

Graph-Based Social Relation Reasoning

Wanhua Li^{1,2,3}, Yueqi Duan^{1,2,3,5,†}, Jiwen Lu^{1,2,3,*}, Jianjiang Feng^{1,2,3}, and
Jie Zhou^{1,2,3,4}

¹ Department of Automation, Tsinghua University, China

² State Key Lab of Intelligent Technologies and Systems, China

³ Beijing National Research Center for Information Science and Technology, China

⁴ Tsinghua Shenzhen International Graduate School, Tsinghua University, China

⁵ Stanford University

li-wh17@mails.tsinghua.edu.cn, duanyq19@stanford.edu,

{lujiwen, jfeng, jzhou}@tsinghua.edu.cn

Abstract. Human beings are fundamentally sociable — that we generally organize our social lives in terms of relations with other people. Understanding social relations from an image has great potential for intelligent systems such as social chatbots and personal assistants. In this paper, we propose a simpler, faster, and more accurate method named graph relational reasoning network (GR²N) for social relation recognition. Different from existing methods which process all social relations on an image independently, our method considers the paradigm of jointly inferring the relations by constructing a social relation graph. Furthermore, the proposed GR²N constructs several virtual relation graphs to explicitly grasp the strong logical constraints among different types of social relations. Experimental results illustrate that our method generates a reasonable and consistent social relation graph and improves the performance in both accuracy and efficiency.

Keywords: Social relation reasoning, Paradigm shift, Graph neural networks, Social relation graph

1 Introduction

Social relations are the theme and basis of human life, where most human behaviors occur in the context of the relationships between individuals and others [32]. The social relation is derived from human social behaviors and defined as the association between individual persons. Social relation recognition from images is in the ascendant in the computer vision community [41, 20], while social relationships have been studied in social psychology for decades [6, 10, 2]. There has been a growing interest in understanding relations among persons in a given still image due to the broad applications such as potential privacy risk warning [41], intelligent and autonomous systems [52], and group activity analysis [19].

[†] Work done while at Tsinghua University.

* Corresponding author




Fig. 1. Examples of how the relations on the same image help each other in reasoning. We observe that social relations on an image usually follow strong *logical constraints*.

In a social scene, there are usually many people appearing at the same time, which contains various social relations. Most existing methods recognize social relations between person pairs separately [41, 20, 52, 13], where each forward pass only processes a pair of bounding boxes on an image. However, as social relations usually form a reasonable social scene, they are not independent of each other, but highly correlated instead. Independently predicting the relations on the same image suffers from the high locality in social scenes, which may result in an unreasonable and contradictory social relation graph. To this end, we consider that jointly inferring all relations for each image helps construct a reasonable and consistent social relation graph with a thorough understanding of the social scene. Moreover, as social relations on the same image usually follow strong *logical constraints*, simultaneously considering all relations can effectively exploit the consistency of them. Two examples are shown in Fig. 1. In the first example, when we know that the relation between A and C is mother-child and that between B and C is father-child, we easily infer that A and B are lovers. In the second example, we can quickly understand the social scene through the relations among A, B, and C, and infer the relations between D and others, even if D is heavily occluded. Clearly, the relations on the same image help each other in reasoning, which is not exploited in existing methods as an important cue.

In this paper, we propose a graph relational reasoning network (GR²N) to jointly infer all relations on an image with the paradigm of treating images as independent samples rather than person-pairs. In fact, the image-based paradigm is closer to the nature of human perception of social relations since human beings do not perceive the social relations of a pair of people in isolation. Many variants of graph neural networks (GNNs) can collaborate with the image-based paradigm by building a graph for each image, where the nodes represent the persons and the edges represent the relations. However, for social relation recognition, different kinds of social relations have strong logical constraints as shown in Fig. 1. Most existing GNNs’ methods simply exploit contextual information via message passing, which fails to explicitly grasp the logical constraints among different types of social relations. To exploit the strong logical constraints, the proposed GR²N constructs different virtual relation graphs for different relation types with shared node representations. Our method learns type-specific messages on each virtual relation graph and updates the node representations by

aggregating all neighbor messages across all virtual relation graphs. In the end, the final representations of nodes are utilized to predict the relations of all pairs of nodes on the graph.

To summarize, the contributions of this paper are three-fold:

- To the best of our knowledge, our method is the first attempt to jointly reason all relations on the same image for social relation recognition. Experimental results verify the superiority of the paradigm shift from person pair-based to image-based.
- We proposed a *simper* method named GR²N to collaborate with the image-based paradigm, which constructs several virtual relation graphs to explicitly model the logical constraints among different types of social relations. GR²N only uses *one input image patch* and *one convolutional neural net* while the existing approaches usually use many more patches and nets.
- The proposed GR²N is *faster* and *more accurate*. Unlike existing methods that only handle the relation of one person pair at a time, our method processes all relations on an image at the same time with a single forward pass, which makes our method computationally efficient. Finally, our method is $2\times \sim 7\times$ faster than other methods due to the paradigm shift. The proposed GR²N effectively exploits the information of relations on the same image to generate a reasonable and consistent social relation graph. Extensive experimental results on the PIPA and PISC datasets demonstrate that our approach outperforms the state-of-the-art methods.

