

# ***Being the center of attention: A Person-Context CNN framework for Personality Recognition***

DARIO DOTTI, Maastricht University, The Netherlands

MIRELA POPA, Maastricht University, The Netherlands

STYLIANOS ASTERIADIS, Maastricht University, The Netherlands

This paper proposes a novel study on personality recognition using video data from different scenarios. Our goal is to jointly model nonverbal behavioral cues with contextual information for a robust, multi-scenario, personality recognition system. Therefore, we build a novel multi-stream Convolutional Neural Network framework (CNN), which considers multiple sources of information. From a given scenario, we extract spatio-temporal motion descriptors from every individual in the scene, spatio-temporal motion descriptors encoding social group dynamics, and proxemics descriptors to encode the interaction with the surrounding context. All the proposed descriptors are mapped to the same feature space facilitating the overall learning effort. Experiments on two public datasets demonstrate the effectiveness of jointly modeling the mutual Person-Context information, outperforming the state-of-the-art results for personality recognition in two different scenarios. Lastly, we present CNN class activation maps for each personality trait, shedding light on behavioral patterns linked with personality attributes.

## **1 INTRODUCTION**

Human personality and its influence on behaviors have been studied extensively in the field of psychology [38]. Independently from the situation or the geographical location [43], humans tend to show consistent behavioral patterns, defined by the psychologists as personality traits. The Big-Five trait model [33] is the most popular one, dividing human behaviors in five general dimensions: Extroversion (i.e. talkative, outgoing), Agreeableness (i.e. considerate, forgiving), Conscientiousness (i.e. efficient, motivated), Neuroticism (i.e. tense, full of worries), and Openness to Experience (i.e. inventive, imaginative). With the recent rise in data usage and data analysis, there has been an increasing interest in the automatic assessment of personality traits using sensory data.

Personality computing [49] is the field in which theories coming from the areas of machine learning and psychology merge to create computational models for personality understanding. However, given the uncertain delineations of human behaviors in different situations, it is extremely challenging to build general systems, which would work in diverse scenarios. As a consequence, recent personality computing systems have mostly focused on context-specific problems (e.g. job screening interviews [37], activities of daily living [18], and work meetings [60]). In contrast, in this study, we focus on building a novel and general architecture for personality recognition suitable for different contexts. In computer vision, contextual information has been shown to improve several challenging tasks such as action recognition [31] and social scene understanding [5]. Building on these findings, we propose a new framework, which maps the mutual relation between individual behaviors and contextual information to personality labels.

Imagine a social scenario in which different interpersonal styles influence the group interaction, for example, individuals that strive to be *the center of attention* will try to actively capture the attention of the rest of the group, resulting in the group acting more passively [47]. By combining

---

Authors' addresses: Dario Dotti, dario.dotti@maastrichtuniversity.nl, Maastricht University, Bouillonstraat 8-10, Maastricht, The Netherlands, 6211 LH; Mirela Popa, mirela.popa@maastrichtuniversity.nl, Maastricht University, Bouillonstraat 8-10, Maastricht, The Netherlands; Stylianos Asteriadis, stelios.asteriadis@maastrichtuniversity.nl, Maastricht University, Bouillonstraat 8-10, Maastricht, The Netherlands, 6211 LH.

Version: Preprint. Under review.




Fig. 1. High level description of the proposed model. Individual motion descriptors as well as two Context descriptors are learned in a novel CNN framework for personality recognition. Individual motion descriptors (red color) indicate the engagement level of every person in the scene, social group descriptors (green color) indicate the engagement level of individuals in conversational groups, and finally, the context proxemics descriptors (purple color) indicate the global attitude of each individual in respect to the others.

information at the individual level with information at the context level, a robust semantic understanding of the situation is perceived. Even if individuals are in a nonsocial scenario (e.g. when an individual is alone at home), it has been shown that people engage with contextual objects as they would engage with other humans (i.e. Anthropomorphism) [19]. Therefore, by examining the interaction between human behaviors and the surrounding scene, important information about human personality can be extracted.

