, the scene-centric spatio-temporal parse graph representation is an interpretable interface of computer vision models to users. In particular, we consider the following properties an explainable interface shall have apart from the correctness of answers:

- *Relevance*: an agent shall recognize the intent of humans and provide information relevant to humans’ questions and intents.
- *Self-explainability*: an agent shall provide information that can be interpreted by humans as how answers are derived. This criterion promotes humans’ trust on an intelligent agent and enables sanity check on the answers.
- *Consistency*: answers provided by an agents shall be consistent throughout an interaction with humans and across multiple interaction sessions. Random or non-consistent behaviors cast doubts and confusions regarding the agent’s functionality.
- *Capability*: an explainable interface shall help humans understand the boundary of capabilities of an agent and avoid blinded trusts.

Potential future directions include quantifying and evaluating the interpretability and user satisfaction by conducting user studies.

References

Aditya, S.; Baral, C. Y. Y. A. Y., and Fermuller, C. 2016. Deepiu: An architecture for image understanding. In *IEEE Conference on Computer Vision and Pattern Recognition*.

- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Zitnick, C. L.; and Parikh, D. 2015. Vqa: Visual question answering. In *IEEE International Conference on Computer Vision*.
- Berclaz, J.; Fleuret, F.; Turetken, E.; and Fua, P. 2011. Multiple object tracking using k-shortest paths optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33(9):1806–1819.
- Biran, O., and McKeown, K. 2014. Justification narratives for individual classifications. In *IEEE International Conference on Machine Learning Workshops*.
- Chen, H.; Seita, D.; Pan, X.; and Canny, J. 2016. An efficient minibatch acceptance test for Metropolis-Hastings. *arXiv preprint arXiv:1610.06848*.
- Core, M. G.; Lane, H. C.; Van Lent, M.; Gomboc, D.; Solomon, S.; and Rosenberg, M. 2006. Building explainable artificial intelligence systems. In *AAAI Conference on Artificial Intelligence*.
- Fleuret, F.; Berclaz, J.; Lengagne, R.; and Fua, P. 2008. Multi-camera people tracking with a probabilistic occupancy map. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 30(2):267–282.
- Geman, D.; Geman, S.; Hallonquist, N.; and Younes, L. 2015. Visual Turing test for computer vision systems. *Proceedings of the National Academy of Sciences* 112(12):3618–3623.
- Han, F., and Zhu, S. 2009. Bottom-up/top-down image parsing with attribute grammar. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31(1):59–73.
- Hendricks, L. A.; Akata, Z.; Rohrbach, M.; Donahue, J.; Schiele, B.; and Darrell, T. 2016. Generating visual explanations. In *European Conference on Computer Vision*.
- Hofmann, M.; Wolf, D.; and Rigoll, G. 2013. Hypergraphs for joint multi-view reconstruction and multi-object tracking. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Koppula, H. S., and Saxena, A. 2016. Anticipating human activities using object affordances for reactive robotic response. *IEEE transactions on Pattern Analysis and Machine Intelligence* 38(1):14–29.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalanditis, Y.; Li, L.-J.; Shamma, D. A.; Bernstein, M.; and Fei-Fei, L. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal on Computer Vision* 123(1):32–73.
- Lane, H. C.; Core, M. G.; Van Lent, M.; Solomon, S.; and Gomboc, D. 2005. Explainable artificial intelligence for training and tutoring. Technical report, Defense Technical Information Center.
- Leal-Taixe, L.; Pons-Moll, G.; and Rosenhahn, B. 2012. Branch-and-price global optimization for multi-view multi-object tracking. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Liebelt, J., and Schmid, C. 2010. Multi-view object class detection with a 3D geometric model. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Liu, H., and Singh, P. 2004. Conceptnets practical commonsense reasoning tool-kit. *BT technology journal* 22(4):211–226.
- Liu, X.; Zhao, Y.; and Zhu, S.-C. 2014. Single-view 3D scene parsing by attributed grammar. In *IEEE Conference Computer Vision and Pattern Recognition*.
- Lomas, M.; Chevalier, R.; Cross II, E. V.; Garrett, R. C.; Hoare, J.; and Kopack, M. 2012. Explaining robot actions. In *ACM/IEEE International Conference on Human-Robot Interaction*.
- Metropolis, N.; Rosenbluth, A.; Rosenbluth, M.; Teller, A.; and Teller, E. 1953. Equation of state calculations by fast computing machines. *Journal of Chemical Physics* 21(6):1087–1092.
- Park, S.; Nie, B. X.; and Zhu, S.-C. 2016. Attribute and-or grammar for joint parsing of human attributes, part and pose. *arXiv preprint arXiv:1605.02112*.
- Pero, L.; Bowdish, J.; Hartley, E.; Kermgard, B.; and Barnard, K. 2013. Understanding bayesian rooms using composite 3D object models. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Pirsiavash, H.; Ramanan, D.; and Fowlkes, C. 2011. Globally-optimal greedy algorithms for tracking a variable number of objects. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Qi, H.; Wu, T.; Lee, M.-W.; and Zhu, S.-C. 2015. A restricted visual Turing test for deep scene and event understanding. *arXiv preprint arXiv:1512.01715*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Annual Conference on Neural Information Processing Systems*.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. Why should I trust you?: Explaining the predictions of any classifier. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144. ACM.
- Tenorth, M.; Bandouch, J.; and Beetz, M. 2009. The TUM kitchen data set of everyday manipulation activities for motion tracking and action recognition. In *IEEE International Conference on Computer Vision Workshop*.
- Tu, K.; Meng, M.; Lee, M. W.; Choe, T. E.; and Zhu, S.-C. 2014. Joint video and text parsing for understanding events and answering queries. *IEEE MultiMedia* 21(2):42–70.
- Utasi, A., and Benedek, C. 2011. A 3-D marked point process model for multi-view people detection. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Van Lent, M.; Fisher, W.; and Mancuso, M. 2004. An explainable artificial intelligence system for small-unit tactical behavior. In *National Conference on Artificial Intelligence*.
- Wang, J.; Nie, X.; Xia, Y.; Wu, Y.; and Zhu, S.-C. 2014. Cross-view action modeling, learning and recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Wei, P.; Zhao, Y.; Zheng, N.; and Zhu, S.-C. 2016. Modeling 4D human-object interactions for joint event segmentation, recognition, and object localization.
- Xu, Y.; Lin, L.; Zheng, W.-S.; and Liu, X. 2013. Human re-identification by matching compositional template with cluster sampling. In *IEEE International Conference on Computer Vision*.
- Xu, Y.; Ma, B.; Huang, R.; and Lin, L. 2014. Person search in a scene by jointly modeling people commonness and person uniqueness. In *ACM Multimedia Conference*.
- Xu, Y.; Liu, X.; Liu, Y.; and Zhu, S.-C. 2016. Multi-view people tracking via hierarchical trajectory composition. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Xu, Y.; Liu, X.; Qin, L.; and Zhu, S.-C. 2017. Multi-view people tracking via hierarchical trajectory composition. In *AAAI Conference on Artificial Intelligence*.
- Yao, B. Z.; Yang, X.; Lin, L.; Lee, M. W.; and Zhu, S.-C. 2010. I2t: Image parsing to text description. *Proceedings of the IEEE* 98(8):1485–1508.
- Zhu, Y.; Groth, O.; Bernstein, M.; and Fei-Fei, L. 2016. Visual7W: Grounded question answering in images. In *Advances of Cognitive Systems*.