2 Related Work

Social Relationship Recognition. Human face and body images contain rich information [26, 27, 45, 8, 22]. In recent years, social relationship recognition from an image has attracted increasing interest in the computer vision community [41, 20, 52, 13]. Some early efforts to mine social information include kinship recognition [23, 28, 44, 25], social role recognition [29, 38] and occupation recognition [37, 39]. Lu *et al.* [28] proposed a neighborhood repulsed metric learning (NRM-L) method for kinship verification. Wang *et al.* [44] explored using familial social relationships as a context for recognizing people and the social relationships between pairs of people. Ramanathan *et al.* [29] presented a conditional random field method to recognize social roles played by people in an event. Personal attributes and contextual information were usually used in occupation recognition [37]. Zhang *et al.* [51] devised an effective multitask network to learn social relation traits from face images in the wild. Sun *et al.* [41] extended the PIPA database by 26,915 social relation annotations based on the social domain theory [2]. Li *et al.* [20] collected a new social relation database and proposed a dual-glance model, which explored useful and complementary visual cues. Wang *et al.* [52] proposed a graph reasoning model to incorporate common sense knowledge of the correlation between social relationships and semantic contextual cues. Arushi *et al.* [13] used memory cells like GRUs to combine human attributes and contextual cues. However, these methods deal with all the relationships on

the same image independently, which ignore the strong logical constraints among the social relationships on an image.

Graph Neural Networks. The notation of GNNs was firstly introduced by Gori *et al.* [14], and further elaborated in [35]. Scarselli *et al.* trained the GNNs via the Almeida-Pineda algorithm. Li *et al.* [24] modified the previous GNNs to use gated recurrent units and modern optimization techniques. Inspired by the huge success of convolutional networks in the computer vision domain, many efforts have been devoted to re-defining the notation of convolution for graph data [1, 7, 18]. One of the most famous works is GCN proposed by Kipf *et al.* [18]. Li *et al.* [21] developed deeper insights into the GCN model and argued that if a GCN was deep with many convolutional layers, the output features might be over-smoothed. Gilmer *et al.* [11] reformulated the existing GNN models into a single common framework called message passing neural networks (MPNNs). Veličković *et al.* [43] presented graph attention networks leveraging masked self-attentional layers. Xu *et al.* [48] characterized the discriminative power of popular GNN variants and developed a simple architecture that was provably the most expressive among the class of GNNs. Hamilton *et al.* [15] presented GraphSAGE to efficiently generate node embeddings for previously unseen data. GNNs have been proven to be an effective tool for relational reasoning [36, 3, 17, 40, 42]. Schlichtkrull *et al.* [36] proposed R-GCNs and applied them to knowledge base completion tasks. Chen *et al.* [4] proposed a GloRe unit for reasoning globally.

3 Approach

In this section, we first illustrate the importance of the paradigm shift. Then we present the details of the graph building and the proposed GR²N. Finally, we provide comparisons between the proposed approach and other GNNs’ methods.

3.1 Revisiting the Paradigm of Social Relation Recognition

Formally, the problem of social relation recognition can be formulated as a probability function: given an input image \mathbf{I} , bounding box values \mathbf{b}_i and queries \mathbf{x} :

$$\mathbf{x} = \{x_{i,j} | i = 1, 2, \dots, N, j = 1, 2, \dots, N\} \quad (1)$$

where $x_{i,j}$ is the social relation between the person i and j , and N is the total number of people in the image \mathbf{I} . The goal of social relation recognition is to find an optimal value of \mathbf{x} :

$$\mathbf{x}^* = \arg \max_{\mathbf{x}} P(\mathbf{x} | \mathbf{I}, \mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_N) \quad (2)$$

As far as we know, all existing methods of social relation recognition organize data as person pair-based, and each sample of these methods consists of an input image \mathbf{I} and a target pair of people highlighted by bounding boxes $\mathbf{b}_i, \mathbf{b}_j$. Therefore, the actual optimization objective of these methods is as follows:

$$x_{i,j}^* = \arg \max_{x_{i,j}} P(x_{i,j} | \mathbf{I}, \mathbf{b}_i, \mathbf{b}_j) \quad (3)$$

The optimization objective in (3) is consistent with that in (2) if and only if the following equation holds:

$$P(\mathbf{x}|\mathbf{I}, \mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_N) = \prod_{1 \leq i, j \leq N} P(x_{i,j}|\mathbf{I}, \mathbf{b}_i, \mathbf{b}_j) \quad (4)$$

It is known that (4) holds if and only if all relations on the input image \mathbf{I} are independent of each other. Unfortunately, as we discussed earlier, this is not true and the relations on an image are usually highly related. Let \mathbf{y} be the ground truths corresponding to queries \mathbf{x} : $\mathbf{y} = \{y_{i,j}|i = 1, 2, \dots, N, j = 1, 2, \dots, N\}$. Then the paradigm of existing methods treats quadruplets $\langle \mathbf{I}, \mathbf{b}_i, \mathbf{b}_j, y_{i,j} \rangle$ as independent samples to optimize the objective in (3), which can not lead to the optimal value of that in (2). On the contrary, we reorganize the data as image-based to directly optimize the objective formulated in (2) with a thorough understanding of the social scene. Our method formulates each sample as an $(N+2)$ -tuple $\langle \mathbf{I}, \mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_N, \mathbf{y} \rangle$ and jointly infers the relations for each image.