In the field of smart monitoring systems, the affective/emotional aspects of human behaviors are often omitted [16], whereas, it has been shown that actions and movements made by the individuals are influenced by their affective state, as well as their personality attributes [59]. For example, authors in [59] show that the pedestrians' walking behaviors in a public environment can be described by affective attributes. Pedestrians are considered aggressive if their walking path is not influenced by the surrounding crowd, whereas pedestrians are considered conservative if their paths are modified to avoid contact with the crowd. Similarly, authors in [16] introduce the use of personal space as a source of information for surveillance systems. By considering mutual distances and mutual orientations, socially relevant material can be derived for assessing unexpected or abnormal situations. In this paper, we build on these findings to create a behavioral monitoring framework that considers the users' personality in different scenarios, and in this way, we aim to fulfill one of the most important systems design principle: "know the user" [21]. There are several ways in which users' information can be gathered through an interactive system, the standard way is the "active interaction", where the user is actively engaged with the system, for instance by pushing buttons or performing some actions. In the last years, with the advance of the technological power as well as sensors precision, studies have been presented in which the user is not directly engaged with the system, but through measurements and observations, the system is able to predict what the user wants. In this direction, we believe that automatically recognizing

users' personality through "passive interactions" can improve the user's experience as well as can increase the system acceptance [1]. For example, the user does not have to actively declare his/her current affective state, mood, or personality to the system, but the system aims to understand and subsequently make actions adapted to the user's personality [30, 48, 50]. Therefore, by making automatic personality recognition systems more accurate and robust to different contexts, we are improving the system capacity to interact and adapt more like a "human" (i.e. perceiving more subtle human characteristics).

In the last few years, deep neural network (DNN) architectures achieved reliable results in the field of Personality Computing, using video [18], audio [27], as well as multimodal data [20, 54]. However, despite the growing effort, many challenges still remain to be tackled. First of all, datasets providing personality labels are still limited, resulting in ad-hoc methods for specific contexts [2, 18], and, secondly, findings about personality are often too restricted to a research field (i.e. either psychology or computer vision). Thus, interdisciplinary research providing both quantitative, as well as qualitative results, are crucial towards more robust behavior understanding systems. In this paper, we propose a novel architecture based on Convolutional Neural Networks (CNNs) [44] for personality recognition, by studying the interaction between Person and Context information in a general manner, both in a social and a nonsocial context. In Fig. 1, we show the high-level description of the proposed system. Apart from the Person motion Descriptor (red color), computed for every individual in the scene, context information is extracted in different levels of granularity. On the one hand, we compute the Social Group Descriptor (green color) on individuals engaging in close social interactions, where, by social interactions we imply the mutual interaction between two or more people [15]. On the other hand, we compute the Context Proxemics Descriptor (purple color), considering the interpersonal distance between all the individuals in the scene. Additionally, we show that our Context Proxemics descriptor is general enough to be applied also in a nonsocial scenario, encoding human-object interactions. As a result, both the local, as well as the global behavioral cues of an individual are mapped to personality labels.

Our contributions are as follows:

First, we create a novel, end-to-end, multi-stream CNN framework to analyze individual, as well as context information, for personality recognition using video data. This method is robust to different types of data (e.g. RGB or Depth image), different types of camera settings (e.g. fps and resolution), as well as different scenarios (e.g. social as well as nonsocial situations).

Second, in addition to individual motion descriptors [24], that are transferred from the activity recognition field, novel CNN descriptors are proposed to encode the surrounding context in different scenarios. Transforming different sources of information (e.g. Person-Context) into CNN descriptors with the same backbone structure has the great advantage of facilitating the discovery of common latent representations.

Third, to demonstrate the generalizability of our algorithm, experiments are carried out on two public datasets. Results show that our Person-Context model outperforms the state-of-the-art methods. Additionally, to demonstrate the robustness of the learned personality patterns, we evaluate our framework using two different sets of personality classes as training labels (e.g. personality traits [33], as well as personality types [9]).

Fourth, by visualizing the CNN activation map for both high and low scores of each personality trait, we pose the following questions: 1) *Is the relation between Person-Context dynamics reflected in the traits?* and 2) *Do the dynamics correspond to the trait attributes defined by psychologists?* Our qualitative results provide new insights into the interaction between context and individuals' behavioral cues for behavior and personality understanding.

