

Dual-Glance Model for Deciphering Social Relationships

Junnan Li

Graduate School for Integrative Sciences and Engineering
 National University of Singapore
 Singapore

lijunnan@u.nus.edu

Qi Zhao

Department of Computer Science and Engineering
 University of Minnesota
 Minneapolis, USA

qzhao@cs.umn.edu

Yongkang Wong

Interactive & Digital Media Institute
 National University of Singapore
 Singapore

yongkang.wong@nus.edu.sg

Mohan S. Kankanhalli

School of Computing
 National University of Singapore
 Singapore

mohan@comp.nus.edu.sg

Abstract

Since the beginning of early civilizations, social relationships derived from each individual fundamentally form the basis of social structure in our daily life. In the computer vision literature, much progress has been made in scene understanding, such as object detection and scene parsing. Recent research focuses on the relationship between objects based on its functionality and geometrical relations. In this work, we aim to study the problem of social relationship recognition, in still images. We have proposed a dual-glance model for social relationship recognition, where the first glance fixates at the individual pair of interest and the second glance deploys attention mechanism to explore contextual cues. We have also collected a new large scale People in Social Context (PISC) dataset, which comprises of 22,670 images and 76,568 annotated samples from 9 types of social relationship. We provide benchmark results on the PISC dataset, and qualitatively demonstrate the efficacy of the proposed model.

1. Introduction

Social relationships derived from each individual fundamentally form the basis of social structure in our daily life. Naturally, we perceive and interpret a scene with an understanding of the social relationships of the people in the scene. Sociology research shows that such social understanding of people permits inference about their characteristics and their possible behaviors [32].

In computer vision, social information has been exploited to improve several analytics tasks, including human



Figure 1: Example images from the new People in Social Context (PISC) dataset.

trajectory prediction [1, 30], multi-target tracking [4, 26], and group activity recognition [9, 22, 23]. In image understanding task, visual concepts recognition is gaining more attention, which include visual attribute [19] and visual relationship [25]. On the other hand, social attribute and social relationship [35] are equally important concepts for scene understanding, but have received less attention in the research community. In this work, we aim to address the problem of social relationship recognition. Understanding such relationship can enable a well designed algorithm to gen-

erate better descriptions for a scene. For instance, the first image in Figure 1 can be described as ‘Grandma is holding her grandchild’, rather than ‘A person is holding a baby’.

With reference to the *relational models* theory [12], we define a hierarchical social relationship categories which embed the coarse-to-fine characteristic of common social relationships (as illustrated in Figure 2). Our definition follows a prototype-based approach, where we are interested in finding exemplars that most parsimoniously describe the most common situations, rather than an ambiguous definition that could cover all possible cases. The presented recognition problem differs from the visual relationship detection problem in [25]. We argue that inferring social relationship requires a higher level of understanding about the scene. This is because humans make such inferences not only based on the physical appearance (*e.g.*, color of clothes, gender, age, *etc.*), but also from subtler cues (*e.g.*, expression, action, proximity, and context) [2, 27, 42].

Recognizing social relationships from still images is challenging due to the wide variations in scale, scene, pose, and appearance. In this work, we propose a dual-glance model, which exploits information from a target individual pair as well as the surrounding contextual cues. The key contributions can be summarized as:

- The proposed dual-glance model mimics human visual system to explore useful and complementary visual cue for social relationship analysis. The first glance fixates at the individual pair of interest, and performs coarse prediction based on its appearance and geometrical information. The second glance exploits contextual cues from regions generated from Region Proposal Network (RPN) [28] to refine the coarse prediction.
- We propose Attentive RCNN, where attention is allocated for each contextual region. The attention mechanism is guided by both bottom-up and top-down signals. Better performance is achieved by selectively focusing on the relevant regions.
- To enable this study, we collected a novel People in Social Context (PISC) dataset¹. It consists of 22,670 images and 76,568 manually annotated labels from 9 types of social relationship. In addition, PISC also consists of 66 annotated occupation categories. To the best of our knowledge, PISC is the first public dataset for social relationship analysis.

The remaining of the paper is organized as follows. Section 2 reviews the related work. Section 3 delineates the details of the new PISC dataset. Section 4 elaborates on the details of the proposed framework, and the empirical evaluation is shown in Section 5. Section 6 concludes the paper.

¹<https://doi.org/10.5281/zenodo.832013>

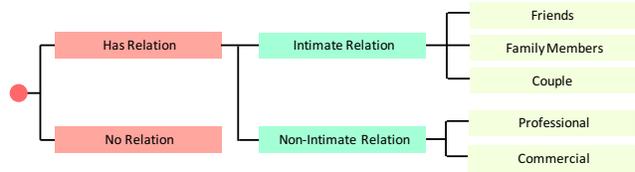


Figure 2: Defined hierarchical social relationship categories.

2. Related Work

2.1. Social Relationship

The study of social relationships lies at the heart of social sciences. There are two forms of representations for relational cognition. The first approach represents relationship with a set of theorized or empirically derived dimensions [5]. The other form of representation proposes implicit categories for relation cognition [17]. One of the most extensively accepted categorical theory is the relational models theory [12]. It offers a unified account of social relations by proposing four elementary prototypes, namely *communal sharing*, *equality matching*, *authority ranking*, and *market pricing*.