3.2 From Image to Graph

We model each person on an image as a node in a graph, and the edge between two nodes represents the relation between the corresponding two people. Then graph operations are applied to perform joint reasoning. So the first step is to create nodes in the graph using the input image \mathbf{I} and bounding boxes $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_N$. As commonly used in detection [12, 33, 30, 31], a convolutional neural network takes the input image \mathbf{I} as input, and the features of object proposals are extracted directly from the last convolutional feature map \mathbf{H} . More specifically, we first map the bounding boxes $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_N$ to $\mathbf{b}'_1, \mathbf{b}'_2, \dots, \mathbf{b}'_N$ according to the transformation from \mathbf{I} to \mathbf{H} . Then the feature \mathbf{p}_i of person i is obtained by employing an RoI pooling layer on the feature map \mathbf{H} : $\mathbf{p}_i = f_{RoIPooling}(\mathbf{H}, \mathbf{b}'_i)$. In this way, we obtain all the features of people on the image \mathbf{I} : $\mathbf{p} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N\}$, $\mathbf{p}_i \in \mathbb{R}^F$, where F is the feature dimension for each person. Subsequently, these features are set as the initial feature representations of nodes in the graph: $\mathbf{h}^0 = \{\mathbf{h}_1^0, \mathbf{h}_2^0, \dots, \mathbf{h}_N^0\}$, where $\mathbf{h}_i^0 = \mathbf{p}_i$.

3.3 Graph Relational Reasoning Network

Having created the nodes in the graph, we propose a graph relational reasoning network to perform relational reasoning on the graph. We begin with a brief description of the overall framework of our GR²N. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denotes a graph with node feature vectors \mathbf{h}_v for $v \in \mathcal{V}$ and edge features \mathbf{e}_{vw} for $(v, w) \in \mathcal{E}$. The framework of GR²N can be formulated as a message passing phase and a readout phase following [11]. The message passing phase, which is defined in terms of message function M_t and vertex update function U_t , runs for T time steps. During the message passing phase, the aggregated messages \mathbf{m}_v^{t+1} are used

to update the hidden states \mathbf{h}_v^t :

$$\mathbf{m}_v^{t+1} = \sum_{w \in N(v)} M_t(\mathbf{h}_v^t, \mathbf{h}_w^t, e_{vw}) \quad (5)$$

$$\mathbf{h}_v^{t+1} = U_t(\mathbf{h}_v^t, \mathbf{m}_v^{t+1}) \quad (6)$$

where $N(v)$ denotes the neighbors of v in graph \mathcal{G} . The readout phase computes feature vectors for edges using the readout function R . The message function M_t , vertex update function U_t , and readout function R are all differentiable functions.

In each time step of the message passing phase, the state of the node is updated according to the messages sent by its neighbor nodes, which requires us to know the topology of the entire graph. In social relation recognition, we use edges to represent social relations, and the edge type represents the category of the relation. However, the task of social relation recognition is to predict the existence of edges and the type of edges, which means that we do not know the topology of the graph at the beginning. To address this issue, the proposed GR²N first models the edges in the graph.

The social relations on an image have strong *logical constraints* and our GR²N aims to grasp these constraints to generate a reasonable and consistent social relation graph. As each kind of social relations has its specific logical constraints, we let GR²N propagate different types of messages for different types of edges. Specifically, we construct a virtual relation graph for each kind of relationship, in which the edge represents the existence of this kind of relationship between two nodes. To achieve mutual reasoning and information fusion among different relationships, the feature representations of nodes are shared across all created virtual relation graphs.

Mathematically, we use K to denote the number of social relationship categories, and K virtual relation graphs $(\mathcal{V}, \mathcal{E}^1), (\mathcal{V}, \mathcal{E}^2), \dots, (\mathcal{V}, \mathcal{E}^K)$ are created to model the K social relationships separately. Assuming that we have obtained the node feature representations at time step t : $\mathbf{h}^t = \{\mathbf{h}_1^t, \mathbf{h}_2^t, \dots, \mathbf{h}_N^t\}$, we first model the edge embedding in the virtual relation graph:

$$\mathbf{e}_{i,j}^{k,t+1} = f_{r_k}(\mathbf{h}_i^t, \mathbf{h}_j^t) = [\mathbf{W}_{r_k} \mathbf{h}_i^t \parallel \mathbf{W}_{r_k} \mathbf{h}_j^t] \quad (7)$$

where $\mathbf{e}_{i,j}^{k,t+1}$ denotes the embedding of edge (i, j) in virtual relation graph $(\mathcal{V}, \mathcal{E}^k)$ at time $t + 1$ and \parallel is the concatenation operation. $f_{r_k}(\cdot)$ is parameterized by a weight matrix $\mathbf{W}_{r_k} \in \mathbb{R}^{F \times F}$. In standard GNNs, the existence of edges is binary and deterministic. In this paper, we consider rather a probabilistic soft edge, *i.e.* edge (i, j) in the virtual relation graph $(\mathcal{V}, \mathcal{E}^k)$ exists according to the probability:

$$\alpha_{i,j}^{k,t+1} = \sigma(\mathbf{a}_{r_k}^\top \mathbf{e}_{i,j}^{k,t+1}) \quad (8)$$

where $\mathbf{a}_{r_k} \in \mathbb{R}^{2F}$ is a weight vector and \cdot^\top denotes transposition. The sigmoid function $\sigma(\cdot)$ is employed as the activation function to normalize the values between 0 and 1. In the end, the \mathcal{E}^k can be represented as the set $\{\alpha_{i,j}^k | 1 \leq i, j \leq N\}$ if we ignore the time step.