## 2 RELATED WORKS

In this section, we describe works that are closely related to ours, in both the fields of computer vision and personality computing.

### 2.1 Skeleton motion features

Human motion recognition is one of the most important challenges in the computer vision community, because of its great applicability in several real-world challenges, such as video surveillance and activity recognition [52]. In the past years, a multitude of methods have been proposed spanning across RGB-based features, Depth-based features, RGB-Depth-based features and skeleton-based features, combined with a plethora of Deep Neural Network (DNN) models [52]. Since in this paper we utilize skeleton-based features, we will mainly focus on this type of information.

Skeleton joint estimation has shown promising results in the tasks of pose estimation [12, 61] and activity recognition [24, 28]. However, by using only coordinates from skeleton joints, a considerable portion of RGB information is discarded, limiting the applicability of CNN methods. To overcome this issue, in the last years, novel strategies have been proposed to create feature representations similar to images [18, 24, 28]. Authors in [24] proposed a spatio-temporal representation of skeleton joints called “image clips” for action recognition. These features allowed the use of the VGG pre-trained model [45] for feature extraction. Similarly, in [28], a skeleton image of varying dimensionality is constructed to take advantage of the 2D CNN receptive field. These approaches brought a new way of encoding the skeleton joints, showing promising results for action recognition.

The multitude of methods, as well as the important amount of available data proposed for human motion recognition opens the way to transfer the acquired knowledge to personality computing. For this reason, in this paper, we chose to explore the relation between skeleton motion features and personality labels, due to their compact representation, as well as robustness to occlusions or inter-personal variations. We show that our approach, inspired by [24], and optimized for a CNN architecture, outperforms the state-of-the-art for personality recognition.

### 2.2 Personality recognition

Several datasets for personality recognition have been proposed in the recent years, both in controlled scenarios [13, 37], and real life scenarios [2, 18]. These datasets are recorded using video data [18], multimodal data ( video and audio [37] or video and wearable sensors [2]), as well as neuro-physiological data [34].

A popular approach towards personality computing is to link face and audio features to personality traits [54], and recently, authors in [62] showed that facial and audio features can be used to link personality and emotion cues. However, systems that apply face analysis are not robust to camera positions, as well as to privacy restrictions [10]. To overcome this obstacle, it has been shown that body motion patterns are also a major indicator of personality traits [2, 18]. Authors in [2] found that high extroversion trait is highly correlated with body postures predisposed towards face to face conversations, while [25] showed that upper body motion of public speakers is a good predictor for personality. Body pose features in combination with Deep Boltzmann Machines (DBMs) proved their efficiency at discovering the emergent leadership in a small group [7]. In [18], posture dynamics are learned using an Autoencoder-LSTM architecture, discovering motion patterns for personality labels. In this direction, we propose a model for personality recognition, to discover discriminative nonverbal behavioral cues for each personality trait.




Fig. 2. Framework architecture. The Person-Context streams (a, b, c) are extracted from video sequences (0) and processed separately (1, 2, 3). The output descriptors are fused in FC layers (4) and finally, a softmax layer is employed for personality recognition.

### 2.3 Social and nonsocial interaction

The analysis of social interactions is an important challenge for several applications, like video surveillance [17] and activity recognition [8]. Social interactions are rich in information, revealing characteristics of both single individuals as well as groups. In [8], authors proposed an LSTM based system to model, at the same time, individual and group activities. In [51], an LSTM system is proposed to model intra-group dynamics (person within a group) and inter-group (group to group) interactions. When it comes to nonsocial interaction, human-object interaction has been studied by several authors. In [31], the correlation between actions and context has been explored, showing how actions are constrained by certain scenes. In [58], the mutual context between human poses and objects has been modeled for Human-Object Interactions. Following these studies, this paper studies the relation between Person-Context interaction dynamics and personality behavioral patterns.

