

ContextHOI: Spatial Context Learning for Human-Object Interaction Detection

Mingda Jia¹, Liming Zhao², Ge Li^{1*}, Yun Zheng²,

¹Guangdong Provincial Key Laboratory of Ultra High Definition Immersive Media Technology, Shenzhen Graduate School, Peking University

²Alibaba Group

mingdaj@stu.pku.edu.cn, geli@ece.pku.edu.cn, {lingchen.zlm, zhengyun.zy}@alibaba-inc.com

Abstract

Spatial contexts, such as the backgrounds and surroundings, are considered critical in Human-Object Interaction (HOI) recognition, especially when the instance-centric foreground is blurred or occluded. Recent advancements in HOI detectors are usually built upon detection transformer pipelines. While such an object-detection-oriented paradigm shows promise in localizing objects, its exploration of spatial context is often insufficient for accurately recognizing human actions. To enhance the capabilities of object detectors for HOI detection, we present a dual-branch framework named ContextHOI, which efficiently captures both object detection features and spatial contexts. In the context branch, we train the model to extract informative spatial context without requiring additional hand-craft background labels. Furthermore, we introduce context-aware spatial and semantic supervision to the context branch to filter out irrelevant noise and capture informative contexts. ContextHOI achieves state-of-the-art performance on the HICO-DET and v-coco benchmarks. For further validation, we construct a novel benchmark, HICO-ambiguous, which is a subset of HICO-DET that contains images with occluded or impaired instance cues. Extensive experiments across all benchmarks, complemented by visualizations, underscore the enhancements provided by ContextHOI, especially in recognizing interactions involving occluded or blurred instances.

1 Introduction

Human-object interaction (HOI) detection (Gao et al. 2018) involves identifying instance locations and their interactions, typically represented as $\langle \text{Human}, \text{Interaction}, \text{Object} \rangle$ triplets. Visual uncertainties in real-world scenes, such as occlusion, significantly affect the discernibility of the subjects (Sun et al. 2022). Thus, a primary challenge of HOI detection is accurately inferring the interactions with limited visual cues.

Recent advancements in HOI detection are driven by various approaches, including those based on convolutional networks (Wang et al. 2019), and graphs (Gao et al. 2020). HOI detectors built upon transformers generally demonstrate superior performance (Ning et al. 2023). Both one-stage and two-stage HOI detectors perform inference depending on

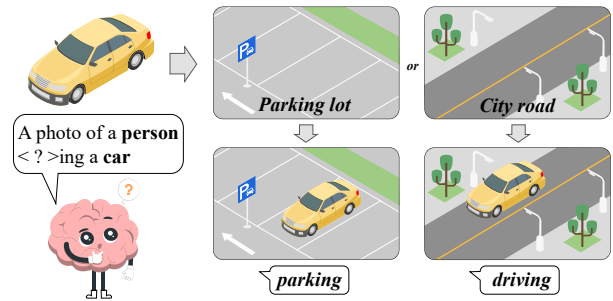


Figure 1: **The role of context learning in HOI Detection.** Spatial context, like a *parkinglot* or a *cityroad* helps little with identify the salient car. However, context is critical in distinguishing human interactions. Both *parking* and *driving* are highly related to the context information.

the captured instance-centric attributes, while the two-stage methods rely more heavily on the confidence of their pre-trained object detection backbones (Zhang et al. 2022).

Despite significant advancements, existing HOI detectors face a major difficulty. They are vulnerable in **identifying HOI scenes with limited foreground visual cues**, such as when subjects are blurred or occluded. In contrast, Humans can accurately recognize HOIs even when instances are unclear or completely unseen. For example, we can easily infer the presence of a driver driving in a car on the highway, though the driver is obscured by the tinted windows of the car, as illustrated in Figure 1.

Upon analysing such difficulties, we observe a potential tendency to **overemphasize instance-centric attributes** in previous HOI detectors. Detection transformers (DETR), foundational to many HOI detectors, often capture minimal context information while neglecting most backgrounds as negative or irrelevant. Although this approach is a consensus of *object detection*, it may not be sufficient enough for robust *action recognition*, which requires a more comprehensive scene understanding. Therefore, the core issue appears to be the misalignment between the minimal context modelling inherent in object detection and the extensive contextual requirements essential for HOI detection. While

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

acknowledging the impressive capabilities of transformer-based object detectors, our goal is to adapt them to HOI detection better.

A practical adaption strategy is augmenting instance features from object detectors with additional contextual information. Hence, we propose ContextHOI, a dual-branch framework that integrates object detection features with learned context to enhance the robustness of HOI predictions. Specifically, we design multiple novel modules in the context learning branch to efficiently capture informative contextual features.

In the context branch, we propose a context extractor that grounds context features within the image to capture the context regions. Recent HOI detectors have started to capture global visual features and learn backgrounds heuristically. However, these approaches might degrade into mere replications of instance-centric modules due to the lack of context-aware supervision. To mitigate such risk, we incorporate explicit spatial and semantic supervision into the context extractor. These supervisions enable efficient context learning without additional hand-craft context labels or segmentation priors.

For the **spatial supervision**, we propose a series of self-supervised *spatially contrastive constraints*. These constraints provide multi-level coarse-to-fine constraints, aiming to increase the margin between context and instance regions. Such an objective guides the context branch to focus on regions out of RoIs while minimising the background noise for the instance branch. Furthermore, we propose a dynamic distance weight in one of the constraints, which allows the context extractor to capture contexts with suitable spatial distributions rather than regressing regions near the image margins.

For the **semantic supervision**, we introduce a *semantic-guided context explorer*. Fascinated by the significant capability of pre-trained VLMs (Radford et al. 2021; Sun et al. 2023) in representing the context (An et al. 2023), we integrate their prior knowledge into our context explorer. Additionally, we introduce an adaptive knowledge distillation approach to bridge the gap between the embedding space of VLMs and the feature space of our visual encoder.

We introduce a context aggregator to extract further more crucial context features that can complete the detection features. The context aggregator grounds the sampled context features according to the object detection features while enhancing the communication between the local and global information. Finally, these features are fused for the interaction prediction.