In the computer vision literature, social information has been widely adopted as supplementary cues to other tasks. Gallagher *et al.* [13] extract features describing group structure to aid demographic recognition. For group activity recognition, social roles and relationship information have been implicitly embedded into the inference model [4, 6, 9, 22, 23]. Alletto *et al.* [2] define ‘social pairwise feature’ based on F-formation and use it for group detection in egocentric videos. Recently, [1, 30] model social factor for human trajectory prediction.

There have been studies that explicitly focus on recognition of social attributes and social structures. Wang *et al.* [35] first study familial social relationship recognition in personal image collections. Kinship verification [7, 11, 37] and kinship recognition [3, 15] have been extensively studied. Zhang *et al.* [42] study facial traits (*e.g.*, friendly, dominant, *etc.*) that are informative of social relationships. For video based analysis, Ding and Yilmaz discover social communities formed by actors in movies [8]. Ramanathan *et al.* [27] study weakly supervised social role discovery in events.

Our study partially overlaps with the field of social signal processing [34], which aims to understand social signals and social behaviors using multiple sensors, such as role recognition, influence ranking, and dominance detection in group meeting [20, 29, 31]. Our work substantially differs from the aforementioned studies. Unlike facial attributes based social relationship study [3, 15, 35, 42], we study people in complex daily scenes with uncontrolled poses and orientations. Furthermore, we focus on general social relationships, rather than kinship in family photos [3, 7, 11, 35, 37].

Table 1: Instructions provided to annotators.

Relationship	Description	Examples
Professional	The people are related based on their professions	co-worker; coach & player; boss & staff
Commercial	One person is paying money to receive goods/service from the other	salesman & customer; tour guide & tourist



Figure 3: Example of social relationship labels that are not agreed among annotators.

Different from video-based studies [8, 27], we focus on visual information from a single image.

2.2. Multiple-Instance Learning

The proposed Attentive RCNN is inspired by Multiple-Instance Learning (MIL). MIL is a weakly-supervised learning approach which trains a classifier with bags of instances and bag-level labels. Recently, researchers explored MIL with deep feature representations. Wu *et al.* [36] propose a deep MIL framework to exploit correspondences between keywords and image regions for image classification and annotation, while a similar technique was adopted to detect salient concepts for image captions generation [10]. Inspired by MIL, Gkioxari *et al.* [14] propose R*CNN. Different from previous approaches, it localizes target region for action recognition by exploiting complementary representative cue from a set of candidate regions in an image.

Attention model has been recently proposed and applied to image captioning [39, 41], image question answering [40] and fine-grained classification [38]. We modify R*CNN with attention mechanism to better exploit contextual cues. We treat the attention weights for the contextual regions as latent variable, which can be inferred with a forward pass of the model.

3. People in Social Context Dataset

The People in Social Context (PISC) dataset is the first of its kind that focuses on social relationships. It was collected through a pipeline of three stages. In the first stage, we collected around 40k images containing people from a variety of sources, including Visual Genome [21], MS-COCO [24], YFCC100M [33], Flickr, Instagram, Twitter and commercial search engines (*i.e.* Google and Bing). We used a combination of key words search (*i.e.* co-worker, people, friends, *etc.*) and people detector (Faster RCNN [28]) to collect the image. The collected images have high variation in image resolution, people’s appearance, and scene type.

In the second and third stage, we hired workers from CrowdFlower platform to perform labor intensive task of

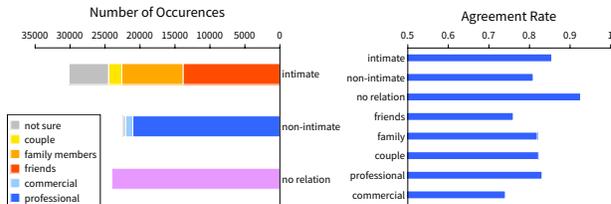


Figure 4: Annotation statistics of the relationship categories.

manual annotation. The second stage focused on the annotation of person bounding box in each image. Following [21], each bounding box is required to strictly satisfy the coverage and quality requirements. To speed up the annotation process, we first deployed Faster RCNN to detect people on all images, followed by asking the annotators to re-annotate the bounding boxes if the computer-generated bounding boxes were inaccurately localized. Overall, 40% of the computer-generated boxes are kept. For images collected from MSCOCO and Visual Genome, we directly used the provided groundtruth bounding boxes.

Once the bounding boxes of all images had been annotated, we selected images consisting of at least two people who occupy a significant amount of region, and avoided images that contain crowds of people where individuals cannot be distinguished. In the final stage, we requested the annotators to identify the occupation of all individuals in the image, as well as the social relationships of all potential individual pairs. To ensure consistency in the occupation categories, the annotation is based on a list of reference occupation categories. The annotators could manually add a new occupation category if it was not in the list.

For social relationships, we formulate the annotation task as multi-choice questions based on the hierarchical structure in Figure 2. We provide instructions (see Table 1) to help the annotators distinguish between professional and commercial relationship. Annotators can choose the option ‘not sure’ at any level if they cannot confidently identify the relationship. Each image was annotated by five workers, and the final decision is determined by majority voting. If the five workers do not reach an agreement (*e.g.* 2-2-1), the annotation will be treated as invalid (see Figure 3). Overall, 7,928 unique workers have contributed to the annotation.

