

A Domain Based Approach to Social Relation Recognition

Qianru Sun Bernt Schiele Mario Fritz

Max Planck Institute for Informatics, Saarland Informatics Campus

{qsun, schiele, mfritz}@mpi-inf.mpg.de

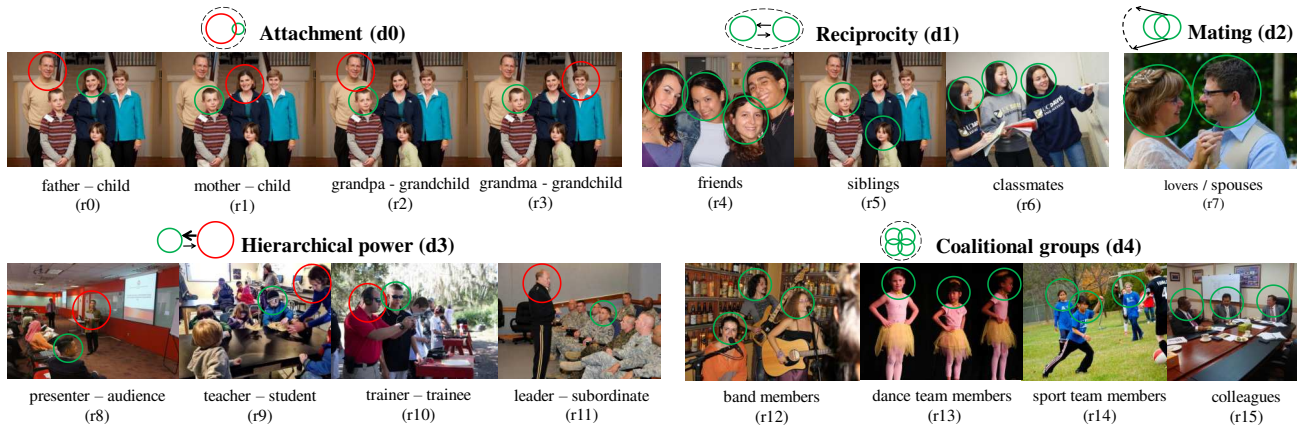


Figure 1: We investigate the recognition of social relations in a domain-based approach. Our study is based on BURGENTAL’s social psychology theory [1] that partitions social life into 5 domains from which we derive 16 social relations.

Abstract

Social relations are the foundation of human daily life. Developing techniques to analyze such relations from visual data bears great potential to build machines that better understand us and are capable of interacting with us at a social level. Previous investigations have remained partial due to the overwhelming diversity and complexity of the topic and consequently have only focused on a handful of social relations. In this paper, we argue that the domain-based theory from social psychology is a great starting point to systematically approach this problem. The theory provides coverage of all aspects of social relations and equally is concrete and predictive about the visual attributes and behaviors defining the relations included in each domain. We provide the first dataset built on this holistic conceptualization of social life that is composed of a hierarchical label space of social domains and social relations. We also contribute the first models to recognize such domains and relations and find superior performance for attribute based features. Beyond the encouraging performance of the attribute based approach, we also find interpretable features that are in accordance with the predictions from social psychology literature. Beyond our findings, we believe that our contributions more tightly interleave visual recognition

and social psychology theory that has the potential to complement the theoretical work in the area with empirical and data-driven models of social life.

1. Introduction

Today, major part of our social life is captured via social media. As we communicate through multi-modal channels such as Facebook or Twitter, we leave traces that explicitly and implicitly capture social relations in texts, blogs, images and video [4]. As we are approaching a future, where intelligent and potential autonomous systems become our assistants and coworkers, we not only want them to be proficient at their task, but also enable them to blend in and act appropriately in different situations of our – human – life. Additionally, through better understanding about such hidden information we would like to inform users about potential privacy risks.

Both tasks require conceptualizations and models of social life that provide coverage of all aspects of social life and lend themselves to computational models from sensory input. While initial efforts have been undertaken to address this challenging problem, progress is hampered by the overwhelming diversity and complexity of social life. Most of-

ten, computational models to recognize social relations are limited only to a handful of adhoc defined classes.

Therefore, in this paper we start from the social psychology literature to identify a conceptualization of human social life that is holistic and thereby encapsulates all aspects of interpersonal interaction, but at the same time is concrete and specific, so that a computational model can be build on top of this theoretic underpinning. We argue to build on Bugental's domain-based theory [1] which partitions social life into 5 domains, namely: Attachment, Reciprocity, Mating, Hierarchical power, and Coalitional groups. Based on these social domains, we derive a list of 16 social relations. This conceptualization of social life is illustrated in Figure 1 with specific photo examples.

In order to study social domains and relations, we build a dataset based on this domain-based theory. We extend the dataset called People in Photo Album (PIPA) [36] by 26,915 person pair annotations. The label space is hierarchical, by assigning social domain labels that partition the data into 5 domain classes as well as assigning a label for the particular relation that two persons appear to be in.

Finally, we build the first computation models based on this theory that predicts social domain and relation from image data. We analyze fully data-driven models as well as semantic attribute based models whose attributes are based on the domain-based theory. In experiments, we find that these semantic attributes perform best in our dataset and that in addition lead to interpretability of the empirically found concepts of social life. Importantly, our empirical models correlate with the predictions of the domain-based theory.

Our contribution is three-fold: (1) we adapt the social domain theory as a framework for computer vision based analysis of social relations; (2) by annotating a large scale Flickr photo dataset with both relation and domain labels, we provide a more comprehensive dataset than previous work; (3) we collect semantic attributes from human body and head images according to the social psychology study [1]. An in-depth analysis of attribute significance is performed to bridge the gap between social psychology theories and our computational models.

2. Related works

Social relation is a significant part of social network study [4, 12, 13]. This section focuses on the related work in computer vision while the next section outlines different theories in the psychology literature.