Moreover, we propose a new benchmark, HICO-DET (ambiguous), for evaluating HOI detectors on scenes with unclear foreground visual contents. ContextHOI demonstrates competitive performance on the regular setting of HICO-DET and v-coco while demonstrating significantly stronger robustness than previous HOI detectors on HICO-DET (ambiguous). Additionally, ContextHOI also obtains zero-shot ability.

In this paper, our contributions are threefold:

- We revisit and explore the significance of spatial contexts in HOI detection, which is not sufficiently discussed in

recent HOI methods.

- To the best of our knowledge, we are the first to systematically learn spatial contexts for enhancing human-object interaction detection. Our model can be trained to capture informative context features without additional background labels or segmentation priors.
- ContextHOI achieves competitive results on HICO-DET and v-coco while obtaining significant state-of-the-art performance on HICO-DET (ambiguous), containing HOI scenes with unclear or occluded foregrounds, highlighting the robustness of our approach.

2 Related Works

HOI detection. Based on transformers, one-stage HOI detectors (Tamura et al. 2021; Ning et al. 2023) predict HOIs with shared (Kim et al. 2021), parallel (Kim, Jung, and Cho 2023), or sequential (Zhang et al. 2021) transformers. In contrast, two-stage detectors (Zhang et al. 2022; Lei et al. 2023) utilize a pre-trained object detector backbone to capture instance locations and foreground attributes at the first stage. In the second stage, they predict interactions using specialized modules.

However, both one-stage and two-stage HOI methods depend heavily on transformer-based object detection pipelines, which limit their ability to model spatial contexts. Thus, these HOI detectors often struggle to recognize images with unclear foregrounds. This limitation is more critical in two-stage approaches, as detailed in Table 2. Such issues are foreseen by CDN (Zhang et al. 2021). However, the exploration of systematically integrating context learning into HOI remains unexplored.

Spatial context learning in HOI. Spatial contexts are crucial in relative tasks like object detection (Chen, Huang, and Tao 2018), scene-graph generation (Yang et al. 2018; Zhai et al. 2023) and group recognition (Yuan and Ni 2021).

In HOI detection, several works capture contexts with customized spatial modules. Early methods used dense graphs (Gao et al. 2020; Frederic Z. Zhang and Gould 2021) for global context. Later approaches focused on instance-centric contexts (Gao et al. 2018; Wang et al. 2019). Recently, BCOM (Wang et al. 2024) attempted to enhance instance features through occluded part extrapolation (OPE). However, they still depended heavily on object detection features. In contrast, our method extracts context features independently of instances. Such an approach ensures the focused context regions are not limited to instance-centric RoIs, providing a broader understanding of the scene. Secondly, it mitigates the error accumulation from failures in object detection backbones.

As for the semantic-guided learning approaches, some work (Yuan et al. 2022, 2023) leverages utilize learnable text embeddings to represent pseudo background categories, while the others (Lei et al. 2023; Wang et al. 2024) learn a global concept memory to explore background information. However, their approaches are implicit or heuristic. The lack of explicit spatial constraints might downgrade their context-aware modules to repetitions of instance-centric modules. To address this, we propose explicit spatial supervision.

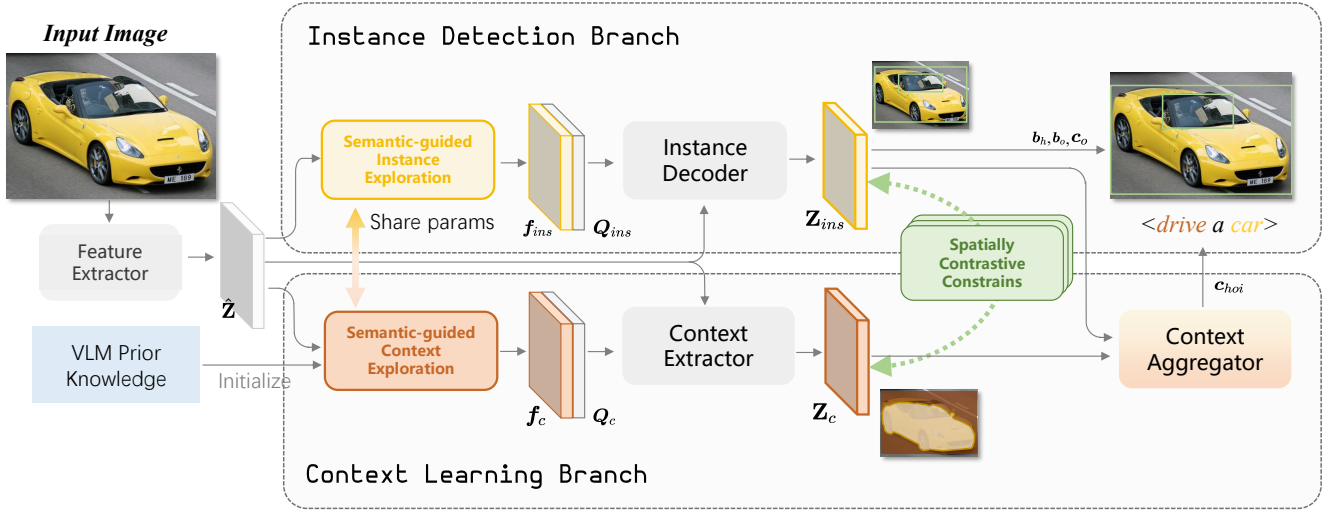


Figure 2: **The overall architecture of ContextHOI.** ContextHOI has a dual-branch and fusion structure, with instance detection and context learning branches. The instance detection branch captures instance-centric attributes, while the context learning branch focuses on instance-independent context features. We introduce a semantic-guided instance/context exploration module to distil prior knowledge from VLM to help ground informative visual content. A set of spatially contrastive constraints supervises the learned instances and contexts to focus on different visual aspects. Finally, a context aggregator will fuse the instance and context feature for HOI prediction.