The PISC dataset consists of 22,670 images. The average number of people per image is 3.11. For the social relationships, we consider each individual pair as one sample. In total, we collected 76,568 valid samples. The distribution for each types of relationships and their agreement rate is shown in Figure 4. The agreement rate is calculated by dividing the number of correct human judgments (judgments that agree with the majority) with the total number of judgments. For occupations, 10,034 images contain people that have recognizable occupations. In total, there are 66 identified occupation categories. The number of occupation occurrence and the workers’ agreement rate for the

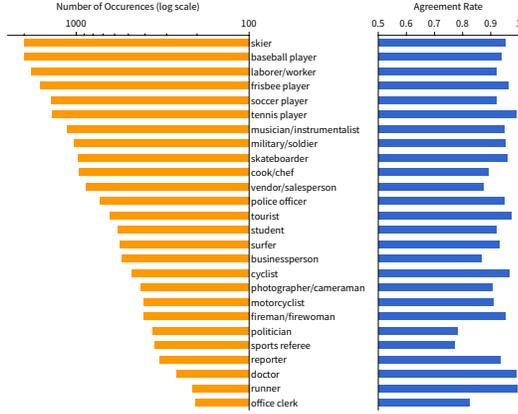


Figure 5: Annotation statistics of the top 26 occupations.

26 most frequent occupation categories are shown in Figure 5. A lower agreement rate indicates that the occupation is harder to visually discriminate (e.g. ‘politician’ and ‘office clerk’). Since two source datasets, *i.e.* MS-COCO and Visual Genome, are highly biased towards ‘baseball player’ and ‘skier’, we limit the total number of instances per occupation to 2000 based on agreement rate ranking to ensure there are no bias towards any particular occupation.

4. Proposed Dual-Glance Model

Given an image I and a target pair of people highlighted by bounding boxes $\{b_1, b_2\}$, our goal is to infer their social relationship r . In this work, we propose a dual-glance relationship recognition model, where the first glance fixates at b_1 and b_2 , and the second glance explores contextual cues from multiple region proposals \mathbf{R} . The final score over possible relationships, \mathbf{S} , is a weighted sum of the two scores via

$$\mathbf{S} = \mathbf{S}_1(I, b_1, b_2) + \alpha \mathbf{S}_2(I, b_1, b_2, \mathbf{R}). \quad (1)$$

We use softmax to transform the final score into a probability distribution. Specifically, the probability that a given pair of people having relationship r is given as

$$p_r = \frac{\exp(S_r)}{\sum_r \exp(S_r)}. \quad (2)$$

An overview of the proposed model is shown in Figure 6.

4.1. First Glance

The first glance takes in input I and two bounding boxes. We first crop three patches from I , where the first two cover each person, p_1 and p_2 , and one for the union region, p_U , that tightly covers both people. These patches are resized to 224×224 pixels and fed into three CNNs. The outputs from the last convolutional layer are flattened and concatenated. p_1 and p_2 are processed by CNNs that share the same weights.

We denote the geometry feature of the bounding box i as $\mathbf{b}_i^{\text{pos}} = \{x_i^{\text{min}}, y_i^{\text{min}}, x_i^{\text{max}}, y_i^{\text{max}}, \text{area}_i\} \in \mathbb{R}^5$, where all the parameters are relative values, normalized with zero mean and unit variance. $\mathbf{b}_1^{\text{pos}}$ and $\mathbf{b}_2^{\text{pos}}$ are concatenated and processed by a fully-connected (fc) layer. We concatenate its output with the CNN features for p_1 , p_2 and p_U to form a single feature vector, which is subsequently passed through another two fc layers to produce first glance score, s_1 . We use $\mathbf{v}_{\text{top}} \in \mathbb{R}^k$ to denote the output from the penultimate fc layer. \mathbf{v}_{top} serves as a top-down signal to guide the attention mechanism in the second glance. We experimented with different values of k , and set k as 4096.

4.2. Attentive RCNN for Second Glance

For the second glance, we adapt Faster RCNN [28] to make use of multiple contextual regions. Faster RCNN processes the input image I with Region Proposal Network (RPN) to generate a set of region proposals \mathbf{P}_I with high objectness. For each target pair with bounding boxes b_1 and b_2 , we select the set of contextual regions $\mathbf{R}(b_1, b_2; I)$ from \mathbf{P}_I as

$$\mathbf{R}(b_1, b_2; I) = \{c \in \mathbf{P}_I : \max(G(c, b_1), G(c, b_2)) < \tau_u\} \quad (3)$$

where $G(b_1, b_2)$ computes the Intersection-over-Union (IoU) between two regions, and τ_u is the upper threshold for IoU. The threshold encourages the second glance to explore cues different from the first glance. It’s effect is reported in Section 5.4.