## 3 PERSON-CONTEXT FRAMEWORK

In this paper, we jointly model the nonverbal behavioral cues from single individuals, with surrounding context information, for personality recognition using video data. Theoretical models such as the Laban Movement Analysis (LMA) [22] conceptualize the relation between human motion and the surrounding space. Although this model was proposed initially to describe dance movements, authors of [39] and [40] adopted certain concepts for Human to Human and Human-Robot Interactions understanding. Moreover, while psychological studies have demonstrated the tight relation between contextual information and personality patterns [26], works considering both sources of information are quite limited in the computer vision community. Therefore, inspired by these interdisciplinary research themes, we propose an automatic personality recognition system using a person-context deep model.

### 3.1 Architecture

The proposed architecture is shown in Fig. 2. For every frame  $I^t$  in a given sequence, we first run a pose estimation algorithm [12], with the role of defining semantic components of the scene (Fig. 2

(a), (b) and (c)). In case some of the body joints are not detected correctly, we estimate their positions using the neighbours' joints coordinates. As the algorithm proposed by [12] does not track the detected objects, a frame by frame tracking algorithm is added [11]. Given the detected skeletons  $p_{subjs}$  ( $subjs \in 1, \dots, n$ ) where  $n$  is the total number of subjects in a social  $sc^s$  or nonsocial scene  $sc^{ns}$ , we aim to map Person as well as Context information to personality labels. One disadvantage of jointly modeling different sources of data is that their intrinsic structure is too diverse, making the learning process more complex. Differently, one key aspect of the proposed method is to use the same backbone structure (i.e. cylindrical coordinates descriptors, Fig. 2 (1)), to describe the various types of information. In this way, different feature modalities are mapped onto the same feature space, allowing the creation of the same network structure (Fig. 2 (2) and Fig. 2 (3)). Note, the extracted components (Fig. 2 (a), (b) and (c)) are general enough to deal with nonsocial scenes  $sc^{ns}$ , where, the Social Group Interaction (Fig. 2 (b)) will be set to 0, and the Context Proxemics Descriptor (Fig. 2 (c<sup>ns</sup>)) will encode interactions with semantic objects in the scene. Finally, Fully-Connected Layers (FC) are adopted to fuse the features extracted by the distinct streams, and a Softmax layer is added for the personality recognition task. As follows, detailed explanation of every component is provided: the Person descriptor is described in Sec. 3.2, then, novel Context CNN descriptors are presented in Sec. 3.3 and Sec. 3.4. The learning component is introduced in Sec. 4, and, finally, quantitative as well as qualitative results are reported in Sec. 5, Sec. 6 and Sec. 7 respectively.

### 3.2 Person motion

Nonverbal behavioral cues have been studied extensively as personality patterns descriptors [18, 25]. In this section, skeleton pose information, extracted using [12], is employed to create the "Person" descriptor (Fig. 3(a, 0)), in the proposed Person-Context framework. Inspired by [24], we create temporal "skeleton clips", to describe the relative motion of each skeleton joint  $J^{joints}$ , where  $joints$  is the number of skeleton joints detected. In particular, for every frame in a sequence  $t$ , the relative position of all skeleton joints is computed with respect to the selected reference joints  $J^{ref}$ , where  $ref$  denotes the number of reference joints. Similar to [24], we set  $ref = 4$ , resulting in  $ref$  matrices of size  $t \times J^{joints-1}$ .

To make our descriptor more compact, unlike [24], the computed distances are stacked vertically (Fig. 3(a, 1)), to form a final matrix that describes along the vertical axis, the temporal evolution of the joints in a sequence  $t$ , and along the horizontal axis, the relative distance of the skeleton joints in one frame. Note that, the relative distances are computed for each coordinate independently, resulting in a 3D matrix. Since the joints positions are declared in the Cartesian coordinates, following [24] and [55], we transform the 3D temporal matrix into cylindrical coordinates (Fig. 3(a, 1)). Authors in [55], created motion features for action recognition using cylindrical coordinates to obtain features invariant to location, orientation and scale. This is particularly important in unconstrained multi-camera datasets, like the Salsa dataset [2], in which the scene is crowded, and occlusions make the skeleton detection challenging.

$$(\rho, \theta, z) = (\sqrt{x^2 + y^2}, \tan^{-1} \left( \frac{y}{x} \right), y_1 - y_2) \quad (1)$$