3 Method

This section details ContextHOI, a framework that learns informative spatial context for HOI detection. In section 3.1, we outline the overall architecture. In sections 3.2 and 3.3, we deliver the design of the two supervisions on the context extractor, the spatially contrastive constraints, and the semantic-guided context explorer. In section 3.4, we present the context aggregator. Finally, we introduce the training method in section 3.5.

3.1 Overall Architecture

The network design of ContextHOI is illustrated in Figure 2, our framework follows a parallel dual-branch architecture, with an instance detection branch and a context learning branch. Given an input image, we first utilize a pre-trained visual encoder (Carion et al. 2020) to summarize the image feature map $\hat{\mathbf{Z}}$. The image feature will be fed into both branches as the visual content memory.

In the instance detection branch, an instance decoder with several transformer decoder layers will utilize a set of instance queries $\mathbf{Q}_{ins} \in \mathbb{R}^{2N_q \times C}$ to ground $\hat{\mathbf{Z}}$ and generate instance-centric features \mathbf{Z}_{ins} . Several prediction heads take the captured feature and generate predictions for detection-oriented tasks, including the human bounding box $\mathbf{B}_h \in \mathbb{R}^{N_q \times 4}$, object bounding box $\mathbf{B}_o \in \mathbb{R}^{N_q \times 4}$ and object categories $\mathbf{C}_o \in \mathbb{R}^{N_q \times N_o}$, where N_q is the number of queries, N_o is the number of object classes.

In the context learning branch, our context extractor, which shares the same architecture as the instance decoder, captures context feature \mathbf{Z}_c by the context queries $\mathbf{Q}_c \in \mathbb{R}^{N_q \times C}$. Then, our context aggregator integrates the instance-centric feature and contexts. A single prediction

head will take the aggregated feature and predict the hoi categories $\mathcal{C}_{hoi} \in \mathbb{R}^{N_q \times N_{hoi}}$, where N_{hoi} are the number of HOI triplet combinations.

3.2 Spatially Contrastive Constraints

In the following part, we propose a set of spatially contrastive constraints in three coarse-to-fine supervision levels, including feature-level constraint \mathcal{L}_{FC} , region-level constraint \mathcal{L}_{RC} and instance-level constraint \mathcal{L}_{IC} . These constraints work as training loss on the outputs of the instance decoder and context extractor simultaneously. The main objective of these constraints is to diverge the attention areas of the instance decoder and context extractor. Furthermore, we introduce dynamic distance weight into \mathcal{L}_{IC} to maintain the suitable spatial shape of the learned contexts.

Feature-level constraint. Given the instance features \mathbf{Z}_{ins} and the context features \mathbf{Z}_c , we calculate the absolute value of their mean cosine similarity along the query dimension. The feature patch with much higher similarity to the corresponding instance patches is seen as a propagation negative through the context features. \mathcal{L}_{FC} can be expressed by:

$$\mathcal{L}_{FC} = \frac{1}{|\Phi|} \sum_{k=1}^{N_q} \frac{|\hat{\mathbf{z}}_{ins}^{kT} \hat{\mathbf{z}}_c^k|}{\|\hat{\mathbf{z}}_{ins}^{kT}\|_1 \|\hat{\mathbf{z}}_c^k\|_1 + \epsilon}, k \notin \Phi, \quad (1)$$

in which the summarized instance and context features from dual-decoders are reshaped to patched feature tokens $\{\hat{\mathbf{z}}_{ins}^k\}_{k=1}^{N_q} \in \mathbb{R}^{L_{dec} \times C}$ and $\{\hat{\mathbf{z}}_c^k\}_{k=1}^{N_q} \in \mathbb{R}^{L_{dec} \times C}$. L_{dec} is the number of decoder layers shared by instance decoder and context extractor. We do not calculate \mathcal{L}_{FC} on Φ , which are features indicated by zero-padded queries.

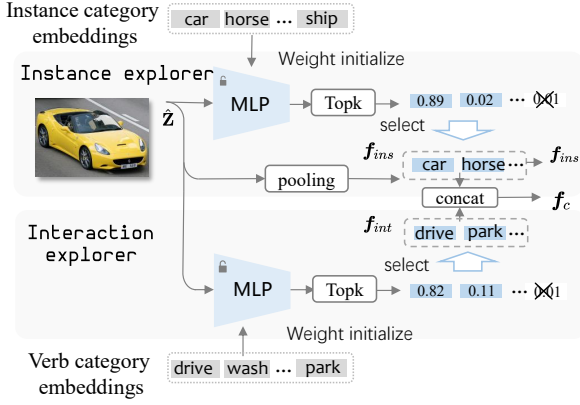


Figure 3: **Inner design of semantic-guided context exploration** module. pooling refers to mean pooling on the spatial dimension of $\hat{\mathbf{Z}}$, concat refers to concatenation.

Region-level constraint. At the region level, we constrain the learned positional guided embeddings $\mathbf{P}_{ins} \in \mathbb{R}^{N_q \times C}$ and $\mathbf{P}_c \in \mathbb{R}^{N_q \times C}$ predicted together with instance and context features by corresponding decoders. As the guided embeddings take the role of directing transformer attention to specified local or global regions (Liao et al. 2022), we constrain \mathbf{P}_{ins} and \mathbf{P}_c to be distinct in the query position domain by reversed L1-distance. This constraint includes an exp operation to maintain a non-negative loss form. The constraint can be formulated as:

$$\mathcal{L}_{RC} = \frac{1}{N_q} \sum_{k=1}^{N_q} \exp(-\|\mathbf{p}_{ins}^k - \mathbf{p}_c^k\|_1), \quad (2)$$

in which \mathbf{p}_{ins}^k and \mathbf{p}_c^k are instance-guided embedding and context-guided embedding at query position k .