Once obtained, we propagate messages on each virtual relation graph according to the soft edges and aggregate messages across all virtual relation graphs:

$$\mathbf{m}_i^{t+1} = \sum_{j \in N(i)} \sum_{k=1}^K \alpha_{i,j}^{k,t+1} f_m(\mathbf{h}_i^t, \mathbf{h}_j^t, \mathbf{e}_{i,j}^{k,t+1}) \quad (9)$$

For simplicity, we reuse the previous weight \mathbf{W}_{r_k} and let $f_m(\mathbf{h}_i^t, \mathbf{h}_j^t, \mathbf{e}_{i,j}^{k,t+1}) = \mathbf{W}_{r_k} \mathbf{h}_j^t$. So we reformulate \mathbf{m}_i^{t+1} as:

$$\mathbf{m}_i^{t+1} = \sum_{j \in N(i)} M_t(\mathbf{h}_i^t, \mathbf{h}_j^t, \mathbf{e}_{i,j}^{k,t+1}) = \sum_{j \in N(i)} \left(\sum_{k=1}^K \alpha_{i,j}^{k,t+1} \mathbf{W}_{r_k} \mathbf{h}_j^t \right) \quad (10)$$

Finally, we use the aggregated messages to update the hidden state of node:

$$\mathbf{h}_i^{t+1} = U_t(\mathbf{h}_i^t, \mathbf{m}_i^{t+1}) = f_{GRU}(\mathbf{h}_i^t, \mathbf{m}_i^{t+1}) \quad (11)$$

where $f_{GRU}(\cdot)$ is the Gated Recurrent Unit update function introduced in [5].

Having repeated the above process for T time steps, we obtain the final node feature representations: $\mathbf{h}^T = \{\mathbf{h}_1^T, \mathbf{h}_2^T, \dots, \mathbf{h}_N^T\}$. Then the readout function R is applied to the features \mathbf{h}^T for social relation recognition. As mentioned above, the task of social relation recognition is to predict the existence and type of edges in the graph, which is exactly what virtual social graphs accomplish. So the readout function R simply reuses the functions that create the virtual social graphs:

$$\hat{x}_{i,j}^k = R(\mathbf{h}_i^T, \mathbf{h}_j^T, k) = \sigma(\mathbf{a}_{r_k}^\top [\mathbf{W}_{r_k} \mathbf{h}_i^T \parallel \mathbf{W}_{r_k} \mathbf{h}_j^T]) \quad (12)$$

where $\hat{x}_{i,j}^k$ indicates the probability of the existence of the k -th social relation between person i and j . During the test stage, the social relation between person i and j is chosen as the category with the greatest probability: $\arg \max_k \hat{x}_{i,j}^k$.

The overall framework including the backbone CNN and the proposed GR²N is jointly optimized end-to-end. If we extend the ground truths \mathbf{y} to one hot, then the loss function for the sample $\langle \mathbf{I}, \mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_N, \mathbf{y} \rangle$ is:

$$\mathcal{L} = \sum_{1 \leq i, j \leq N} \sum_{k=1}^K -[y_{i,j}^k \log(\hat{x}_{i,j}^k) + (1 - y_{i,j}^k) \log(1 - \hat{x}_{i,j}^k)] \quad (13)$$

where $y_{i,j}^k$ is the k -th element of one hot form of $y_{i,j}$.

3.4 Discussion

Another intuitive way to jointly model all the social relationships among people on an image is to use some common variants of GNNs, such as GGNN [24] and GCN [18]. They can be utilized to perform graph operations on the graphs to replace the proposed GR²N. In addition, GNNs are used in some other tasks such as group activity recognition [9], and scene graph generation [47] to model

the relationships between objects. In the following, we will elaborate on the key difference between the proposed GR²N and those GNNs’ methods.

Different from group activity recognition, scene graph recognition, and other tasks, different kinds of social relations have *strong logical constraints* in the social relation recognition task. For example, knowing that A-C is mother-child and B-C is father-child, we can directly infer that A-B is spouses even without checking the relation A-B itself. However, most existing methods of jointly reasoning all relations such as GGNN [24], GCN [18], Structure Inference Machines [9] and Iterative Message Passing [47] simply exploit *contextual information* via message passing, which fails to *explicitly* grasp the *logical constraints* among different types of social relations. Instead, our GR²N constructs a virtual social relation graph for each social relation type, where each graph learns a *type-specific* message passing matrix \mathbf{W}_{r_k} . As each type of social relation has its *specific logical constraints*, the proposed GR²N is aware of *relation types* and better exploits the *type-specific logical constraints* in social relation recognition. In this way, our method can better perform relational reasoning and generate a reasonable and consistent social relation graph.

4 Experiments

In this section, we first introduce the datasets and present some implementation details of our approach. Then we evaluate the performance of our method using quantitative and qualitative analysis.

4.1 Datasets

We evaluated the GR²N and existing competing methods on the People in Photo Album (PIPA) [41] database and People in Social Context (PISC) database [20].

PIPA: The PIPA database is collected from Flickr photo albums for the task of person recognition [50]. Then the dataset is extended by Sun *et al.* with 26,915 person pair annotations based on the social domain theory [2]. The PIPA dataset involves two-level recognition tasks: 1) social domain recognition focuses on five categories of social domain, i.e., attachment domain, reciprocity domain, mating domain, hierarchical power domain, and coalitional group domain; 2) social relationship recognition focuses on 16 finer categories of relationship, such as friends, classmates, father-child, leader-subordinate, band members, and so on. For fair comparisons, we adopt the standard train/val/test split released by Sun *et al.* [41], which uses 13,672 domain/relation instances in 5,857 images for training, 709 domain/relation instances in 261 images for validation, and 5,106 domain/relation instances in 2,452 images for testing. The top-1 classification accuracy is used for evaluation.