We then process I with a CNN to generate a convolutional feature map $\text{conv}(I)$. For each contextual region $c \in \mathbf{R}$, ROI pooling is applied to extract a fixed-length feature vector $\mathbf{v} \in \mathbb{R}^k$ from $\text{conv}(I)$. Denote $\{\mathbf{v}_i | i = 1, 2, \dots, N\}$ as the bag of N feature vectors for \mathbf{R} , also given the high-level feature vector from the first glance \mathbf{v}_{top} , we first combine them into a hidden vector $\mathbf{h}_i \in \mathbb{R}^k$ via

$$\mathbf{h}_i = \mathbf{v}_i + \mathbf{w}_{\text{top}} \otimes \mathbf{v}_{\text{top}}, \quad (4)$$

where $\mathbf{w}_{\text{top}} \in \mathbb{R}^k$, and \otimes is the element-wise multiplication of two vectors. Then, we calculate the attention $a_i \in [0, 1]$ over the i th region proposal as

$$a_i = \frac{1}{1 + \exp(-(\mathbf{W}_{h,a} \mathbf{h}_i + b_a))}, \quad (5)$$

where $\mathbf{W}_{h,a} \in \mathbb{R}^{1 \times k}$ is the weight matrix, and $b_a \in \mathbb{R}$ is the bias term. The attention over each contextual region is guided by both bottom-up signal from local region \mathbf{v}_i and top-down signal from the first glance model \mathbf{v}_{top} . Hence, the weighted feature vector for region i is computed via

$$\mathbf{v}_i^{\text{att}} = a_i \mathbf{v}_i. \quad (6)$$

The obtained $\mathbf{v}_i^{\text{att}}$ is processed by the last fc layer to generate the output score for the i th region proposal

$$\mathbf{s}_i = \mathbf{W}_s \mathbf{v}_i^{\text{att}} + \mathbf{b}_s. \quad (7)$$

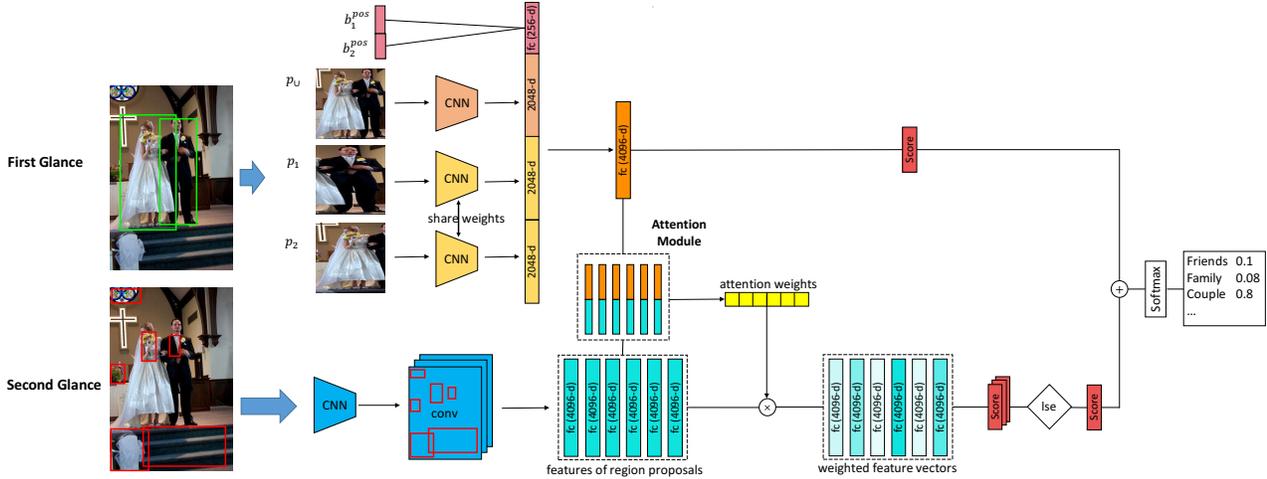


Figure 6: An overview of the proposed dual-glance model. The first glance looks at the pair of people in question and makes a coarse prediction. The second glance looks at region proposals, allocates attention to each region, and aggregates their outputs to refine the score. The attention is guided by both top-down signal from the first glance, and bottom-up signal from the local context.

Functions to aggregate scores, *i.e.*, $f(s_i)$, from a bag of instances $i = 1, 2, \dots, N$ include $\max(s_i)$, $\text{avg}(s_i)$, and $\log[1 + \sum_{i=1}^N \exp(s_i)]$ (log-sum-exp, denoted as $\text{lse}(s_i)$). In our experiment (Section 5.3), we evaluated all three variants of $f(s_i)$ and the results show that $\text{lse}(s_i)$ provides the best performance. Hence, the score of the second glance model is

$$S_2 = \log[1 + \sum_{i=1}^N \exp(s_i)]. \quad (8)$$

5. Experiment

5.1. Dataset and Training Details

In this work, we conducted experiments on the proposed PISC dataset (Section 3) and evaluated our proposed method with two recognition tasks. The first task, denoted as 3-relationship recognition, focuses on three coarse-level relationship categories, namely *No Relation*, *Intimate Relation*, and *Non-Intimate Relation*. We randomly select 4,000 images (14,852 samples) as test set, and use the remaining images as training set. The second task, denoted as 6-relationship recognition, focuses on finer relationships listed in Figure 2. Since the data label is unbalanced (fewer images with *Couple* or *Commercial relationship*), we split 1,500 images (3,517 samples) into test set and ensure it contains around 600 samples for each of the six relationships.

The relationship imbalance reflects their frequency of occurrence, which is also observed in [25]. To address this, we adopt oversampling and undersampling strategies. Specifically, we oversample the minority labeled samples by reversing the pair of people (*i.e.*, if p_1 and p_2 are a couple, then p_2 and p_1 are also a couple), and by horizontally flipping the image. We undersample the majority labeled samples using stratified sampling scheme to ensure the sam-

ples in each batch is balanced.