Kinship recognition. Relationships among family members are the most basic social relations for human. There exist a large number of studies about family member recognition and kinship verification [24, 25, 26, 27, 28, 29, 30, 31, 32]. Most of these works focus on familial relations: husband-wife, parents-children, siblings, grandparents-grandchildren. Researchers leverage certain

visual patterns exhibited in these relations. For instance, for two people in a wife-husband relation, husband's face is usually in a higher position than wife's [24, 26]. Not only the location information but also the facial appearance, attributes and landmarks are essential features to verify family members. Dehghan *et al.* [31] learn the optimal face features to answer the question of "Do offspring resemble their parents?". Singla *et al.* [25] propose some attribute-related assumptions, *e.g.*, two people of similar age and opposite gender appearing together are spouses.

Based on the social domain definition in [1], familial relations between adults and offspring are in Attachment domain, for which attribute categories such as age, gender and emotion are essential cues. Sibling relation is categorized in Reciprocity domain, which shows more functional and appearance equality than Attachment domain. This is also consistent with the visual pattern of siblings.

Social role recognition in events. In social events, there are immediate social roles and inherent relations among participants. The notion of "social roles" here models the expected behaviors of certain people [14, 15, 16, 17]. For example, in a child birthday party, social roles are birthday child, parents, friends and guests [14]. Instead of immediate roles, we focus on the identity-specified interpersonal relations, which naturally derive permanent social roles. For example, if "leader and subordinate" is confirmed, then it is easy to define the leader's social role as a manager/boss which is much more permanent than "the guest in a party". More importantly, our social relation definition is based on psychology studies that suggest comprehensive social scopes in people's long life.

Social categorization and occupation recognition. Social life endows various social appearances to people. Some research focus on urban tribes in daily life [18], social categories defined by Wikipedia [19, 20], and popular groups such as "Loli", "Syota" and "Goddess" which are mostly derived from social networks [21]. These fine-grained categorization uses body/face positions and attributes such as age, face appearance, hair style, clothing style and so on. Occupation recognition studies [22, 23] not only use personal attributes but also leverage the contextual information in a semantic level, *e.g.*, a waiter is more likely to stand beside sitting consumers in a restaurant.

Relation prediction. Social relation is different with the visual relation of "subject-predicate-object" [33, 34], but it is derived from human social behaviors. A relevant topic is intimacy prediction based on interactive poses [10, 11]. In [10], human poses combined with relative distance, leaning direction and orientation are proposed rich representation. Another interesting work is relation traits estimation by faces [39]. It predicts relation traits such as "warm", "friendly" and "dominant" in face images. Our work is different that we aim to do relation categorization and analyze

social domains covering people's social life. The social psychology basis is introduced in the following section.

3. Social psychology theories

People organize their social life in terms of their relations with other people [3]. The traditional view is that socialization consists of the individuals' learning of principles that can be applied to all social situations [2]. Due to the diversity of social situations, it is almost impossible to define a completed list of social relations. We study social psychological theories to identify a theory that: (1) provides a broad coverage of our social life, (2) is concrete enough to allow deriving relevant social relations, and (3) lends itself for computational modeling and recognition in images and video. After reviewing on related theories, we argue that Bugental's social domain theory [1] is a suitable candidate.

3.1. Social domain theory [1]

Social domain theory [1] partitions social life into 5 social domains and argues that these cover all relevant aspects of our social interactions. Additionally, these domains manifest themselves in concrete social behavior that can be recognized from visual data. Bugental [1] gives comprehensive definitions for each domain including explanations of social cues like appearances and behaviors. Concrete and exemplary social relations are also proposed to illustrate the high-level concept of each domain. While it is illusive to expect that a comprehensive list of all social relations within a domain can be given, the domain partition along with its clear definition serves as a basis to derive social relations from our dataset (see Section 4.1). Specifically, domain definitions and some of the examples given for social relations are as follows:

Attachment domain, characterized by proximity maintenance within a protective relationship, *e.g.* kinship between parents and children. Human attributes such as age difference, proximity and the activity of seeking protection are social cues which can be visually recognizable.

Reciprocity domain, characterized by the negotiation of matched benefits with functional equality between people. Key features are the matched and mutually beneficial interactions in long-term accounting process, which are quite common among friends and siblings. Typically, age difference among peers is small, which is an important semantic attribute. Also, mutual activities such as "gathering" and "sharing" often appear in this domain. Sequenced exchange of positive effect is another factor, which is hard to predict in an image but might be useful when employing video.

Mating domain, concerned with the selection and protection of access to the sexual partner, *e.g.*, the relationship between lovers. Gender cues and the behavior cue of caring offspring are essential for this domain. Bugental also

emphasized the facial attractiveness of prospective partners suggesting that facial and most likely also full body appearance are important cues.

Hierarchical power domain, characterized by using or expressing social dominance. Dominance appears in resource provision and threatening activities. Concrete examples are leaders, powerful peers and teachers. On the other hand, submissive activities like "listening" and "agreeing" are more adaptive for those who lack dominance.

Coalitional groups domain, concerned with the identification of the lines dividing "us" and "them". The focus is on the grouping and conformity cues which ranges from colleagues at work, over sport team members to band members. Coalitional group members often share similar or identical clothing and perform joint activities.

3.2. Related theories

For completeness, we briefly discuss several related social psychology theories ordered by the time of appearance.