PISC: The PISC database collects 22,670 images and is annotated following the relational model theory [10]. It has hierarchical social relationship categories: coarse-grained relationships (intimate, not-intimate and no relation) and fine-grained relationships (commercial, couple, family, friends, professional and

no-relation). For coarse-grained relationship, 13,142 images with 49,017 relationship instances are used for training, 4,000 images with 14,536 instances are used for validation, 4,000 images with 15,497 instances are used for testing. For fine-grained relationship, the train/val/test set has 16,828 images and 55,400 instances, 500 images and 1,505 instances, 1250 images and 3,961 instances, respectively. We follow the commonly used metrics [13, 20, 52] on the PISC dataset and report the per-class recall for each relationship and the mean average precision (mAP) over all relationships.

4.2 Implementation Details

For fair comparisons, we employed ResNet-101 as the backbone CNN following [20, 52]. The ResNet-101 was initialized with the pre-trained weights on ImageNet [34]. The shape of the output region of the RoI pooling layer was set to be the same as the shape of the last convolution feature map \mathbf{H} , and then a global average pooling layer was used to obtain a 2048-D feature for each person. The time step T was set to 1, which achieved the best result. The model was trained by Adam optimization [16] with a learning rate of 10^{-5} . We trained our method for 10 epochs with a batch size of 32. During training, images were horizontally flipped with probability 0.5 and randomly cropped for data augmentation.

In a mini-batch, different images might have different numbers of people, which made the size of the graphs variable. To deal with this problem, we set the number of people N on an image to the maximum possible persons N_{max} . If the number of people on an image was less than N_{max} , then we set the missing nodes as empty ones and there would be no soft edges between the empty nodes and the real nodes. For those relationships that were not labeled, no loss would occur from them in (13) to avoid having an impact on network training. Another issue was that the social relationships on the PISC dataset were highly imbalanced, especially for the fine-level relationships. To address this, we adopted the reweighting strategy for fine-grained relationship recognition on the PISC dataset. Specifically, the samples were reweighted inversely proportionally to the class frequencies.

4.3 Results and Analysis

Comparisons with the State-of-the-Art Methods. We compare the proposed approach with several existing state-of-the-art methods. The details of these competing methods are as follows:

Finetuned CNN + SVM [41]. Double-stream CaffeNet is used to extract features, then the extracted features are utilized to train a linear SVM.

All attributes + SVM [41]. Many semantic attribute categories including age, gender, location, and activity are used in this method. Then all attribute features are concatenated to train a linear SVM.

Pair CNN [20]. This model consists of two CNNs (ResNet-101) with shared weights. The input is two cropped image patches for the two individuals and the extracted features are concatenated for social relation recognition.

Table 1. Comparisons of the accuracy between our GR²N and other state-of-the-art methods on the PIPA dataset.

methods	# of input image patches	# of CNNs	domain	relation
Finetuned CNN + SVM [41]	2	2	63.2%	48.6%
All attributes + SVM [41]	4	10	67.8%	57.2%
Pair CNN [20]	2	1	65.9%	58.0%
Dual-Glance [20]	4	3	-	59.6%
SRG-GN [13]	3	5	-	53.6%
GRM [52]	4	3	-	62.3%
MGR [49]	4	5	-	64.4%
GR ² N-224	1	1	69.3%	61.3%
GR ² N-448	1	1	72.3%	64.3%

Dual-Glance [20]. Two cropped individual patches and the union region of them are sent to CNNs to extract features. The introduced second glance exploits surrounding proposals as contextual information to refine the predictions.

SRG-GN [13]. Five CNNs are utilized to extract scene and attribute context information (*e.g.*, age, gender, and activity). Then these features are utilized to update the state of memory cells like GRUs.

MGR [49]. It employ five CNNs to extract the global, middle, and fine granularity features to comprehensively capture the multi-granularity semantics.

GRM [52]. It replaces the second glance in Dual-Glance with a pre-trained Faster-RCNN detector [33] and a Gated Graph Neural Network [24] to model the interaction between the contextual objects and the persons of interest.

All of the above methods are person pair-based, which means that they consider the social relations on the same image separately. It is worth noting that although the SRG-GN mentions that they can generate a social relationship graph, they actually still process each relationship independently during the training and testing phases. These methods usually crop the image patches for interested individuals and resize them to 224×224 pixels while our approach takes the entire image as input. In order to make a fair comparison, we should choose the appropriate entire input image size to ensure that the personal area of interest is roughly equal to 224×224 . In the PIPA and PISC dataset, the area of the bounding boxes is on average 1/4 and 1/5 of the area of the images, respectively, so we resize the original image to 448×448 pixels, which is denoted as GR²N-448. Although this is still unfavorable to our approach on the PISC datasets, our method can still achieve promising performance. Besides, we also report the performance of GR²N-224 to show the effectiveness of our method, which resizes the original image to 224×224 pixels.