In this work, we train our model with Stochastic Gradient Descent using backpropagation. First, we train the first-glance model until the loss converges, then we freeze the first-glance model, and train the second-glance model. For the first glance model, we fine-tune the ResNet-101 model pre-trained on ImageNet classification task [18]. For the second glance model, we fine-tune the VGG-16 model pre-trained on ImageNet detection task [28]. We set the learning rate as 0.001, while the fine-tuning model has a lower learning rate of 0.0001. We use a batch size of 32 and a momentum of 0.9 during training.

During the test stage, we found that the performance would slightly improve if we feed the model twice with $\{b_1, b_2\}$ and $\{b_2, b_1\}$, and take their average as the final score. However, the performance gain (0.5%-1%) doubles the time budget, and we do not recommend it in practice.

5.2. Single-Glance vs. Dual-Glance

As there exists limited literature on this problem, we evaluate multiple variants of our model as baseline and compare them to the proposed dual-glance mode to show its efficacy. Formally, the compared methods are as followed:

1. **Union-CNN**: Following the predicate prediction model in [25], a single CNN model is used to classify the union region of the individual pair of interest.
2. **BBox**: We only use the geometry feature of the two bounding boxes to infer the relationship.
3. **Pair-CNN**: The model consists of two CNNs with shared weights. The input is the cropped image patches for the two individuals.

Table 2: Recall-per-class and mean average precision (mAP) of baselines and our proposed dual-glance model on the PISC dataset.

	3-relationship				6-relationship						
	Intimate	Non-Intimate	No Relation	mAP	Friends	Family	Couple	Professional	Commercial	No Relation	mAP
Union-CNN [25]	72.1	81.8	19.2	58.4	29.9	58.5	70.7	55.4	43.0	19.6	43.5
BBox	42.4	33.0	41.9	34.9	20.7	36.4	42.7	31.8	23.2	32.7	28.8
Pair-CNN	70.3	80.5	38.8	65.1	30.2	59.1	69.4	57.5	41.9	34.2	48.2
Pair-CNN+BBox	71.8	80.3	50.6	69.6	30.7	60.2	72.5	58.1	43.7	50.7	54.3
Pair-CNN+BBox+Union	71.1	81.2	57.9	72.2	32.5	62.1	73.9	61.4	46.0	52.1	56.9
Pair-CNN+BBox+Global	70.5	80.9	53.7	70.5	32.2	61.7	72.6	60.8	44.3	51.0	54.6
Pair-CNN+BBox+Scene	71.0	80.6	46.7	68.0	30.2	59.4	71.7	57.6	43.0	49.9	51.7
RCNN	72.9	83.3	14.8	63.5	29.7	61.9	71.2	60.1	45.9	20.7	48.4
Dual-Glance	73.1	84.2	59.6	79.7	35.4	68.1	76.3	70.3	57.6	60.9	63.2

- Pair-CNN+BBox**: We extend Pair-CNN by using the geometry feature of the two bounding boxes.
- Pair-CNN+BBox+Union**: The first glance model as illustrated in Figure 6, which combines Pair-CNN+BBox and Union-CNN.
- Pair-CNN+BBox+Global**: Instead of the union region, we use the entire image as input to Union-CNN.
- Pair-CNN+BBox+Scene**: The Union-CNN is replaced with a Scene-CNN pre-trained on Places [43]. It extracts scene information using the entire image.
- RCNN**: We train a RCNN using the region proposals P_I , and adopt average pooling to combine the features.
- Dual-Glance**: Our proposed model (Section 4).

Table 2 shows the results on the test set for both the 3-relationship recognition task and the 6-relationship recognition task. **Union-CNN** and **RCNN**, are incapable to recognize *No Relation*. This is because these model don't know the pair of people in question, and would recognize other salient relationships. **Pair-CNN+BBox** outperforms **Pair-CNN**, which suggests that peoples' geometric position in an image contains information useful to infer their relationship, especially for *No Relation*. This is supported by the law of *proxemics* defined in the book "The Silent Language" [16]. However, the position of bounding boxes alone cannot be used to predict relationship, as shown by the results of **BBox**.

Adding **Union-CNN** to **Pair-CNN+BBox** improves performance. However, the performance gain is slight if we use the global context (entire image) rather than local context (union region). Furthermore, the performance even degrades when the global context incorporates scene information, suggesting that social relationships are independent of

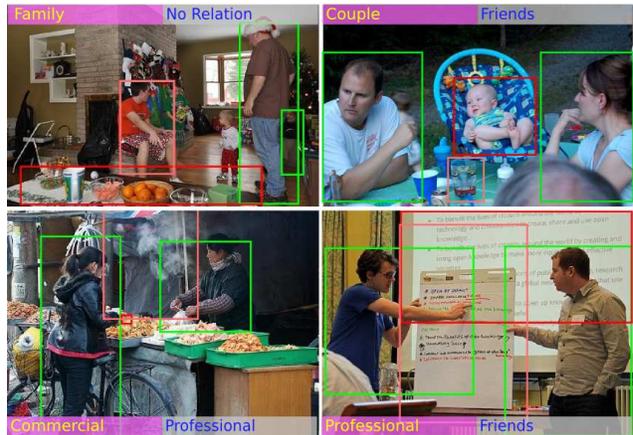


Figure 7: Examples where dual-glance model correctly predict the relationship (yellow label) while the first-glance model fails (blue label). GREEN boxes highlight the pair of people in question, and the top two contextual regions with highest attention are highlighted in RED.

scene types. **RCNN** demonstrates the effectiveness of using contextual regions, particularly for *Intimate Relation* and *Non-Intimate Relation*.