Instance-level constraint. Aiming to capture \mathbf{Z}_c with structured optimization goals, we extend the context extractor with a simple MLP, predicting context region boxes $\mathbf{B}_c \in \mathbb{R}^{N_q \times 4}$. Then we reverse the Generalized Intersection-over-Union (GIoU) (Rezatofighi et al. 2019) to formulate the loss, calculated on predicted $\hat{\mathbf{B}}_c$ and ground-truth \mathbf{B}_h and \mathbf{B}_o . However, these constraints are too strict, which causes the context to shift to the edges of the images. To address this, we introduce a dynamic distance weight to improve context learning. This weight is calculated as:

$$\mathcal{W}_d(\mathbf{b}_i, \mathbf{b}_j) = \exp\left(-\frac{|\mathbf{b}_i - \mathbf{b}_j|}{\tau + \epsilon}\right), \quad (3)$$

where τ is a learnable parameter describing a dynamic margin between the considered boxes. With this weight, the instance-level constraint can be expressed by:

$$\mathcal{L}_{IC} = \frac{1}{2|\Phi|} \sum_{k=1}^{N_q} \left[2 + \mathcal{W}_d(\mathbf{b}_h, \hat{\mathbf{b}}_c^k) \text{GIoU}(\mathbf{b}_h, \hat{\mathbf{b}}_c^k) + \mathcal{W}_d(\mathbf{b}_o, \hat{\mathbf{b}}_c^k) \text{GIoU}(\mathbf{b}_o, \hat{\mathbf{b}}_c^k) \right], k \notin \Phi, \quad (4)$$

in which \mathbf{b}_h , \mathbf{b}_o and $\hat{\mathbf{b}}_c^k$ are ground-truth human box, ground-truth object box and corresponding context coordinate pre-

dictions at query dimension k . Thanks to the dynamic distance weight, the strength of the positive signal from instance-level constraint diminishes significantly as the captured context regions approach the edges of the image.

3.3 Semantic-guided Context Exploration

The semantic-guided context exploration module, detailed in Figure 3, provides semantic guidance to the following context extractor. Specifically, we construct learnable instance- and interaction-explorers in the exploration module. To distill VLM knowledge to these explorers, we initialize their weights using VLM text embeddings of object and verb categories. To mitigate the gap between the VLM text embedding space and the feature space of the visual extractors, we do not utilize explorers directly as classification heads. Instead, these explorers compute the VL similarity between visual features and linguistic representations. This process can be formulated as follows:

$$\begin{aligned} \omega_{ins} &= \sigma \text{MLP}_{vlm}^{ins}(\hat{\mathbf{Z}}, \theta_{ins}), \\ \omega_{int} &= \sigma \text{MLP}_{vlm}^{int}(\hat{\mathbf{Z}}, \theta_{int}), \end{aligned} \quad (5)$$

where θ_{ins} and θ_{int} refer to the weights of the instance explorer and interaction explorer, while σ refers to a Gumbel softmax (Jang, Gu, and Poole 2017), which optimizes the tuning of explorers. ω_{ins} and ω_{int} refer to the category-related similarity of the visual content memory. Then, we select N_q pooled feature maps with the highest pooled similarities $\hat{\omega}_{ins}$ and $\hat{\omega}_{int}$ in a Topk manner. The selected features are used as online instance-aware guidance and interaction-aware guidance $\mathbf{f}_{ins} \in \mathbb{R}^{N_q \times C}$ and $\mathbf{f}_{int} \in \mathbb{R}^{N_q \times C}$. Such process is expressed by:

$$\begin{aligned} \mathbf{f}_{ins} &= \prod_{k=0}^{N_q} \text{Topk}_{\omega_{ins} \in N_o}(\hat{\mathbf{Z}}_{\omega_{ins}}^k), \\ \mathbf{f}_{int} &= \prod_{k=0}^{N_q} \text{Topk}_{\omega_{int} \in N_v}(\hat{\mathbf{Z}}_{\omega_c}^k), \end{aligned} \quad (6)$$

where \prod refers to concatenation along the query dimension. The query dimension is padded with zero if the number of predicted similarities N_o or N_v are less than N_q . To our best practice, we select OpenAI CLIP (Radford et al. 2021) as our semantic teacher for the explorers. As both instance and interaction-related semantic information help capture meaningful contexts, we concatenate \mathbf{f}_{ins} and \mathbf{f}_{int} in the query dimension to construct semantic context guidance \mathbf{f}_c . Finally, we add \mathbf{f}_c to context queries. Observing that such semantic guidance is also helpful for instance detection, we incorporate \mathbf{f}_{ins} with the instance queries, this process can be seen as an additional semantic-guided instance exploration for instance detection branch, shown in Figure 2. In practice, we employ two linear layers to map the semantic guidance to align with the query dimensions of the instance and context queries.

3.4 Context Aggregator

Given the captured instance detection feature \mathbf{Z}_{ins} and informative contexts \mathbf{Z}_c , a learnable aggregation query $\mathbf{Q}_{agg} \in$

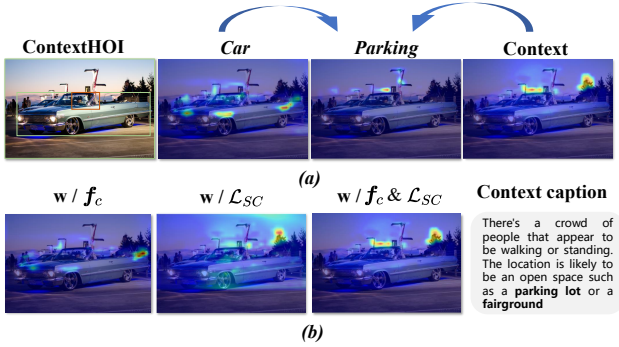


Figure 4: **Visualization analysis on spatial context learning.** (a) The feature map of the last layer of instance decoder, context aggregator and context extractor, indexed by the highest logits. Our instance decoder focuses on the appearance of the car, and the context extractor captures backgrounds and surrounding humans. (b) The features that are captured by context extractors with different component compositions. Both components help capture spatial contexts. Best viewed in color. Please zoom in for details.