The experimental results of social domain recognition and social relationship recognition on the PIPA database are shown in Table 1. We first compare our method with a simple baseline Pair CNN, which, like our method, does not use any scene and attribute context cues, so we can see the benefits brought by the

Table 2. We present the per-class recall for each relationship and the mAP over all relationships (in %) on the PISC dataset. The first and second best scores are highlighted in red and blue colors, respectively. (Int: Intimate, Non: Non-Intimate, NoR: No Relation, Fri: Friends, Fam: Family, Cou: Couple, Pro: Professional, Com: Commercial)

Methods	Coarse relationships				Fine relationships							
	Int	Non	NoR	mAP	Fri	Fam	Cou	Pro	Com	NoR	mAP	
Pair CNN [20]	70.3	80.5	38.8	65.1	30.2	59.1	69.4	57.5	41.9	34.2	48.2	
Dual-Glance [20]	73.1	84.2	59.6	79.7	35.4	68.1	76.3	70.3	57.6	60.9	63.2	
GRM [52]	81.7	73.4	65.5	82.8	59.6	64.4	58.6	76.6	39.5	67.7	68.7	
MGR [49]	-	-	-	-	64.6	67.8	60.5	76.8	34.7	70.4	70.0	
SRG-GN [13]	-	-	-	-	25.2	80.0	100.0	78.4	83.3	62.5	71.6	
GR ² N-224	76.3	65.3	74.0	72.2	54.7	68.5	78.1	78.0	49.7	57.5	68.6	
GR ² N-448	81.6	74.3	70.8	83.1	60.8	65.9	84.8	73.0	51.7	70.4	72.7	

paradigm shift. We observe that our method significantly outperforms Pair CNN. Specifically, the GR²N-448 achieves an accuracy of 64.3% for social relation recognition and 72.3% for social domain recognition, improving the Pair CNN by 6.3% and 6.4% respectively. What’s more, even the GR²N-224 with inferior input image size outperforms the Pair CNN by 3.3% and 3.4% for social relation recognition and social domain recognition respectively, which demonstrates that the paradigm shift and GR²N’s superior relational reasoning ability bring significant performance gains.

Next, we compare our method with state-of-the-art methods. In terms of social relation recognition, the proposed GR²N-448 improves GRM by 2.0%. What is worth mentioning is that GRM uses three individual CNNs to extract features from four image patches including one 448×448 entire image and three 224×224 person patches to exploit key contextual cues, while the GR²N only uses one input image and one convolutional neural network. Our method also achieves competitive results with MGR whereas MGR uses 5 CNNs and 4 patches. For social domain recognition, our method GR²N-448 achieves an accuracy of 72.3%, which improves the performance of All attributes + SVM by 4.5% without using any semantic attribute categories. Clearly, these results illustrate the superiority of our method.

Table 2 shows the experimental comparison with the recent state-of-the-art methods on the PISC database. We observe that our method achieves an mAP of 83.1% for the coarse-level recognition and 72.7% for the fine-level recognition, improving the simple baseline Pair CNN by a large margin: 18.0% for coarse relationships and 24.5% for fine relationships, which further validates the advantage of the paradigm shift. Our approach outperforms SRG-GN by 1.1% with a much simpler framework (1 neural net vs. 5 neural nets, 1 image patch vs. 3 image patches). Compared with GRM, GR²N-448 achieves competitive performance for coarse relationship recognition and superior performance for fine relation-

Table 3. Comparisons of the runtime (seconds/image) under different batch sizes for social relation recognition on the PIPA dataset.

methods	batch size			
	1	2	4	8
GRM	0.294	0.171	0.089	*
Pair CNN	0.077	0.045	0.039	0.037
GR ² N-448	0.046	0.025	0.021	0.021

Table 4. Comparisons with GCN and GGNN for social relation recognition.

methods	accuracy
Pair CNN	58.0%
GCN	59.3%
GGNN	59.8%
GR ² N	61.3%

ship recognition with only one image patch and one neural net, while the GRM uses four image patches and three neural nets. It is worth noting that the experimental results of our approach on the PISC dataset are achieved with inferior input image size, which further illustrates the effectiveness of our GR²N.

Runtime Analysis. In addition to being *simpler* and *more accurate*, another advantage of our GR²N is that it is *faster*. Since our method handles all relationships on an image at the same time, the GR²N is computationally efficient. We conduct experiments under different batch sizes to compare the speed of different methods using a GeForce GTX 1080Ti GPU on the PIPA dataset. To be fair, the version of GR²N used here is GR²N-448. The results of the forward runtime (seconds/image) are reported in Table 3. The * in the table represents memory overflow. Compared with Pair CNN, our method achieves about 2× speed-ups, which shows the benefits of the paradigm shift. The GR²N-448 is 4× ∼ 7× faster than the state-of-the-art method GRM, which further demonstrates the *efficiency* of our approach. Since the average number of bounding boxes per image on the PIPA datasets is only 2.5, we expect higher speedup when there are more individuals on an image.

Ablation Study. To validate the effectiveness of the proposed GR²N, we compare it with two commonly used GNN variants: GGNN [24] and GCN [18]. To apply GGNN and GCN, we connect the nodes together into a fully connected graph. Following [46], we also add a residual connection in every layer of GCN, which observes an improvement in performance in the experiments. For a fair comparison, all three methods use an image of size 224 × 224 as input. The readout function of GGNN and GCN is a two-layer neural network and the edge embedding is obtained by concatenating the features of the corresponding two nodes. The cross-entropy loss function is used to train GGNN and GCN.

The comparisons on the PIPA dataset by accuracy are listed in Table 4. We observe that our proposed GR²N outperforms GCN and GGNN, which demonstrates that the proposed GR²N effectively performs relational reasoning and exploits the strong logical constraints of social relations on the same image. On the other hand, all three methods adopt the image-based paradigm and we see that they all outperform Pair CNN. In fact, these three methods achieve superior or competitive performance compared with some of the state-of-the-art methods which adopt the person pair-based paradigm, such as Dual-Glance and SRG-GN.