The proposed **Dual-Glance** significantly outperforms all baseline models. Figure 7 shows some intuitive illustrations where proposed model correctly classifies relationships misclassified by the first-glance (**Pair-CNN+BBox+Union**) model.

Across all models, *Friends* and *Commercial* are more difficult to recognize. This is consistent with the agreement rate in Figure 4, which indicates that *Friends* and *Commercial* are less visually distinguishable. Figure 8 shows the confusion matrix of 6-relationship recognition task. The three intimate relationships (*Friends*, *Family*, *Couple*) are more often to be confused with each other than with non-intimate relationships, suggesting they share similar visual

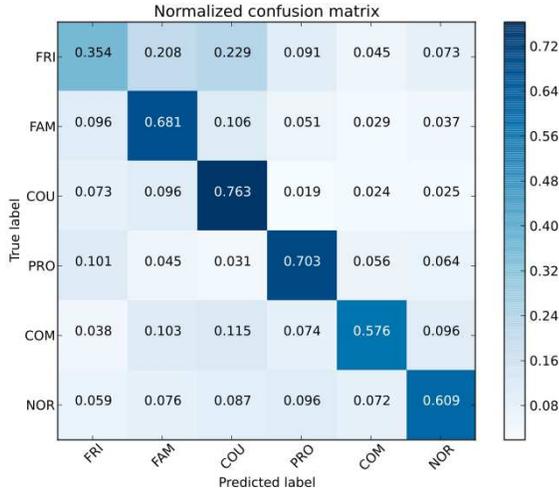


Figure 8: Confusion matrix of 6-relationship recognition task with the proposed dual-glance model.

Table 3: mAP (%) of the proposed dual-glance model with and without attention mechanism using various aggregation functions.

	Without Attention			With Attention		
	avg(\cdot)	lse(\cdot)	max(\cdot)	avg(\cdot)	lse(\cdot)	max(\cdot)
3-relationship	71.7	73.0	74.8	76.9	79.7	77.6
6-relationship	55.9	57.5	58.2	61.8	63.2	62.1

features. However, the non-intimate relationships (*Professional*, *Commercial*) do not tend to be easily confused with each other.

5.3. Analysis on Attention Mechanism

Here, we remove the attention module and compare it with our proposed dual-glance model. For the second glance, we experiment with three widely used aggregation functions $f(\cdot)$, which are $\text{avg}(\cdot)$, $\text{lse}(\cdot)$ and $\text{max}(\cdot)$. The results are shown in Table 3. Adding attention mechanism improves performance for all three aggregation functions. For Dual-glance without attention, $\text{max}(\cdot)$ performs best, which conforms to the results in [14, 36]. While for Dual-glance with attention, $\text{lse}(\cdot)$ performs best.

The reason is that $\text{max}(\cdot)$ uses a single instance to infer the label for an entire bag. It works well in the presence of a ‘strong’ instance, but sometimes there is no strong instance, but several ‘weak’ instances. On the other hand, $\text{lse}(\cdot)$ and $\text{avg}(\cdot)$ consider all instances in a bag, but could be distracted by irrelevant instances. However, with properly guided attention, $\text{lse}(\cdot)$ and $\text{avg}(\cdot)$ can better exploit the collaborative power of relevant instances for more accurate inference.

5.4. Variations of Contextual Regions

Since RPN can generate hundreds of region proposals per image, we suppress those proposals with non-maximum suppression (NMP). We vary m , the maximum number of

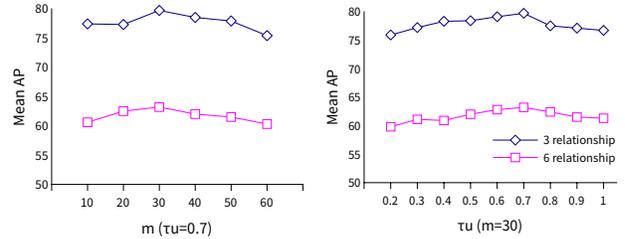


Figure 9: Evaluation of dual-glance model over variations in maximum number of region proposals (Left) and upper threshold of overlap between region proposals and the pair of people (Right).

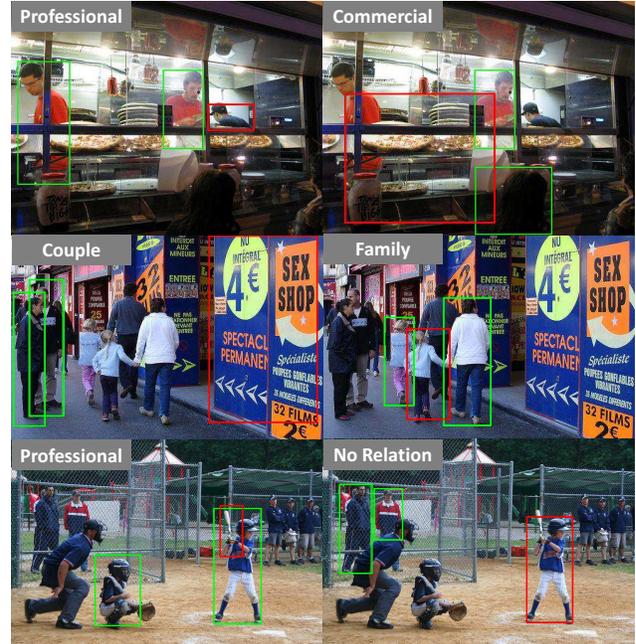


Figure 10: Illustration of the proposed attentive RCNN. GREEN boxes highlight the pair of people in question, and RED box highlights the context region with the highest attention. For each target pair, the attention mechanism fixates on different region.

region proposals used, as well as τ_u , the upper threshold of overlap between a region proposal and the target people. We experimented with different combinations of m and τ_u with the dual-glance model. As shown in Figure 9, $m = 30$ and $\tau_u = 0.7$ produce the best performance.