(1) Parson's theory of role expectations [8] uses five pattern variables to compose a systematic classification of social relations, namely General-categorical, Personal-categorical, General-behavioral and Personal-behavioral. (2) MacCrimmon and Messick's theory of social motives [9] studies individualism and proposes six motives: Altruism, Cooperation, Individualism, Competition and Aggression. (3) Mills and Clark's theory of Communal and Exchange relations [6], focuses on the rules and expectations governing "give" and "take" benefits. (4) Foa and Foa's theory of resource exchange [7], defines social relations based on six social resources: Love, Status, Money, Goods, Services and Information. (5) Fiske's theory of relational models [3] argues that relations can be differentiated into four parts: Communal sharing, Authority ranking, Equality matching, and Market pricing.

Theories (1)-(3) are rather abstract and theoretical in the field of social psychology and thus not concrete enough for our purpose in computer vision. Theory (4) considers social resources which are concrete but are difficult to infer from visual data. Theory (5) is similar to Bugental's theory but focuses on the cognitive individual experiences reflecting personal history, while Bugental concerned more to link the theory to social behaviors, appearances and environments, which are often visually interpretable. Moreover, Bugental's domain partitioning is based on a large number of social cues (*e.g.* the Table 1 of [1]), which helps to devise computational models for his theory.

4. Social domain and relation data set

To study social domains and relations, we start from the PIPA dataset [36]. The following discusses the dataset, how we have derived the social relations for annotation and presents the annotation process itself.

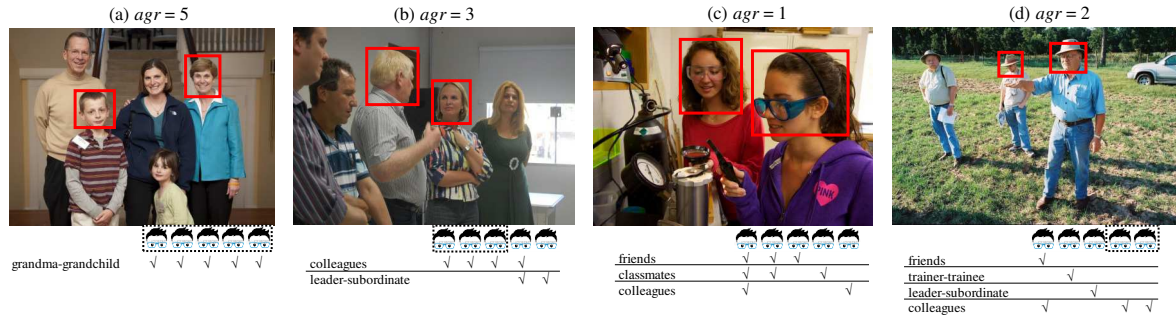


Figure 2: Photo samples of different agreements (agr). Annotators in the dashed box are in agreement. (a) is obvious a family photo and contains grandma-grandchild. In (b), the forth annotator gives a vote to colleagues but is not counted in agreement, since agreement means annotators’ recognitions are exactly the same. (c)(d) contain ambiguous social relations, thus result in very low agreement.

PIPA dataset. The PIPA dataset was collected from Flickr photo albums for the task of person recognition [36]. Photos from Flickr cover a wide range of social situations and are thus a good starting point for our study. PIPA contains 37,107 photos with 63,188 instances of 2,356 identities. For each annotated person, the head bounding box and the identity ID are available. The same person often appears in different social scenarios and interacting with different people which make it ideal for our purpose. Identity information is used for selecting person pairs and defining train-validation-test splits (see Section 6.1).

4.1. Social relation labels

As argued before, the social domain theory is a good basis to derive social relations for annotation. While the domains are precisely described, we still have to obtain a set of relation labels. We proceed in three steps.

- (1) In [1], 12 exemplary social relations are listed for the different domains. We takes these as a starting point.
- (2) As this list is unlikely to cover all social relations in PIPA, we start with a pre-annotation phase of PIPA covering 10% of the data. During this pre-annotation phase, we identify social relations not yet covered and use FrameNet [35] to name the relation and the detailed social domain description to associate the social relation to its domain. This increases the number of relation labels from 12 to 20.
- (3) Finally, We ask five annotators to annotate all of PIPA using the derived list of 20 relations. After the annotation process was concluded, we analyze the statistics and found 4 relations with insufficient occurrences (≤ 20) or insufficient agreement. This results in our final list of 16 relations. For example, kinship like “mother-child” is in Attachment domain, and the relation “leader-subordinate” is in Hierarchical power domain.

4.2. Annotation process

Annotators. Annotating social relations might be subjective and ambiguous. One reason is that a person pair may have multiple plausible relations, as shown in Figure 2. An-

other reason is that the definition of the same social relation might differ, depending on the cultural backgrounds of the annotators. We selected five annotators from Asia, Africa, Europe and America and gave them detailed explanations and photo examples to help them keep some basic consistency (given in supplementary files).

Annotation Protocol. Annotators are asked to individually annotate all person pairs for which we present pairs of head bounding boxes. For each pair the annotator can either pick a relation from our list or, if they are too uncertain, can skip this pair. For example, two people wearing uniforms and working in the factory should be labeled as “colleagues”, as the cues of action “working”, clothing “uniforms” and environment “factory” are obvious. If the annotators are uncertain they are asked to indicate this by clicking “maybe” for this relation. Based on our pre-annotation phase we allowed at most 3 relation labels per person pair which is justified also by the following statistics: 92.3% of the person pairs have 1 relation annotated, 7.5% have 2 relations and only 0.3% have 3 relations.

4.3. Label statistics

After the annotation process we have 26,915 person pairs annotated by five annotators. Given the fact that each annotator can give up to 3 labels per pair or skip a pair, we obtain a total number of 134,556 annotations. For about 8% of annotations, the annotators choose “maybe”. Given that this is a rather small part of the data we include those in the following statistics.

As mentioned before, annotation of social relations using single images might be ambiguous and subjective. Interestingly, for 53% of person pairs at most one relation was chosen across annotators. For 38.8% of pairs two relations, for 7.4% three relations and only for 0.8% four relations are chosen *across* annotators.