Fig. 2. Some examples of qualitative evaluation. The first row shows the input images, the second row is the corresponding predicted results of Pair CNN, and the third row is the corresponding predicted results of our method. The red cross mark indicates a prediction error, while the green check mark indicates that the corresponding error is corrected. We see that our method corrects the prediction errors by performing graph-based relational reasoning.

Note that these GNNs' methods only use one input patch and one convolutional neural net, which illustrates the superiority of the image-based paradigm over the person pair-based paradigm.

Qualitative Evaluation. In this part, we present some examples to illustrate how our GR²N improves performance by performing relational reasoning and exploiting the strong logical constraints of social relations on a graph in Fig. 2. For the first example, Pair CNN independently processes the relations on the image and predicts that the relation between B and C is friends. On the contrary, our method jointly infers all relations for each image. We easily infer that the relation between B and C is classmates when we know that the relation between A and B and the relation between A and C both are teacher-student. This is also the case in the rest of the examples. When the relationships on an image are processed independently, some obviously unreasonable and contradictory relationships may occur simultaneously, and our method is capable of correcting them through relational reasoning with a thorough understanding of the social scene. Clearly, our method better models the interactions among people on the graph, and generates a reasonable and consistent social relation graph. We also present some failure cases in Fig. 3. We see that our method can not correct the prediction error in the first case since the error does not cause any obvious conflict. Although the ground truth of the relationship between person A and B is friends, predicting it as lovers does not lead to a contradictory social rela-




Fig. 3. Some failure cases of our method on the PIPA dataset. The first column shows the input images, the second column is the corresponding predicted results of Pair CNN, and the third column is the corresponding predicted results of our method. The red cross mark indicates a prediction error.

tion graph. Our method may also fail when most predictions are wrong and our model performs relational reasoning based on these false predictions, as shown in the second case.

5 Conclusion

In this paper, we have presented the GR^2N for social relation recognition, which is a *simpler*, *faster*, and *more accurate* method. Unlike existing methods, we simultaneously reason all relations for each image and treat images as independent samples rather than person-pairs. Our method constructs a graph for each image to model the interactions among people and perform relational reasoning on this graph with fully exploiting the strong logical constraints of relations. Experimental results demonstrate that our method generates a reasonable and consistent social relation graph. Moreover, GR^2N achieves better performance with less time cost compared with the state-of-the-art methods, which further illustrates the *effectiveness* and *efficiency* of our method.

Acknowledgments

This work was supported in part by the National Key Research and Development Program of China under Grant 2017YFA0700802, in part by the National Natural Science Foundation of China under Grant 61822603, Grant U1813218, Grant U1713214, and Grant 61672306, in part by Beijing Natural Science Foundation under Grant No. L172051, in part by Beijing Academy of Artificial Intelligence (BAAI), in part by a grant from the Institute for Guo Qiang, Tsinghua University, in part by the Shenzhen Fundamental Research Fund (Subject Arrangement) under Grant JCYJ20170412170602564, and in part by Tsinghua University Initiative Scientific Research Program.

References

1. Bruna, J., Zaremba, W., Szlam, A., LeCun, Y.: Spectral networks and locally connected networks on graphs. arXiv preprint arXiv:1312.6203 (2013)
2. Bugental, D.B.: Acquisition of the algorithms of social life: A domain-based approach. *Psychological Bulletin* **126**(2), 187 (2000)
3. Chen, S., Abhinav, S., Carl, V., Rahul, S., Kevin, M., Cordelia, S.: Relational action forecasting. In: CVPR (2019)
4. Chen, Y., Rohrbach, M., Yan, Z., Shuicheng, Y., Feng, J., Kalantidis, Y.: Graph-based global reasoning networks. In: CVPR. pp. 433–442 (2019)
5. Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using rnn encoder-decoder for statistical machine translation. In: EMNLP (2014)
6. Conte, H.R., Plutchik, R.: A circumplex model for interpersonal personality traits. *Journal of Personality and Social Psychology* **40**(4), 701 (1981)
7. Defferrard, M., Bresson, X., Vandergheynst, P.: Convolutional neural networks on graphs with fast localized spectral filtering. In: NeurIPS. pp. 3844–3852 (2016)
8. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: CVPR. pp. 4690–4699 (2019)
9. Deng, Z., Vahdat, A., Hu, H., Mori, G.: Structure inference machines: Recurrent neural networks for analyzing relations in group activity recognition. In: CVPR. pp. 4772–4781 (2016)
10. Fiske, A.P.: The four elementary forms of sociality: framework for a unified theory of social relations. *Psychological Review* **99**(4), 689 (1992)
11. Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E.: Neural message passing for quantum chemistry. In: ICML. pp. 1263–1272 (2017)
12. Girshick, R.: Fast r-cnn. In: ICCV. pp. 1440–1448 (2015)
13. Goel, A., Ma, K.T., Tan, C.: An end-to-end network for generating social relationship graphs. In: CVPR. pp. 11186–11195 (2019)
14. Gori, M., Monfardini, G., Scarselli, F.: A new model for learning in graph domains. In: IJCNN. pp. 729–734 (2005)
15. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. In: NeurIPS. pp. 1024–1034 (2017)
16. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
17. Kipf, T., Fetaya, E., Wang, K.C., Welling, M., Zemel, R.: Neural relational inference for interacting systems. In: ICML (2018)
18. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: ICLR (2017)
19. Lan, T., Sigal, L., Mori, G.: Social roles in hierarchical models for human activity recognition. In: CVPR. pp. 1354–1361 (2012)
20. Li, J., Wong, Y., Zhao, Q., Kankanhalli, M.S.: Dual-glance model for deciphering social relationships. In: ICCV. pp. 2650–2659 (2017)
21. Li, Q., Han, Z., Wu, X.M.: Deeper insights into graph convolutional networks for semi-supervised learning. In: AAAI (2018)
22. Li, W., Lu, J., Feng, J., Xu, C., Zhou, J., Tian, Q.: Bridgenet: A continuity-aware probabilistic network for age estimation. In: CVPR. pp. 1145–1154 (2019)
23. Li, W., Zhang, Y., Lv, K., Lu, J., Feng, J., Zhou, J.: Graph-based kinship reasoning network. In: ICME. pp. 1–6 (2020)