$\mathbb{R}^{N_q \times C}$ is learned to fuse instance and context feature that complementary to each other. Inspired by the knowledge integration approach in (Ning et al. 2023), we implement the aggregator with a multi-branch transformer decoder with shared cross-attention weight. In each decoder layer, such process can be expressed by:

$$\begin{aligned} \hat{\mathbf{Z}}_{ins} &= \text{CrossAttn}(\mathbf{Q}_{agg}, \mathbf{Z}_{ins}, \theta_{attn}), \\ \hat{\mathbf{Z}}_c &= \text{CrossAttn}(\mathbf{Q}_{agg}, \mathbf{Z}_c, \theta_{attn}). \end{aligned} \quad (7)$$

We also calculate a cross-attention between the VLM visual feature and the aggregation query as an auxiliary semantic guidance, obtaining $\hat{\mathbf{Z}}_{vlm}$. After this step, we concat $\hat{\mathbf{Z}}_{ins}$, $\hat{\mathbf{Z}}_c$ and $\hat{\mathbf{Z}}_{vlm}$ as the aggregated feature for interaction prediction.

3.5 Training

Context-aware training. ContextHOI is trained end-to-end and predict HOIs with set prediction manner following (Kim et al. 2021). During training, we calculate the matching cost between all the predictions and ground truths with a bipartite graph, under the Hungarian Matching algorithm (Chen et al. 2021), and select matched predictions with the lowest matching costs. The matched predictions are combined to a set $\{\mathbf{B}_h, \mathbf{B}_o, \mathbf{C}_o, \mathbf{C}_{hoi}\}$, and the conventional HOI losses \mathcal{L}_{HOI} proposed by QPIC (Tamura et al. 2021) are calculated on them.

Additionally, for the spatially contrastive constraints proposed in section 3.2, the formulated losses are as follows:

$$\mathcal{L}_{SC} = \lambda_{fc} \mathcal{L}_{FC} + \lambda_{rc} \mathcal{L}_{RC} + \lambda_{ic} \mathcal{L}_{IC}. \quad (8)$$

During the training stage, all the parameters share gradient propagation with the following end-to-end form:

$$\mathcal{L} = \mathcal{L}_{HOI} + \mathcal{L}_{SC}. \quad (9)$$

Method	Backbone	HICO-DET (Default)			v-coco
		full	rare	non-rare	$AP_{role}^{\#1}$
iCAN (Gao et al. 2018)	R50	14.84	10.45	16.15	45.3
(Wang et al. 2019)	R50	16.24	11.16	17.75	47.3
DRG (Gao et al. 2020)	R50-F	24.53	19.47	26.04	51.0
HOTR (Kim et al. 2021)	R50	25.10	17.34	27.42	55.2
QPIC (Tamura et al. 2021)	R50	29.07	21.85	31.23	58.8
CPC (Park et al. 2022)	R50	29.63	23.14	31.57	63.1
FGAHOI (Ma et al. 2023)	Swin-t	29.94	22.24	32.24	60.5
RLIP (Yuan et al. 2022)	R50	32.84	26.85	34.63	61.9
GENVLKT (Liao et al. 2022)	R50	33.75	29.25	35.10	62.4
PViC (Zhang et al. 2023)	R50	34.69	32.14	35.45	62.8
HOICLIP (Ning et al. 2023)	R50	34.69	31.12	35.74	63.5
RLIPv2 (Yuan et al. 2023)	R50	35.38	29.61	37.10	65.9
LOGICHOI (Li et al. 2023)	R50	35.47	32.03	36.22	64.4
RmLR (Cao et al. 2023a)	R50	36.93	29.03	39.29	63.7
ViPLO (Part et al. 2023)	ViT-B	37.22	35.45	37.75	62.2
SCTC (Jiang et al. 2024)	R50	37.92	34.78	38.86	67.1
ADA-CM (Lei et al. 2023)	ViT-L	38.40	37.52	38.66	58.5
BCOM (Wang et al. 2024)	R50	39.34	39.90	39.17	65.8
UniHOI (Cao et al. 2023b)	R50	40.06	39.91	40.11	65.5
ContextHOI	R50	41.82	43.91	41.19	66.1
ContextHOI	R101	42.09	42.41	41.99	67.3

Table 1: **Performance comparison with state-of-the-art on benchmark HICO-DET.** R50 and R50-F refers to ResNet50 and ResNet50-FPN, Swin-t refers to Swin-tiny.

4 Experiments

4.1 Benchmarks and Metrics

Traditional benchmarks. We first conduct experiments on two widely-used HOI detection benchmark HICO-DET (Chao et al. 2018) and v-coco (Gupta and Malik 2015). HICO-DET comprises 80 object categories, 117 interaction categories, and 600 HOI triplet categories. HICO-DET contains 38,118 training images and 9,658 validation images. v-coco contains 5,400 train-val images and 4,946 validation images, consisting of 80 object categories, 29 verb categories, and 263 interaction triplet combinations.

Benchmark with ambiguous scenes. In this section, we introduce a specialized benchmark designed to evaluate the robustness of models when handling HOI samples with unclear instance attributes, termed HICO-DET (ambiguous). We select a subset from the test set of HICO-DET (Chao et al. 2018) that includes images with unclear visual attributes. Specifically, we engaged independent human volunteers to select images featuring unseen subjects, occluded subjects, blurred subjects, and instances too small to distinguish. A total of 659 images were selected, and together with their original annotations, they comprise the ambiguous benchmark.

Metrics. The mean Average Precision (mAP) is utilized for performance evaluation, following (Chao et al. 2018). For HICO-DET, mAP on both Default is reported. It has three different category settings, including *full* for all 600 HOI categories, *rare* for long-tail categories with less than 10 training samples, and the last *non-rare* categories. For v-coco, the role mAP in Scene 1 with 29 verb categories is reported.

Method	Context	HICO-DET (default)			HICO-DET (Ambiguous)		
		Full	Rare	Non-rare	Full	Rare	Non-rare
Two-stage HOI Detectors							
UPT (Zhang et al. 2022)	no context	31.66	25.94	33.36	16.53 (-15.13)	10.70 (-15.24)	18.27 (-15.09)
ADA-CM (Lei et al. 2023)	heuristic	38.40	37.52	38.66	18.12 (-20.28)	9.76 (-27.76)	20.62 (-18.94)
One-stage HOI Detectors							
QPIC (Tamura et al. 2021)	no context	29.07	21.85	31.23	9.25 (-19.82)	6.18 (-15.67)	9.61 (-21.62)
ContextHOI	ours	41.82	43.91	41.19	46.99 (+5.17)	60.57 (+16.66)	45.37 (+4.18)

Table 2: **Performance comparison with state-of-the-art on HICO-DET (default) and HICO-DET (ambiguous).** ContextHOI shows robustness on images with unclear instances.