Three examples where multiple annotations are chosen are shown in Figure 2 (b)(c)(d). For the image in (d), there are four annotations which are all plausible: the two men might be friends, colleagues, in a leader-subordinate rela-

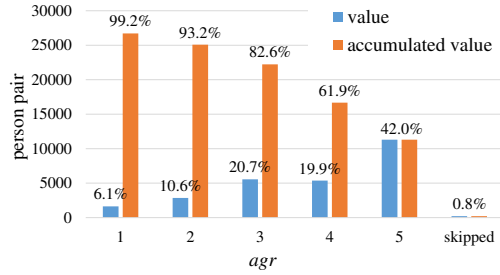


Figure 3: Person pair counting in terms of agreement (*agr*). “Skipped” denotes the pairs skipped by all annotators when they couldn’t recognize any relation. This paper uses annotations with $agr \geq 3$, i.e., useful annotations take 82.6% of all person pairs.

tion, or a trainer-trainee relation. Such highly ambiguous cases however are less prominent in our dataset than one might expect, and for a significant number of person pairs there are at most two relations chosen which indicates that a visual recognition approach is indeed feasible.

4.4. Consistency analysis

We define $consistency \in [1, 5]$ to be the level of complete agreement (*agr*) among the 5 annotators. For instance, $consistency=3$ means $agr \geq 3$ that at least 3 annotators give the exact same labels to a person pair. For examples see figure 2 from left to right: as all annotators give just one and the same label, the first image has $agr=5$; the second has $agr=3$ only, as the fourth annotator not only gives the colleagues relation but also a second relation and is thus not in complete agreement with the first three annotators; the third and the fourth image correspond to $agr=1$ and $agr=2$. It is noted that $agr=1$ is the lowest possible value as each annotator is always in complete agreement with herself/himself.

Figure 3 shows agreement statistics for our dataset. Even though we require complete agreement as discussed above, for 42% of our person pairs $agr=5$. This reinforces the observation that the annotations are less ambiguous than one might have expected. In 19.9% of cases $agr=4$ and in 20.7% of cases $agr=3$. Given these encouraging agreement levels we decided to use the annotations, where $agr \geq 3$ as groundtruth. We refer to it as $consistency=3$ in the following, corresponding to 82.6% of our annotations.

In Figure 4, we show all numbers for photos, identities, person pairs, social relations and domains in terms of $consistency$. As each relation is exactly in one domain, the domain samples is the sum of its relations.

5. Modeling social domains and relations

To model semantic head/body attributes, we choose two image regions based on the groundtruth annotation of PIPA: the head region corresponding to the groundtruth annotation and the body region, defined as $3 \times \text{head width}$, $6 \times \text{head}$

height, following [44]. For recognition, a pair of head or body regions are fed into the model (see Figure 5). We experiment with two types of models: the first type of models are CNN models trained end-to-end; the second type trains CNN models for semantic attributes derived from the social domain theory, then uses the concatenated feature to learn linear SVM.

5.1. End-to-end model

In order to model pairwise relations, we introduce a double-stream CaffeNet, which learns an end-to-end mapping from an image pair to either 5 domain classes or 16 relation classes. This double-stream model is similar to multi-region CNNs [37], based on LeNet [43]. Other similar models can also be considered, such as Siamese-like architectures used for face modeling [41, 38, 39], and multi-channel CNNs used for person identification [40].

The double-stream CaffeNet architecture is shown in Figure 5. Each stream has CONV layers as in CaffeNet [42]. After 5 CONV layers, features are concatenated into one vector that is fed into *fc6* layer. After *fc7* layer, we add an *fc8* to predict either 5 domains or 16 relations.

This model is used to compute baseline recognition results. The same model can be used to predict semantic attributes. This attribute adaption method involves the social cues such as age, emotion and proximity, mentioned in the social psychology article [1].

5.2. Semantic attributes

The second type of model we propose first predicts an intermediate semantic attribute representation and uses this intermediate representation to predict both social domains and relations. In our view there are three advantages that make this representation attractive. First, semantic attributes lend themselves for interpretability. Second, relevant semantic attributes can be derived from the social domain theory and thus can be seen as a way to derive a computational model which in turn allows to link back to the theory. Third, semantic attributes allow to leverage other datasets annotated with such attributes which we consider crucial as it seems illusive to expect large amounts of training data for recognizing fine-grained social relations.

To start, we collected the semantic head/body attribute categories mentioned in the definition of social domains [1]. For example, in the Attachment domain, a mother and her child have a large age difference and some proximity behaviors, therefore, age and proximity are included. Then, for each attribute category, we either leverage existing datasets to train attribute classifiers using our double-stream CaffeNet (default) or use pre-trained models published in previous works [10, 47]. Following presents all attributes used. **Age** infant, child, young, middleAge, senior, unknown are annotated in PIPA by [44]. Additionally, we add small-

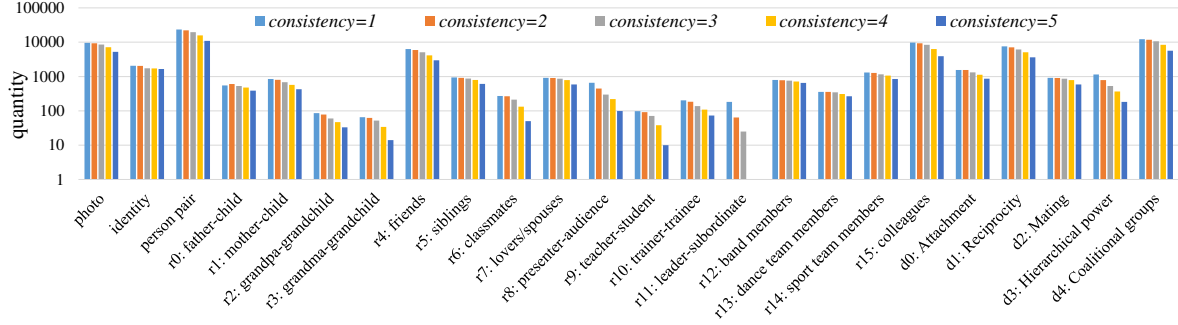


Figure 4: Person pair counting in every social relation (r^*) and domain (d^*) with $consistency \in [1, 5]$. We also list the total counts of photo, identity and person pair. Person pairs with $consistency=3$ (i.e., $agr \geq 3$) are used in our experiments.