Cylindrical coordinates (Eq. 1) are composed by three terms:  $\rho$ ,  $\theta$ , and  $z$ . The term  $\rho$  is defined as the euclidean distance between the reference joints  $J^{ref}$  and the observed joints  $J^{joints}$ , the term  $\theta$  is defined as the angle between the reference joints  $J^{ref}$  and the observed joints  $J^{joints}$ , and the term  $z$  is the difference considering only the vertical axis between the reference joints  $J^{ref}$  and the observed joints  $J^{joints}$ . Finally, the three descriptors values are converted between 0 – 255, suitable for a CNN architecture.

### 3.3 Social Group Motion

Given a social scene  $sc^s$ , the modeling of social group interactions [16] is a challenging task for behavior understanding. Since, interaction dynamics are affected by the behaviors of single individuals involved in the group, the understanding of each personality is of great importance to unfold the social group evolution. For example, an extroverted person will tend to be actively involved in the group, while an agreeable person will tend to be more passive [2]. In this section, we explain how the social group motion is encoded using the cylindrical descriptors (Eq. 1).

The code provided by [15], for social interaction discovery during social events, is used to detect the members of each social group in the scene. We define each member as  $m_M^{group}$ , where  $M$  is the total number of individuals in a given social group *group*. Since social groups dynamics are affected by single individual's motion, we compute the motion dynamics (Sec. 3.2) for every member  $m_M^{group}$ . However, this introduces a new challenge: groups vary the number of members frequently, resulting in descriptors of different sizes. Therefore, a pooling strategy is adopted to encode the social motion dynamics of each group in a compact representation. This strategy ensures that the obtained features are more robust to the effect of group dynamics variation, while still preserving important data patterns.

**3.3.1 Social Motion Pooling.** In order to integrate the motion dynamics of  $m_M^{group}$ , we focus on the two most conventional pooling operations, max pooling and average pooling. Choosing the right pooling strategy is important to “downstream” the descriptors to a fixed dimension independently from the social group size, as well as encoding the group motion dynamics.

Pooling operations are performed on the three cylindrical descriptors independently  $(\rho, \theta, z)$ . Average pooling is defined as  $f_{ave}(x) = \frac{1}{M} \sum_{m=1}^M \sum_{i=1}^N x_{m,i}$ , and max pooling is defined as  $f_{max}(x) = \max_{m,i} x_{m,i}$ , where, the vector  $x$  contains the activation values from pooling regions  $i = [1, \dots, N]$ , computed on every member of the social group  $m = [1, \dots, M]^{group}$ .

An exploratory experiment is conducted to test the performance variation when using the described pooling strategies on the Salsa dataset poster session [2]. In particular, we aim to analyze the information carried by the social group motion descriptor in the task of personality recognition. In Fig. 4(a), accuracy results for the classification of the big 5 personality traits, using the described pooling strategies are visualized. The results indicate that average pooling constantly reaches higher accuracy, and therefore is selected as the social group pooling strategy in the rest of the experiments. However, it is interesting to notice that the performances are not statistically significant (indicated by the asterisks), when predicting two personality traits: Conscientiousness and Neuroticism.

Average motion pooling encodes the motion from all the members in the social group, whereas max pooling encodes the highest motion value in the group. As individuals with high Conscientiousness or Neuroticism traits are less likely to stand out or influence the social group, the average social group motion dynamics may acquire more importance than the individual motion, and therefore confuse the classifier decision.

### 3.4 Context proxemics

Human behaviors are often correlated with the environment they are performed in, hence, in this section, proxemics descriptors encoding the interaction with different contextual information are introduced depending on the analyzed scenario. For example, during a social event, it is likely that individuals show social behaviors, whereas, when individuals are at home alone, it is likely that behaviors are related with objects in the scene. In the next sections, two strategies are proposed towards encoding proxemics to different entities.