4.2 Implementation Details

We implement ContextHOI with transformer detectors introduced by DETR (Carion et al. 2020), and both ResNet5 and ResNet101 backbone. Our detector query dimension N is 64 for both HICO-DET and v-coco. The encoders in the feature extractor have 6 layers, and the instance decode context extractor and context aggregator are all implemented by adopting a 3-layer transformer decoder. The transform hidden dimension C of all the components is 256.

We select CLIP ViT-L/14 (Radford et al. 2021) as our semantic teacher following (Lei et al. 2023), several single linear layers with LayerNorm is learned to match the size of transformer hidden dimension and CLIP visual linguistic dimensions. Our prediction heads follow the setting of (Ning et al. 2023). The simple MLP predicting context region boxes is a 3-layer MLP. The similarity predictors in the semantic-guided context explorer are implemented by 3-Layer MLP, we initialize their weights with the CLIP text embedding of category prompts. The learnable parameter τ in \mathcal{L}_{IC} is initialized to 0.5.

ContextHOI is trained for 60 epochs with an AdamW optimizer with an initial learning rate of $1e-4$ and a 10 times weight decay at 40 epochs. We train the model initialized with DETR (Carion et al. 2020) params pre-trained on MSCOCO. The Hungarian matching process and cost coefficients follow the setting of (Tamura et al. 2021). We remain the coefficient setting of conventional HOI prediction losses in QPIC (Tamura et al. 2021), for the spatial constraint losses, we set the loss coefficients λ_{fc} , λ_{rc} and λ_{ic} to 4, 1 and 4, respectively. ContextHOI is trained on a single Tesla A100 GPU with batch size 16.

4.3 Comparison with State-of-the-Art Methods

Regular HOI prediction. We provide performance comparisons with the state-of-the-art on regular settings of HICO-DET and v-coco. As detailed in Table 1, ContextHOI with a ResNet50 backbone has 41.82 *full* mAP and a higher 43.91 mAP on the *rare* split, achieves state-of-the-art. We hypothesize that the spatial context information helps to discover the long-tail HOIs. ContextHOI with a larger ResNet101 backbone has a higher 42.09 *full* mAP. As for v-coco, ContextHOI achieves 66.1 mAP on AP role Scene 1 with ResNet50 and 67.3 mAP with ResNet101. Both performance outperforms existing HOI detectors under the same backbones.



Figure 5: **Visualizations of the visual feature captured by ContextHOI from HICO-DET (ambiguous).** We mask the predicted instance boxes and let GPT-4V (OpenAI 2023) describe the left images; the words describing contexts in GPT captions are selected and shown as the yellow light texts.

Ambiguous benchmarks. We evaluate ContextHOI and several fundamental one-stage and two-stage detectors on HICO-DET (ambiguous). As demonstrated in Table 2, traditional one-stage and two-stage HOI detectors based on detection transformers show vulnerability in unclear and ambiguous scenes. The two-stage methods UPT (Zhang et al. 2022) and ADA-CM (Lei et al. 2023), which depend on pre-trained detection backbones, experience more performance decline than one-stage detectors, though they learn background information heuristically.

By integrating these techniques with our spatial context learning, our model further improves performance, showing a 5.17 increase in overall mAP and a 16.66 increase in rare mAP for ambiguous scenarios. This highlights the capability of our model to recognize interactions even with limited visual cues effectively. As ContextHOI also has a competitive performance on the regular HICO-DET test set, it indicates the spatial context learning do not excessively introduce irrelevant noise to the instance detection branch.

4.4 Effectiveness of Spatial Context Learning

Ablations of proposed modules. First, we provide performance ablation about our proposed spatial and semantic supervision in Table 3 (a). Directly equipping the base model (Ning et al. 2023) with an additional context branch can enhance the performance. The spatially contrastive constraints and the semantic-guided context exploration significantly further enhance the performance. Each type of supervision complements the other, with their combination yielding the highest mAP.

Visualizations of spatial context. We further validate the

(a)				(b)				(c)				(d)					
Context	\mathcal{L}_{SC}	SCE	Full	Rare	Non-rare	\mathcal{L}_{SC}	full	rare	non-rare	Prior	full	rare	non-rare	Aggregator	full	rare	non-rare
			34.69	31.12	35.75	None	38.91	39.73	38.66	None	36.75	35.47	37.12	w/ \mathbf{Z}_c	39.92	40.18	39.84
✓			35.81	33.27	36.87	w/ \mathcal{L}_{FC}	40.93	41.56	40.74	EVA-01	40.28	40.02	40.36	w/ $\mathbf{Z}_c + \mathbf{Z}_v$	41.82	43.91	41.19
✓	✓		36.75	35.47	37.12	w/ \mathcal{L}_{RC}	41.41	41.90	41.26	CLIP-L/14	41.82	43.91	41.19				
✓		✓	38.91	39.73	38.66	w/ \mathcal{L}_{IC}	40.59	40.77	40.54								
✓	✓	✓	41.82	43.91	41.19	w/ $\mathcal{L}_{IC} + \mathcal{W}_d$	41.46	42.84	41.05								
						all	41.82	43.91	41.19								

Table 3: **Ablations** of (a) Proposed context-aware supervision, SCE refers to Semantic-guided explorers, (b) components of spatially contrastive constraints, (c) utilizing prior knowledge from different VLMs, and (d) the implementation of the context aggregator. All experiments are conducted with ResNet50 backbone.

effectiveness of spatial context learning with visualizations, as illustrated in Figure 4. First, in (a), we demonstrate the captured features of the last layer of the instance decoder, the context aggregator and context extractor, presented from left to right. Evidently, each component consistently focuses on the intended feature regions, aligning well with their respective design objectives.