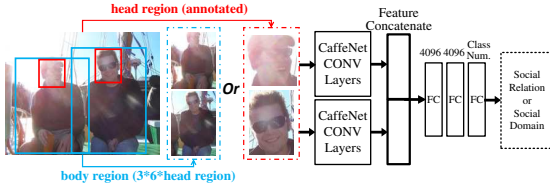


Figure 5: Architecture of double-stream CaffeNet used for modeling social relations and domains. CONV layers are same with those in CaffeNet [42]. Either head image pair or body image pair are fed into the network. The weights of CONV layers are independently trained for double streams.

AgeDiff, middleAgeDiff, largeAgeDiff, because age difference is important for distinguishing social relations [1]. We use head age and body age respectively trained on the head and body regions of PIPA images.

Gender male, female are annotated in [44]. We add sameGender, diffGender. Same with the age case, we use head gender and body gender.

Location & Scale directly collected from head/body regions of a person pair in PIPA. It is composed of 4-dim location coordinates (x, y, width, height), relativeDistance (far, close) and relativeSizeRatio (large, small).

Head appearance 40 classes such as straight hair, wavy hair, wearing earring, wearing hat and so on. This attribute model is trained on the CelebA dataset [46], which contains 202,599 head photos of 10,177 identities.

Head pose, Face emotion. Poses are frontal, left, right, up and down. Emotions are anger, happiness, sadness, surprise, fear, disgust and neutral. Both models are trained on the IMFDB dataset [48], which contains 34,512 head images collected from movies.

Clothing longHair, glasses, hat, tShirt, shorts, jeans, longPants and longSleeves. We learn the model on Berkeley People Attribute dataset [45] containing 8,035 body images, then we extract the features in PIPA using body regions.

Additionally, we utilize the published models of proximity [10] and activity [47] for extracting body features.

Proximity holdingFromBehind, holdingHands, highFive, hug, armOverShoulder, shoulderToShoulder and armIn-

Arm. We use the published Multi-task RNN model [10], which was trained on 10,000 images. We first extract the features for PIPA from its $fconv9$ layer which is a $338 \times 50 \times 50$ feature tensor containing 338 feature maps. In order to reduce the feature dimension, we use a max pooling operator (to select the most significant feature) along the channel axis, and flatten the feature into 2500 dimensions.

Activity 504 activities such as adjusting, ailing, applauding, arranging, attacking, ballooning, baptizing and so on. We use the published CNN-CRF model, which was trained on a dataset of 126,102 images [47]. We extract the features for PIPA of the $fc7$ layer which is a 1024-dim feature vector.

6. Experiments

Our experiments analyze three aspects. The first set of experiments compare end-to-end training with the proposed semantic attribute model to recognize both social domains and relations. Additionally, motivated by the fact that it is illusive to have a comprehensive dataset of all relations within a domain, we also analyze domain recognition in a leave-one-relation-out setting that measures domain generalization performance. Finally, we give insights into our semantic attribute models, in particular in the light of the social domain theory that we base our investigation on.

6.1. Data splits

We introduce two data splitting methods. All-class (AC) splits are used for traditional recognition, where training data cover all classes of social relations (or domains). Single-relation(SR) splits each time leave the data of a relation out of training, then predict the domain label of these data. The goal is to evaluate the model ability to generalize unseen relation classes within a social domain.

All-class(AC) splits. We keep the test split of PIPA dataset unchanged. For train and validation sets, there are strong data imbalance issues on relation classes, as PIPA was originally proposed for person identification [36]. In order to lessen this problem, we take 3 steps: (1) merge the original train and validation sets; (2) build a new validation set using person pairs from 8 random albums; (3) use the

rest data for training. In summary, there are 13,729 person pairs in train set, 709 in validation, and 5,106 in test.

Single-relation(SR) splits. In order to test domain generalization, each time we pick one relation for testing, and randomly divide other data (of 15 relations) into 10 folders by averaging identities. One folder is used for validation, and others for training. Please note, (1) “10-folder identity averaging” can not be exactly reached, because an identity may compose multiple person pairs. We discard a pair when one of its identities has samples in train set, meanwhile, the other one has samples in validation set; (2) there is not training data for Mating domain when leaving lovers/spouses out. Therefore, lovers/spouses data are always in train set. Finally, we have 15 groups of train-validation-test splits, thus we run the model for 15 epoches to get testing results.

6.2. Recognition

To our knowledge, this is the first work to recognize social domains and test domain generalization across relations. Both are challenging problems. The data of a social domain has large intra-variation, due to diverse relations included. Generalization test is even more challenging, since it is to predict the domain of an unseen relation class. This is difficult but very essential to validate whether social domain theory can cover all its relations, and to what extent our model achieves the coverage of social domain.

To evaluate the semantic attribute based model, we have 5 settings¹ as shown in Table 1: (1) end-to-end double-stream CaffeNet (default) trained from scratch; (2) end-to-end finetuned model from pre-trained in ImageNet; (3) Extract features (from *fc7* layer) by the ImageNet pre-trained model, then train linear SVM; (4) Replace the pre-trained model in setting (3) with a finetuned model; (5) Concatenate features extracted from attribute finetuned models, then train linear SVM (ours). In setting (5), we test head features, body features and concatenating both. Except that the initial learning rate 10^{-4} is used for scratch, 10^{-5} is fixed for others. Each time we run 30 epoches of training samples.