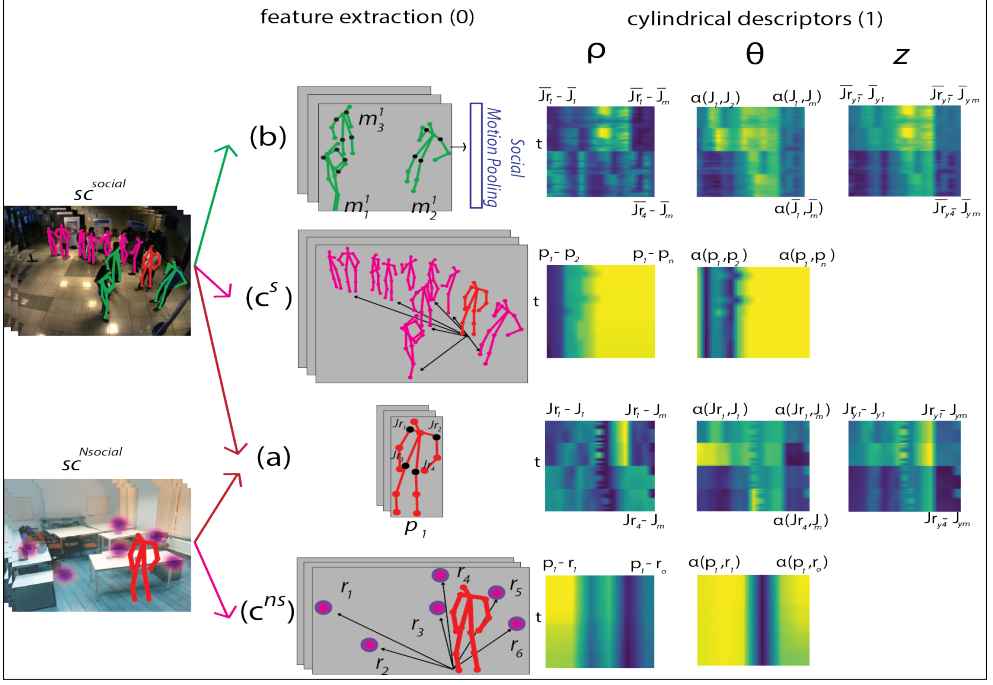


Fig. 3. CNN descriptors of the proposed framework: (a) Person motion descriptor, (b) Social group motion, ( $c^s$ ) Social proxemics, and ( $c^{ns}$ ) Nonsocial proxemics. The descriptors values (0 – 255) range in color from blue (low distances between skeleton joints) to yellow (large distances between skeleton joints).

**3.4.1 Social Proxemics.** Several works highlighted the correlation between interpersonal distances (i.e. proxemics) and personality traits [60], [15] in social scenarios. Thus, in this section, a proxemics descriptor encoding the global position of  $p_{subjs}$  in respect to all the other individuals in the scene, is computed.

In the analyzed scene  $sc^s$ , proxemics is intended as the way people use their personal space in relation to other individuals. The process is illustrated in Fig. 3 ( $c^s$ , 0), where we compute  $\rho$  and  $\theta$  descriptors between  $p_{subjs}$  and the rest of the individuals using Eq. 1, for a frame sequence of length  $t$ . To avoid dimensionality discrepancy due to the different number of individuals in the scene, we always consider the maximum number  $n$  of subjects for a given dataset. Note that the  $z$  descriptor is not considered, as the interpersonal distance does not occur on the vertical axis. Thus, two descriptors of size  $t \times n - 1$  are obtained, and to be consistent with the other generated features, they are transformed into grey images by scaling the values between 0 – 255 (Fig. 3 ( $c^s$ , 1)).

These descriptors have two advantages: Firstly, the global position of every subject with respect to all the other subjects in the scene is represented, reflecting behaviors that have been already correlated with personality display [2]. For example, individuals who like to be engaged in group conversations (Agreeableness personality trait) will have a higher number of neighbours, while on the other hand, shy individuals will have fewer neighbours. Secondly, due to the common cylindrical coordinates structure, these descriptors can be fused with the other generated features, optimizing the learning capacity of our CNN architecture.