5.5. Visualization of Examples

The attention mechanism enables different pairs of people in question to exploit different contextual cues. Some examples are shown in Figure 10. In the second row, the little girl in red box is useful to infer that the other girl on her left and the woman on her right are family members, but her existence indicates little of the couple in black.

Figure 11 shows examples of correct recognition for each relationship category in the test set. We can observe that the proposed model learns to recognize social relationship from a wide range of visual cues including cloth-



Figure 11: Example of correct predictions on PISC dataset. Green boxes highlight the targets, and red box highlights the contextual region with highest attention.



Figure 12: Examples of incorrect predictions on PISC dataset. Yellow labels are the ground truth, and blue labels are the model's predictions.

ing, environment, surrounding people/animals, contextual objects, etc. For intimate relationships, the contextual cues varies from *beer* (friends), *gamepad* (friends), *TV* (family), to *cake* (couple) and *flowers* (couple). In terms of non-intimate relationships, the contextual cues are related to the occupations of the individuals. For instance, *goods shelf* and *scale* indicate commercial relationship, while *uniform* and *documents* imply professional relationship. Figure 12 shows the misclassified cases. The proposed model fails to recognize the gender (misclassifies *friends* as *couple* in the image at row 3 column 3), or picks up the wrong cue (the white board instead of the vegetable in the image at row 2 column 3).

6. Conclusion

In this study, we aim to address pairwise social relationship recognition, a key challenge to bridge the social gap towards higher-level social scene understanding. To this end, we propose a dual-glance model, which exploits useful information from the individual pair of interest as well as multiple contextual regions. We incorporate attention mechanism to assess the relevance of each region instance with respect to the target pair. We evaluate the proposed model on PISC dataset, a large-scale image dataset we collected to facilitate research in social scene understanding. We demonstrate both quantitatively and qualitatively the efficacy of the proposed model. We also experiment with a few variants of the proposed system to explore information useful for social relationship inference.

Our work is the first step towards general social scene understanding in a data-driven fashion. The PISC dataset provides further potential in this line of research, including but not limited to group social relation analysis, occupation recognition, and joint inference of social role and social relationships. We intend to address some of those challenges in future work.

Acknowledgment

This research is supported by the National Research Foundation, Prime Minister's Office, Singapore under its International Research Centre in Singapore Funding Initiative.

References

- [1] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese. Social LSTM: Human trajectory prediction