MODEL	RELATION	DOMAIN	GENERALIZATION
END-TO-END SCRATCH	34.4%	41.9%	–
END-TO-END FINETUNED	46.2%	59.0%	18.5%
PRE-TRAINED, SVM	35.9%	53.3%	27.7%
FINETUNED, SVM	48.6%	63.2%	27.1%
HEAD ATTRIBUTES, SVM	44.8%	59.4%	21.5%
BODY ATTRIBUTES, SVM	57.2%	67.7%	32.8%
ALL ATTRIBUTES, SVM	57.2%	67.8%	33.3%

Table 1: Accuracies of relation/domain recognition (AC splits), and domain generalization (SR splits). “ALL” means concatenating all body and head attribute features.

“End-to-end finetuned” gets more than 10% improvement over “scratch” for each recognition. However, it fails

¹The dataset and trained models can be downloaded in our project page: www.mpi-inf.mpg.de/social-relation

in the harder task of generalization test (last column), since 18.5% is around chance level of 20%. Using semantic attribute based models, we get the highest 14.8% improvement over “end-to-end finetuned” for the generalization test. In the recognition tasks, our best results are 57.2% and 67.8%, respectively 8.6% and 4.6% higher than best baseline numbers. On the one hand, recognizing relations is much harder due to the larger class number than domains (16 vs. 5). Our attribute model gains a larger improvement for this harder task. On the other hand, the gap between recognizing relation and domain is not very significant (only 10.6%), due to the fact that the intra-variance of a domain is quite larger than that of a relation. In particular, Hierarchical power is the hardest domain to recognize. Its relations such as “teacher-student” and “leader-subordinate” are quite different in both behaviors and appearances.

We can conclude from these improvements that semantic attributes proposed in the social psychology study are very helpful to model high-level social concepts, even though half of the attribute models were trained on other datasets. In the next section, we further analyze our attribute model to get insights into the contribution of specific attributes.

6.3. Analysis of semantic attributes

Firstly, we compare the contribution to the overall performance of each attribute category. Then, we present qualitative examples to understand which detail attributes help to improve the recognition.

Attribute categories. In Figure 6, we present the relative recognition contribution of each single attribute category in relation vs. domain models. Taking body age as an example, its X-Y coordinates is computed as follows: (1) we train a model using only one feature: bodyAge; (2) we evaluate the performance for relation and domain, denoted as $acc(\text{bodyAge}, \text{relation})$ and $acc(\text{bodyAge}, \text{domain})$. Accuracies of using all attributes, 52.7% and 67.8% in Table 1, are denoted as $acc(\text{all}, \text{relation})$ and $acc(\text{all}, \text{domain})$; (3) normalized results of $acc(\text{bodyAge}, \text{domain})/acc(\text{all}, \text{domain})$ and $acc(\text{bodyAge}, \text{relation})/acc(\text{all}, \text{relation})$ are used as X, Y coordinates, respectively.

Overall, we can observe in Figure 6 that most attributes are below the diagonal. This indicates that the relative, individual contribution of attributes is stronger for recognizing domains. We conclude that currently more attributes are needed for the relation classification as it is a more challenging task due to more classes and finer granularity.

In terms of the attribute contribution, activity and clothing are ranked in top 2 for both relation and domain (upper right corner). This is consistent to our social domain interpretation in Section 3.1 that social relations regulate our behavior and “communicate” relations via appearance. For example, band team members and dance team members in Coalitional groups share similar or identical cloth-

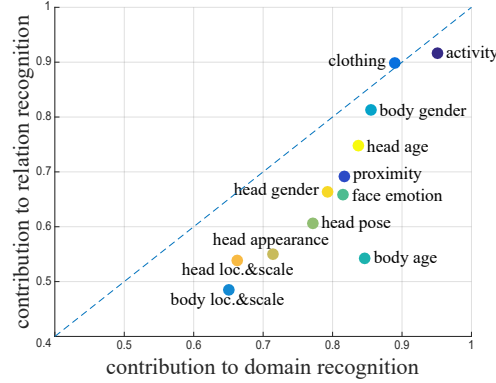


Figure 6: Normalized contributions from 12 attributes.

ing and perform joint activities both in the level of the specific relation and in the level of the shared domain. Although these two attribute models are transferred from other datasets [47, 45], they still make greater influence than age and gender which are learned on PIPA dataset. This reflects that our computational model successfully transfer the social cues from social psychology study to visual data. In terms of the third and forth contributors, age and gender, we can see body gender and head age are relatively important. This is reasonable that in the social psychology definition of Mating domain and Attachment domain, age and age difference, gender and gender difference are dominant features. Another indication is that age is better learned in head images and gender is better learned in body images.

Detail attributes. To figure out detail attributes contributing to social relation recognition, we show some positive and negative examples in Figure 7. Predictions by all-attribute model (our best) and top 4 single-attribute models (ranked in Figure 6) are listed under images.

Among positive examples, (d)(e)(f) are correctly predicted by all-attribute and 4 single-attribute models. These images contain many strong social cues, *e.g.* in (e), “large age difference”, “senior-aged woman”, “little girl”, “daily clothing” and “intimate hugging” contribute to the recognition of grandma-grandchild. It is noted that the person pair in (d) are almost in the hazy background but are correctly recognized, validating the ability of our attribute models to handle tough images. Other positive examples in (a)(b)(c)(g)(h) are correctly recognized by all-attribute model but get noisy predictions by single-attribute models, *e.g.* in (g), clothing model makes a prediction of colleagues, probably because this couple wear unified clothes for taking part in an event. Other strong cues like the intimate activity of “arm on shoulder”, body proximity, age and gender features contribute to the prediction of lovers/spouses when combining all attributes.