**3.4.2 Nonsocial proxemics.** As humans are by nature social beings, personality displays have been well explored during social interaction. However, as reported by [18], few efforts have been made



towards the understanding of personality in a nonsocial context. This problem has been shown to be important for applications in domains like Ambient Assisted Living (AAL) and Smart Homes, where it often occurs that individuals are spending their time at home alone. Thus, in this section, we aim to learn the relation between human behaviors and the surrounding nonsocial context, defining proxemics as the way people use their personal space in relation to objects (Fig. 3 ( $c^{ns}$ )).

We build on the philosophical and psychological theorem which states that people engage with objects as they engage with other humans (e.g. Anthropomorphism [19]). In this section, a novel descriptor is proposed to capture the way individuals engage with the scene. As shown by [19], especially in case of “lonely” situations, subjects are more likely to be subjected to anthropomorphism with nonhuman elements like objects, robots etc.

Since no a priori information is given about the scenario, firstly, semantic regions of the scene are discovered in an unsupervised way and, secondly, we compute a temporal descriptor optimized for a CNN architecture, encoding Person-Context interactions. Spatial information is extracted by dividing the scene into fined-grained patches. In every patch, we assume that the arm motions are the most important joints for person-context interaction location. Therefore, we quantize motion, as well as orientation information of the arms in a histogram of oriented tracklets descriptor [35]. However, movements such as walking generate motion from the arms but are not interesting for our Person-Context Interaction investigation. To avoid this, the average body motion is subtracted from the arm motion. To locate the regions of the scene where the interaction occurs in an unsupervised way, Gaussian Mixture Model (GMM) is applied [6]. GMM clustering is a method which attempts to find a mixture of multidimensional Gaussian probability distributions that best model the input data. Because GMM is a probabilistic model, we can treat the clustering task as “soft” assignment, being more robust to outliers and noisy data. In Fig. 3 ( $c^{ns}$ ), the discovered semantic regions  $r^{regions}$  are shown, where  $regions = 6$ .

To extract descriptors that reflect the engagement of individuals in the generated semantic regions, we adopt the same strategy as in Sec. 3.4.1, computing  $\rho$  and  $\theta$  (Eq. 1) descriptors between  $p_{subjs}$  and the generated regions, for every frame sequence  $t$  (Fig. 3 ( $c^{ns}$ , 0)). Finally, two descriptors of size  $t \times regions$  are obtained, and transformed into grey images by scaling the values between 0 – 255 (Fig. 3 ( $c^{ns}$ , 1)).

## 4 PERSON-CONTEXT INTERACTION LEARNING USING A CNN ARCHITECTURE

The integration of multiple feature representations, is a challenging task due to the heterogeneity of their distributions [63]. However, CNN architectures showed great ability in discovering common latent representations. Therefore, in this paper, a CNN framework is proposed for the fusion and learning of features extracted from multiple sources: individuals’ motion (Sec. 3.2), social group motion (Sec. 3.3) and context proxemics (Sec. 3.4). In order to reduce the structural difference of different distributions, a common feature representation based on cylindrical coordinates is created (Fig. 2(1)). This strategy allows the CNN model to use the same parameters (e.g. number of hidden units, pooling layers etc.) for all the feature representations, decreasing the duration and complexity of the training phase (Fig. 2(2), Fig. 2(3), Fig. 2(4)).

### 4.1 Feature Learning

The temporal resolution of each descriptor varies depending on the fps of the analyzed dataset ( $t \in 15, 30$  frames), while, in order to be fed to the CNN network, their final size is scaled to  $(68 \times 68 \times 3)$ . This dimensionality is chosen to preserve the mapping of the skeleton joints indexes on the original vector, with the rescaled descriptors. This is particularly useful, for example, for the visualization of the CNN layers (Sec. 7). Given the CNN descriptors as input, the VGG19 model [45] pre-trained on ImageNet [41], is used as a first step in the learning phase. Note that we utilize




Fig. 4. (a) Exploratory experiment investigating the recognition accuracy using two different pooling strategies, encoding the social groups’ motion. Average pooling demonstrates better performances over all the big 5 traits, however the results are not significant over two traits: Conscientiousness and Neuroticism. The p-values less than 0.001 are summarized with three asterisks, p-values less than 0.01 are summarized with two asterisks, and finally, no asterisks show not significant p-values. (b) Exploratory experiment investigating the recognition accuracy using different fusion methods. Feature concatenation shows better results than the other methods, being significantly superior on four out of five traits. The p-values less than 0.05 are summarized using one asterisk.