- in crowded spaces. In *CVPR*, pages 961–971, 2016. 1, 2
- [2] S. Alletto, G. Serra, S. Calderara, F. Solera, and R. Cucchiara. From ego to nos-vision: Detecting social relationships in first-person views. In *CVPR Workshops*, pages 594–599, 2014. 2
- [3] Y. Chen, W. H. Hsu, and H. M. Liao. Discovering informative social subgraphs and predicting pairwise relationships from group photos. In *ACMMM*, pages 669–678, 2012. 2
- [4] W. Choi and S. Savarese. A unified framework for multi-target tracking and collective activity recognition. In *ECCV*, pages 215–230, 2012. 1, 2
- [5] H. R. Conte and R. Plutchik. A circumplex model for interpersonal personality traits. *Journal of Personality and Social Psychology*, 40(4):701, 1981. 2
- [6] Z. Deng, A. Vahdat, H. Hu, and G. Mori. Structure inference machines: Recurrent neural networks for analyzing relations in group activity recognition. In *CVPR*, pages 4772–4781, 2016. 2
- [7] H. Dibeklioglu, A. A. Salah, and T. Gevers. Like father, like son: Facial expression dynamics for kinship verification. In *ICCV*, pages 1497–1504, 2013. 2
- [8] L. Ding and A. Yilmaz. Learning social relations from videos: Features, models, and analytics. In *Human-Centered Social Media Analytics*, pages 21–41. 2014. 2, 3
- [9] C. Direkoglu and N. E. O’Connor. Team activity recognition in sports. In *ECCV*, pages 69–83, 2012. 1, 2
- [10] H. Fang, S. Gupta, F. N. Iandola, R. K. Srivastava, L. Deng, P. Dollár, J. Gao, X. He, M. Mitchell, J. C. Platt, C. L. Zitnick, and G. Zweig. From captions to visual concepts and back. In *CVPR*, pages 1473–1482, 2015. 3
- [11] R. Fang, K. D. Tang, N. Snavely, and T. Chen. Towards computational models of kinship verification. In *ICIP*, pages 1577–1580, 2010. 2
- [12] A. P. Fiske. The four elementary forms of sociality: framework for a unified theory of social relations. *Psychological review*, 99(4):689, 1992. 2
- [13] A. C. Gallagher and T. Chen. Understanding images of groups of people. In *CVPR*, pages 256–263, 2009. 2
- [14] G. Gkioxari, R. B. Girshick, and J. Malik. Contextual action recognition with R*CNN. In *ICCV*, pages 1080–1088, 2015. 3, 7
- [15] Y. Guo, H. Dibeklioglu, and L. van der Maaten. Graph-based kinship recognition. In *ICPR*, pages 4287–4292, 2014. 2
- [16] E. T. Hall. *The silent language*, volume 3. Doubleday New York, 1959. 6
- [17] N. Haslam. Categories of social relationship. *Cognition*, 53(1):59–90, 1994. 2
- [18] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 5
- [19] C. Huang, C. C. Loy, and X. Tang. Unsupervised learning of discriminative attributes and visual representations. In *CVPR*, pages 5175–5184, 2016. 1
- [20] H. Hung, D. B. Jayagopi, C. Yeo, G. Friedland, S. O. Ba, J. Odobez, K. Ramchandran, N. Mirghafori, and D. Gatica-Perez. Using audio and video features to classify the most dominant person in a group meeting. In *ACMMM*, pages 835–838, 2007. 2
- [21] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L. Li, D. A. Shamma, M. S. Bernstein, and L. Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2016. 3
- [22] T. Lan, L. Sigal, and G. Mori. Social roles in hierarchical models for human activity recognition. In *CVPR*, pages 1354–1361, 2012. 1, 2
- [23] T. Lan, Y. Wang, W. Yang, S. N. Robinovitch, and G. Mori. Discriminative latent models for recognizing contextual group activities. *TPAMI*, 34(8):1549–1562, 2012. 1, 2
- [24] T. Lin, M. Maire, S. J. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. In *ECCV*, pages 740–755, 2014. 3
- [25] C. Lu, R. Krishna, M. S. Bernstein, and L. Fei-Fei. Visual relationship detection with language priors. In *ECCV*, pages 852–869, 2016. 1, 2, 5, 6
- [26] Z. Qin and C. R. Shelton. Improving multi-target tracking via social grouping. In *CVPR*, pages 1972–1978, 2012. 1
- [27] V. Ramanathan, B. Yao, and L. Fei-Fei. Social role discovery in human events. In *CVPR*, pages 2475–2482, 2013. 2, 3
- [28] S. Ren, K. He, R. B. Girshick, and J. Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *NIPS*, pages 91–99, 2015. 2, 3, 4, 5
- [29] R. Rienks, D. Zhang, D. Gatica-Perez, and W. Post. Detection and application of influence rankings in small group meetings. In *ICMI*, pages 257–264, 2006. 2
- [30] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese. Learning social etiquette: Human trajectory understanding in crowded scenes. In *ECCV*, pages 549–565, 2016. 1, 2
- [31] H. Salamin, S. Favre, and A. Vinciarelli. Automatic role recognition in multiparty recordings: Using social affiliation networks for feature extraction. *IEEE Trans. Multimedia*, 11(7):1373–1380, 2009. 2
- [32] E. R. Smith and M. A. Zarate. Exemplar and prototype use in social categorization. *Social Cognition*, 8(3):243, 1990. 1
- [33] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L. Li. YFCC100M: the new data in multimedia research. *Commun. ACM*, 59(2):64–73, 2016. 3
- [34] A. Vinciarelli, M. Pantic, D. Heylen, C. Pelachaud, I. Poggi, F. D’Errico, and M. Schröder. Bridging the gap between social animal and unsocial machine: A survey of social signal processing. *IEEE Trans. Affective Computing*, 3(1):69–87, 2012. 2
- [35] G. Wang, A. C. Gallagher, J. Luo, and D. A. Forsyth. Seeing people in social context: Recognizing people and social relationships. In *ECCV*, pages 169–182, 2010. 1, 2
- [36] J. Wu, Y. Yu, C. Huang, and K. Yu. Deep multiple instance learning for image classification and auto-annotation. In *CVPR*, pages 3460–3469, 2015. 3, 7
- [37] S. Xia, M. Shao, J. Luo, and Y. Fu. Understanding kin relationships in a photo. *IEEE Trans. Multimedia*, 14(4):1046–1056, 2012. 2
- [38] T. Xiao, Y. Xu, K. Yang, J. Zhang, Y. Peng, and Z. Zhang. The application of two-level attention models in deep convolutional neural network for fine-grained image classification. In *CVPR*, pages 842–850, 2015. 3

- [39] K. Xu, J. Ba, R. Kiros, K. Cho, A. C. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, pages 2048–2057, 2015. 3
- [40] Z. Yang, X. He, J. Gao, L. Deng, and A. Smola. Stacked attention networks for image question answering. In *CVPR*, pages 21–29, 2016. 3
- [41] Q. You, H. Jin, Z. Wang, C. Fang, and J. Luo. Image captioning with semantic attention. In *CVPR*, pages 4651–4659, 2016. 3
- [42] Z. Zhang, P. Luo, C. C. Loy, and X. Tang. Learning social relation traits from face images. In *ICCV*, pages 3631–3639, 2015. 2
- [43] B. Zhou, A. Khosla, A. Lapedriza, A. Torralba, and A. Oliva. Places: A 10 million image database for scene recognition. *TPAMI*, 2017. 6