Overall, negative examples show more unusual behaviors and confused visual cues than positive ones. For example, the groundtruth of (i) is grandma-grandchild, while the activities of “grandma crawling and trying to hold baby’s



Figure 7: Positive and negative examples of relation recognition. Please note, “positive” and “negative” here refer to the best results using all-attribute model. Image titles are the groundtruth. Predicted relations using either all-attribution model or top 4 attribute models (ranked in Figure 6) are listed under images.

hand” seem very unusual. Although using age or gender finntuned model make correct predictions, activity and clothing are so dominant that the prediction by all-attribute model is wrong. In another example (l), human annotators can recognize the official colleagues in a press conference, but the prediction of our model is failure due to unclear body/head appearances in the image.

7. Conclusion

In this paper, we explore a challenging problem of recognizing social relations in daily life photos. While prior work remained partial in terms of the covered social relations, we argue for a social domain based approach in order to investigate relations covering all aspects of social life. We evaluate recognition performance of social relations, social domains as well as generalization performance of recognizing domain across relations. Our experiments highlight the importance of using semantic attributes, which in turn lends to inspectable models that connect to the social psychology theory. Beyond the presented work, we belief that our work can pave the way to a more empirical study of social relations that is yet grounded and interpretable in the context of social psychology theories.

Acknowledgements

This research was supported by the German Research Foundation(DFG CRC 1223).

References

- [1] Bugental, D.B.: Acquisition of the Algorithms of Social Life: A Domain-Based Approach. *Psychological Bulletin*, Vol. 126, No. 2, pp. 187-219, 2000. 1, 2, 3, 4, 5, 6
- [2] Reis, H.T., Collins, W.A. and Berscheid, E.: The relationship context of human behavior and development. *Psychological Bulletin*, 126(6), pp. 844-872, 2000. 3
- [3] Fiske, A. P.: The four elementary forms of sociality: framework for a unified theory of social relations. *Psychological Review*, 99(4): 689, 1992. 3
- [4] Fairclough, N.: *Analysing discourse: Textual analysis for social research*. Psychology Press. 2003. 1, 2
- [5] Haslam, N., Fiske, A. P.: Implicit relationship prototypes: Investigating five theories of the cognitive organization of social relationships. *Journal of Experimental Social Psychology*, 28(5), pp. 441-474, 1992.
- [6] Clark, M. S., Mills, J.: Interpersonal attraction in exchange and communal relationships. *Journal of personality and social psychology*, 37(1): 12, 1979. 3
- [7] Foa, E.B., Foa, U.G.: *Resource theory. Social exchange*. Springer US, pp. 77-94, 1980. 3
- [8] Parsons, T., Shils, E. A. and Smelser, N. J. (Eds.): *Toward a general theory of action: Theoretical foundations for the social sciences*. Transaction publishers, 1965. 3
- [9] MacCrimmon, K. R., Messick, D. M.: A framework for social motives. *Behavioral Science*, 21(2), pp. 86-100, 1976. 3
- [10] Chu, X., Ouyang, W., Yang, W. and Wang, X.: Multi-task recurrent neural network for immediacy prediction. *IEEE International Conference on Computer Vision*. pp. 3352-3360, 2015. 2, 5, 6
- [11] Yang, Y., Baker, S., Kannan, A. and Ramanan, D., 2012, June. Recognizing proxemics in personal photos. *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3522-3529, 2012. 2
- [12] Barr, J.R., Cament, L.A., Bowyer, K.W. and Flynn, P.J.: Active clustering with ensembles for social structure extraction. *IEEE Winter Conference on Applications of Computer Vision*. pp. 969-976, 2014. 2
- [13] Li, L.J., Shamma, D.A., Kong, X., Jafarpour, S., Van Zwol, R. and Wang, X.: CelebrityNet: A Social Network Constructed from Large-Scale Online Celebrity Images. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 12(1), No. 3, 2015. 2
- [14] Ramanathan, V., Yao, B. and Fei-Fei, L.: Social role discovery in human events. *IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2475-2482, 2013. 2
- [15] Lan, T., Sigal, L. and Mori, G.: Social roles in hierarchical models for human activity recognition. *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1354-1361. 2012. 2
- [16] Shu, T., Xie, D., Rothrock, B., Todorovic, S. and Zhu, S.C.: Joint inference of groups, events and human roles in aerial videos. *IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4576-4584, 2015. 2
- [17] Zhang, J., Hu, W., Yao, B., Wang, Y. and Zhu, S.C.: Inferring social roles in long timespan video sequence. *IEEE International Conference on Computer Vision Workshops*. pp. 1456-1463, 2011. 2
- [18] Murillo, A.C., Kwak, I.S., Bourdev, L., Kriegman, D. and Belongie, S.: Urban tribes: Analyzing group photos from a social perspective. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. pp. 28-35, 2012. 2
- [19] Kwak, I.S., Murillo, A.C., Belhumeur, P.N., Kriegman, D.J. and Belongie, S.J.: From Bikers to Surfers: Visual Recognition of Urban Tribes. *British Machine Vision Conference*, 2013. 2
- [20] Shu, H., Gallagher, A., Chen, H. and Chen, T.: Face-graph matching for classifying groups of people. *IEEE International Conference on Image Processing*. pp. 2425-2429, 2013. 2
- [21] Hong, R., Hu, Z., Liu, L., Wang, M., Yan, S. and Tian, Q.: Understanding blooming human groups in social networks. *IEEE Transactions on Multimedia*, 17(11), pp.1980-1988, 2015. 2
- [22] Song, Z., Wang, M., Hua, X.S. and Yan, S.: Predicting occupation via human clothing and contexts. *IEEE International Conference on Computer Vision*. pp. 1084-1091, 2011. 2
- [23] Shao, M., Li, L. and Fu, Y.: What do you do? occupation recognition in a photo via social context. *IEEE International Conference on Computer Vision*. pp. 3631-3638, 2013. 2
- [24] Wang, G., Gallagher, A., Luo, J. and Forsyth, D.: Seeing people in social context: Recognizing people and social relationships. *European Conference on Computer Vision*. pp. 169-182, 2010. 2
- [25] Singla, P., Kautz, H., Luo, J. and Gallagher, A.: Discovery of social relationships in consumer photo collections using markov logic. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. pp. 1-7, 2008. 2
- [26] Dai, Q., Carr, P., Sigal, L. and Hoiem, D.: Family Member Identification from Photo Collections. *IEEE Winter Conference on Applications of Computer Vision*. pp. 982-989, 2015. 2
- [27] Shao, M., Xia, S. and Fu, Y.: Identity and kinship relations in group pictures. In *Human-Centered Social Media Analytics*. Springer International Publishing. pp. 175-190, 2014. 2
- [28] Xia, S., Shao, M., Luo, J. and Fu, Y.: Understanding kin relationships in a photo. *IEEE Transactions on Multimedia*, 14(4), pp.1046-1056, 2012. 2
- [29] Guo, Y., Dibeklioglu, H. and van der Maaten, L.: Graph-Based Kinship Recognition. *IEEE International Conference on Pattern Recognition*. pp. 4287-4292, 2014. 2
- [30] Chen, Y.Y., Hsu, W.H. and Liao, H.Y.M.: Discovering informative social subgraphs and predicting pairwise relationships from group photos. *The 20th ACM international conference on Multimedia*. pp. 669-678, 2012. 2