24. Li, Y., Tarlow, D., Brockschmidt, M., Zemel, R.: Gated graph sequence neural networks. ICLR (2016)
25. Lu, J., Hu, J., Tan, Y.P.: Discriminative deep metric learning for face and kinship verification. TIP **26**(9), 4269–4282 (2017)
26. Lu, J., Tan, Y.P.: Cost-sensitive subspace analysis and extensions for face recognition. TIFS **8**(3), 510–519 (2013)
27. Lu, J., Wang, G., Deng, W., Jia, K.: Reconstruction-based metric learning for unconstrained face verification. TIFS **10**(1), 79–89 (2014)
28. Lu, J., Zhou, X., Tan, Y.P., Shang, Y., Zhou, J.: Neighborhood repulsed metric learning for kinship verification. TPAMI **36**(2), 331–345 (2013)
29. Ramanathan, V., Yao, B., Fei-Fei, L.: Social role discovery in human events. In: CVPR. pp. 2475–2482 (2013)
30. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: CVPR. pp. 779–788 (2016)
31. Redmon, J., Farhadi, A.: Yolo9000: better, faster, stronger. In: CVPR. pp. 7263–7271 (2017)
32. Reis, H.T., Collins, W.A., Berscheid, E.: The relationship context of human behavior and development. Psychological Bulletin **126**(6), 844 (2000)
33. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NeurIPS. pp. 91–99 (2015)
34. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. IJCV **115**(3), 211–252 (2015)
35. Scarselli, F., Gori, M., Tsoi, A.C., Hagenbuchner, M., Monfardini, G.: The graph neural network model. TNNLS **20**(1), 61–80 (2008)
36. Schlichtkrull, M., Kipf, T.N., Bloem, P., Van Den Berg, R., Titov, I., Welling, M.: Modeling relational data with graph convolutional networks. In: ESWC. pp. 593–607 (2018)
37. Shao, M., Li, L., Fu, Y.: What do you do? occupation recognition in a photo via social context. In: ICCV. pp. 3631–3638 (2013)
38. Shu, T., Xie, D., Rothrock, B., Todorovic, S., Chun Zhu, S.: Joint inference of groups, events and human roles in aerial videos. In: CVPR. pp. 4576–4584 (2015)
39. Song, Z., Wang, M., Hua, X.s., Yan, S.: Predicting occupation via human clothing and contexts. In: ICCV. pp. 1084–1091 (2011)
40. Sun, C., Karlsson, P., Wu, J., Tenenbaum, J.B., Murphy, K.: Stochastic prediction of multi-agent interactions from partial observations. In: ICLR (2019)
41. Sun, Q., Schiele, B., Fritz, M.: A domain based approach to social relation recognition. In: CVPR. pp. 3481–3490 (2017)
42. Tacchetti, A., Song, H.F., Mediano, P.A., Zambaldi, V., Rabinowitz, N.C., Graepel, T., Botvinick, M., Battaglia, P.W.: Relational forward models for multi-agent learning. In: ICLR (2019)
43. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph Attention Networks. In: ICLR (2018)
44. Wang, G., Gallagher, A., Luo, J., Forsyth, D.: Seeing people in social context: Recognizing people and social relationships. In: ECCV. pp. 169–182 (2010)
45. Wang, T., Gong, S., Zhu, X., Wang, S.: Person re-identification by video ranking. In: ECCV. pp. 688–703 (2014)
46. Wang, X., Gupta, A.: Videos as space-time region graphs. In: ECCV. pp. 399–417 (2018)
47. Xu, D., Zhu, Y., Choy, C.B., Fei-Fei, L.: Scene graph generation by iterative message passing. In: CVPR. pp. 5410–5419 (2017)

48. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? In: ICLR (2019)
49. Zhang, M., Liu, X., Liu, W., Zhou, A., Ma, H., Mei, T.: Multi-granularity reasoning for social relation recognition from images. In: ICME. pp. 1618–1623 (2019)
50. Zhang, N., Paluri, M., Taigman, Y., Fergus, R., Bourdev, L.: Beyond frontal faces: Improving person recognition using multiple cues. In: CVPR. pp. 4804–4813 (2015)
51. Zhang, Z., Luo, P., Loy, C.C., Tang, X.: From facial expression recognition to interpersonal relation prediction. *IJCV* **126**(5), 550–569 (2018)
52. Zhouxia, W., Tianshui, C., Jimmy, R., Weihao, Y., Hui, C., Liang, L.: Deep reasoning with knowledge graph for social relationship understanding. In: IJCAI (2018)