In (b), we evaluate the effectiveness of our designed supervision by comparing the heatmaps of the context extractor under different supervision settings. From left to right, the first heatmap shows that when the context extractor is guided solely by semantic supervision, it primarily focuses on the instance regions and duplicates the function of the instance decoder. With only spatial constraints, the context extractor focuses on regions close to the margin, which introduces unexpected noise. However, when both types of supervises are applied, the context extractor captures informative context regions, which are well-shaped.

Interpretability of spatial contexts. In the final block of Figure 4 (b), we present captions generated by GPT-4V (OpenAI 2023) that describe the spatial contexts depicted in the image. GPT-4V identifies the surrounding environment as either a parking lot or a fairground. These labels align with the context map shown in (a) and intuitively support the prediction of *parking*. This alignment highlights the implicit semantic meanings of the captured context regions.

4.5 Ablations

In this part, we provide detailed ablations on the model design and parameter efficiency of ContextHOI.

Spatially contrastive constraints. Table 3 (b) provides the ablation on the different settings of spatially contrastive constraints. All these constraint losses boost ContextHOI, while the naive \mathcal{L}_{IC} only provides light enhancement. With our proposed dynamic distance weight, \mathcal{L}_{IC} provides the highest enhancement. Combining all of these constraints, we obtain the best performance.

Different VLMs as knowledge teacher. Table 3 (c) compares different pre-trained VLMs utilized as our prior knowledge teacher. We try to train ContextHOI with no prior knowledge, with the visual-linguistic knowledge of EVA-CLIP-01 (Fang et al. 2022; Sun et al. 2023) and with CLIP ViT-L/14 (Radford et al. 2021). The OpenAI CLIP is the best VLM fitting our model.

Implementation of context aggregator. We evaluate different implementations of mechanisms integrating the instance and context featured in Table 3 (d). The term add denotes a gated addition approach between \mathbf{Z}_{ins} and \mathbf{Z}_c , enhancing feature integration by applying a learned gate to control the contribution from each feature set. The cross-attn term indicates shared cross-attention, as detailed in Section 3.4. In the last row, we further enhance the feature aggregation by incorporating visual features, \mathbf{Z}_v , obtained from a pre-trained VLM visual encoder into the shared cross-attention mechanism. The auxiliary semantic guidance provided by \mathbf{Z}_v also contributes to the robustness of the feature aggregation process.

Model efficiency. We evaluate model efficiency by conducting inference on the first 100 HICO-DET (Chao et al. 2018) test images with resolution 800×1333 , batch size 1, on a single A800 GPU. ContextHOI maintains an acceptable 176.43 GFLOPs and 13.37 FPS, compared with HOICLIP (Ning et al. 2023) with 159.93 GFLOPs and 16.47 FPS. Our context-aware supervision, as training losses, contributes a significant performance increase with light additional parameters, and it is not required during inference.

More visualizations. We show additional visualizations of the fused feature map for interaction prediction of ContextHOI in Figure 5. While recognizing images with tiny objects, unclear subjects and occluded subjects, ContextHOI tends to focus on spatial backgrounds containing informative environments. For instance, the sand and cloud for *fly a kite*, wave for *hold a surfboard*, and the dirt road for *drive a car*.

5 Conclusion & Discussion

We introduce a novel spatial context learning paradigm tailored for transformer-based HOI detectors, featuring a dual-branch and fusion architecture. This framework is augmented with novel spatially contrastive constraints and a semantic-guided context explorer. The proposed components combine our network ContextHOI, achieving a new state-of-the-art performance on HICO-DET with ResNet50. ContextHOI demonstrates significant robustness in HOI scenarios with unclear and occluded instance clues. In the future, we plan to explore more fine-grained context-learning approaches to mitigate background noise for robust HOI detection.

Acknowledgements

This work was supported by the National Natural Science Foundation of China under Grant 62172021 and the Guangdong Provincial Key Laboratory of Ultra High Definition Immersive Media Technology under Grant 2024B1212010006.