- [31] Dehghan, A., Ortiz, E.G., Villegas, R. and Shah, M.: Who do i look like? determining parent-offspring resemblance via gated autoencoders. IEEE Conference on Computer Vision and Pattern Recognition. pp. 1757-1764, 2014. [2](#)
- [32] Zhang, Z., Chen, Y. and Saligrama, V.: Group membership prediction. IEEE International Conference on Computer Vision. pp. 3916-3924, 2015. [2](#)
- [33] Zhang, H., Zawlin, K., Chang, S., and Chua, T. Visual Translation Embedding Network for Visual Relation Detection. IEEE Conference on Computer Vision and Pattern Recognition, 2017. [2](#)
- [34] Lu, C., Ranjay K., Michael B., and Li F.F.: Visual relationship detection with language priors. European Conference on Computer Vision. pp. 852-869, 2016.
- [35] Baker, C.F., Fillmore, C.J. and Lowe, J.B.: The berkeley framenet project. The 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics. Vol. 1, pp. 86-90, 1998. [2](#)
- [36] Zhang, N., Paluri, M., Taigman, Y., Fergus, R. and Bourdev, L.: Beyond frontal faces: Improving person recognition using multiple cues. IEEE Conference on Computer Vision and Pattern Recognition. pp. 4804-4813, 2015. [4](#)
- [37] Zhang, X., Sugano, Y., Fritz, M., Bulling, A.: Appearance-based gaze estimation in the wild. IEEE Conference on Computer Vision and Pattern Recognition. pp. 4511-4520, 2015. [2](#), [3](#), [4](#), [6](#)
- [38] Chopra, S., Hadsell, R., LeCun, Y.: Learning a similarity metric discriminatively, with application to face verification. IEEE Conference on Computer Vision and Pattern Recognition. Vol. 1, pp. 539-546, 2005. [5](#)
- [39] Zhang, Z., Luo, P., Loy, C. C., Tang, X.: Learning social relation traits from face images. IEEE International Conference on Computer Vision. pp. 3631-3639, 2015. [5](#)
- [40] Cheng, D., Gong, Y., Zhou, S., Wang, J. and Zheng, N.: Person re-identification by multi-channel parts-based CNN with improved triplet loss function. IEEE Conference on Computer Vision and Pattern Recognition. pp. 1335-1344, 2016. [2](#), [5](#)
- [41] Taigman, Y., Yang, M., Ranzato, M. A., Wolf, L.: Deep-face: Closing the gap to human-level performance in face verification. IEEE Conference on Computer Vision and Pattern Recognition. pp. 1701-1708, 2014. [5](#)
[5](#)
- [42] Jia, Y., *et al.*: Caffe: Convolutional architecture for fast feature embedding. ACM international conference on Multimedia. pp. 675-678, 2014. [5](#), [6](#)
- [43] LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. In Proceedings of the IEEE 86(11), pp. 2278-2324, 1998. [5](#)
- [44] Oh S., Benenson R., Fritz M., Schiele B.: Person Recognition in Personal Photo Collections. IEEE International Conference on Computer Vision. pp. 3862-3870, 2015. [5](#), [6](#)
- [45] Bourdev L., Maji S., Malik J.: Describing People: Poselet-Based Approach to Attribute Classification. IEEE International Conference on Computer Vision. pp. 1543-1550, 2011. [6](#), [8](#)
- [46] Liu, Z., Luo, P., Wang, X. and Tang, X.: Deep learning face attributes in the wild. IEEE International Conference on Computer Vision. pp. 3730-3738, 2015. [6](#)
- [47] Yatskar M., Zettlemoyer L., A. Farhadi: Situation recognition: Visual semantic role labeling for image understanding. IEEE Conference on Computer Vision and Pattern Recognition, 2015. [5](#), [6](#), [8](#)
- [48] S. Setty, et al.: Indian Movie Face Database: A Benchmark for Face Recognition Under Wide Variations. National Conference on Computer Vision, Pattern Recognition, Image Processing and Graphics, 2013. [6](#)