References

- An, B.; Zhu, S.; Panaitescu-Liess, M.-A.; Mummadi, C. K.; and Huang, F. 2023. More Context, Less Distraction: Visual Classification by Inferring and Conditioning on Contextual Attributes. arXiv:2308.01313.
- Cao, S.; Yin, Y.; Huang, L.; Liu, Y.; Zhao, X.; Zhao, D.; and Huang, K. 2023a. Efficient-VQGAN: Towards High-Resolution Image Generation with Efficient Vision Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 7368–7377.
- Cao, Y.; Tang, Q.; Su, X.; Chen, S.; You, S.; Lu, X.; and Xu, C. 2023b. Detecting Any Human-Object Interaction Relationship: Universal HOI Detector with Spatial Prompt Learning on Foundation Models. In *Advances in Neural Information Processing Systems*, volume 36, 739–751.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-End Object Detection with Transformers. In *Computer Vision – ECCV 2020*, 213–229.
- Chao, Y.-W.; Liu, Y.; Liu, X.; Zeng, H.; and Deng, J. 2018. Learning to Detect Human-Object Interactions. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 381–389.
- Chen, M.; Liao, Y.; Liu, S.; Chen, Z.; Wang, F.; and Qian, C. 2021. Reformulating HOI Detection As Adaptive Set Prediction. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, 9004–9013.
- Chen, Z.; Huang, S.; and Tao, D. 2018. Context Refinement for Object Detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Fang, Y.; Wang, W.; Xie, B.; Sun, Q.; Wu, L.; Wang, X.; Huang, T.; Wang, X.; and Cao, Y. 2022. EVA: Exploring the Limits of Masked Visual Representation Learning at Scale. *arXiv preprint arXiv:2211.07636*.
- Frederic Z. Zhang, D. C.; and Gould, S. 2021. Spatially Conditioned Graphs for Detecting Human-Object Interactions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 13319–13327.
- Gao, C.; Xu, J.; Zou, Y.; and Huang, J.-B. 2020. DRG: Dual Relation Graph for Human-Object Interaction Detection. In *European Conference on Computer Vision*.
- Gao et al. (2018) Chen, G.; Yuliang, Z.; and Huang, J. 2018. iCAN: Instance-Centric Attention Network for Human-Object Interaction Detection. In *BMVC2018*.
- Gupta, S.; and Malik, J. 2015. Visual Semantic Role Labeling. arXiv:1505.04474.
- Jang, E.; Gu, S.; and Poole, B. 2017. Categorical Reparameterization with Gumbel-Softmax. arXiv:1611.01144.
- Jiang, W.; Ren, W.; Tian, J.; Qu, L.; Wang, Z.; and Liu, H. 2024. Exploring Self- and Cross-Triplet Correlations for Human-Object Interaction Detection. arXiv:2401.05676.
- Kim, B.; Lee, J.; Kang, J.; Kim, E.; and Kim, H. J. 2021. HOTR: End-to-End Human-Object Interaction Detection With Transformers. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, 74–83.
- Kim, S.; Jung, D.; and Cho, M. 2023. Relational Context Learning for Human-Object Interaction Detection. arXiv:2304.04997.
- Lei, T.; Caba, F.; Chen, Q.; Jin, H.; Peng, Y.; and Liu, Y. 2023. Efficient Adaptive Human-Object Interaction Detection with Concept-guided Memory. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 6480–6490.
- Li, L.; Wei, J.; Wang, W.; and Yang, Y. 2023. Neural-Logic Human-Object Interaction Detection. In *Advances in Neural Information Processing Systems*, volume 36, 21158–21171.
- Liao, Y.; Zhang, A.; Lu, M.; Wang, Y.; Li, X.; and Liu, S. 2022. GEN-VLKT: Simplify Association and Enhance Interaction Understanding for HOI Detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, 20091–20100.
- Ma, S.; Wang, Y.; Wang, S.; and Wei, Y. 2023. FGAHOI: Fine-Grained Anchors for Human-Object Interaction Detection. arXiv:2301.04019.
- Ning, S.; Qiu, L.; Liu, Y.; and He, X. 2023. HOICLIP: Efficient Knowledge Transfer for HOI Detection with Vision-Language Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023*, 23507–23517.
- OpenAI. 2023. GPT-4: Enhancements and Capabilities. <https://openai.com/blog/gpt-4>. Accessed: 2024-08-01.
- Park, J.; Lee, S.; Heo, H.; Choi, H. K.; and Kim, H. J. 2022. Consistency Learning via Decoding Path Augmentation for Transformers in Human Object Interaction Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022*.
- Part et al. (2023) Park, J.; Park, J.-W.; and Lee, J.-S. 2023. ViPLO: Vision Transformer Based Pose-Conditioned Self-Loop Graph for Human-Object Interaction Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 17152–17162.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139, 8748–8763.
- Rezatofighi, H.; Tsoi, N.; Gwak, J.; Sadeghian, A.; Reid, I. D.; and Savarese, S. 2019. Generalized Intersection Over Union: A Metric and a Loss for Bounding Box Regression. In *IEEE Conference on Computer Vision and Pattern*

Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019, 658–666.

Sun, Q.; Fang, Y.; Wu, L.; Wang, X.; and Cao, Y. 2023. EVA-CLIP: Improved Training Techniques for CLIP at Scale. *arXiv preprint arXiv:2303.15389*.

Sun, S.; Chen, Y.; Zhu, Y.; Guo, G.; and Li, G. 2022. Sk-flow: Learning optical flow with super kernels. *Advances in Neural Information Processing Systems*, 35: 11313–11326.

Tamura et al.(2021) Masato, T.; Hiroki, O.; and Tomoaki, Y. 2021. QPIC: Query-Based Pairwise Human-Object Interaction Detection With Image-Wide Contextual Information. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, 10410–10419.

Wang, G.; Guo, Y.; Xu, Z.; and Kankanhalli, M. 2024. Bilateral Adaptation for Human-Object Interaction Detection with Occlusion-Robustness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 27970–27980.

Wang, T.; Anwer, R. M.; Khan, M. H.; Khan, F. S.; Pang, Y.; Shao, L.; and Laaksonen, J. 2019. Deep Contextual Attention for Human-Object Interaction Detection. In *ICCV*.

Yang, J.; Lu, J.; Lee, S.; Batra, D.; and Parikh, D. 2018. Graph r-cnn for scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 670–685.

Yuan, H.; Jiang, J.; Albanie, S.; Feng, T.; Huang, Z.; Ni, D.; and Tang, M. 2022. RLIP: Relational Language-Image Pre-training for Human-Object Interaction Detection. In *Advances in Neural Information Processing Systems*, volume 35, 37416–37431.

Yuan, H.; and Ni, D. 2021. Learning Visual Context for Group Activity Recognition. *AAAI*, 35.

Yuan, H.; Zhang, S.; Wang, X.; Albanie, S.; Pan, Y.; Feng, T.; Jiang, J.; Ni, D.; Zhang, Y.; and Zhao, D. 2023. RLIPv2: Fast Scaling of Relational Language-Image Pre-Training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 21649–21661.

Zhai, Y.; Liu, Z.; Wu, Z.; Wu, Y.; Zhou, C.; Doermann, D.; Yuan, J.; and Hua, G. 2023. SOAR: Scene-debiasing Open-set Action Recognition. *arXiv:2309.01265*.

Zhang, A.; Liao, Y.; Liu, S.; Lu, M.; Wang, Y.; Gao, C.; and Li, X. 2021. Mining the Benefits of Two-stage and One-stage HOI Detection. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 17209–17220.

Zhang, F. Z.; Yuan, Y.; Campbell, D.; Zhong, Z.; and Gould, S. 2023. Exploring Predicate Visual Context in Detecting of Human-Object Interactions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 10411–10421.

Zhang et al. (2022) Zhang, F. Z.; Campbell, D.; and Gould, S. 2022. Efficient Two-Stage Detection of Human-Object Interactions with a Novel Unary-Pairwise Transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 20104–20112.