the VGG19 model only for feature extraction discarding the last 3 fully-connected layers. Early convolutional layers showed to learn more generic features, whereas, deeper layers are more influenced by the task they are trained for [24]. Since our CNN representations are very different from the images contained in ImageNet, we extract the features from the convolutional layers  $conv5_1$ . We justify the adoption of this layer, as authors in [24] applied it on similar image clips. Specifically, the generated cylindrical descriptors are fed as input to the VGG model. As output, from the  $conv5_1$  we obtain feature representations of size  $4 \times 4 \times 512$  (512 feature maps of size  $4 \times 4$ ). Since the VGG feature extractor projected the Person-Context information in a lower-dimensional feature space, a pooling strategy is adopted to encode semantic spatio-motion features. Specifically, as in [24], we apply temporal mean pooling (TMP) to extract temporal information from the sequences.

After the pooling step, the local features, defined as  $D_{pers}$ ,  $D_{group}$ , and  $D_{prox}$  in Fig. 2 (3), are fused to form the input to two Fully-Connected Layers (FC) and the Softmax layer for the final classification. Between the two FC layers there is a rectified linear unit (ReLU) [36] to introduce an additional non-linearity.

**4.1.1 Fusion.** The fusion of different input cues can improve the recognition performance, as it combines different sources of information that are relevant to discover personality patterns. In our framework, fusion can be performed as early fusion (feature fusion) or late fusion (decision fusion) [46].

As feature fusion, we implemented the following methods: 1) concatenation of the features coming from different modalities. 2) Principal Component Analysis (PCA) [57] was applied on the features extracted from different modalities, for dimensionality reduction, as well as to remove correlations between the features. Following [4], 98% of variance was kept. 3) PCA was applied on the concatenated features. As decision fusion, the network was trained on the three descriptors (e.g.  $D_{pers}$ ,  $D_{group}$ , and  $D_{prox}$ ) independently and the recognition decision was fused (sum-rule) [4]. Fig. 4(b) shows the personality traits recognition results using the described fusion methods on

Salsa dataset poster session. The concatenation method exhibits the best results on all the big 5 traits, being significantly higher on four out of five traits (p-value less than 0.05 is summarized using one asterisk). Therefore, this method is used for the reminder of the experiments.

## 5 EXPERIMENTS

In order to examine the strength of our framework, in this section, we present personality recognition experiments on two public datasets (Sec. 5.1). The choice of these two datasets justifies the evaluation of the system in two different scenes, where, depending on the scenario, the deep Person stream is combined with distinct Context streams (Fig. 3). The experiments are organized as follows: in Sec. 5.3, we compare our work with state-of-the-art systems for personality recognition, and in Sec. 5.4 we evaluate our framework using the Big-Five personality traits as labels, and finally, qualitative results are explained in Sec. 6.

### 5.1 Datasets

The Salsa dataset [2] contains multimodal data from two social events (30 minutes each) in an unconstrained indoor scenario. Video data was recorded from 4 cameras placed at each corner of the room, and ground truth positions of the subjects' movements were provided every 45 frames. It consists of two parts, the first part was recorded during a poster presentation session, and the second one was recorded during the coffee break, where all the participants were allowed to freely interact with each other (this part is named cocktail party). The two parts contain the same participants and their personality scores were collected using the Big-Five personality questionnaire [23]. The personality in a nonsocial context [18] is a recent dataset, providing video data of 45 participants in an unconstrained indoor scenario. Every subject performed six tasks resembling Activities of Daily Living (ADL) and filled the short version of the Big-Five personality questionnaire [38].

### 5.2 Experimental Setup

Since the research in [18] is the most related to our work, the same experimental procedures are followed. In particular, the personality trait scores from the two datasets are normalized (between 0, . . . , 1), and a hierarchical clustering technique [14] is applied to explore the traits configuration. Confirming the findings in [18],  $clusters = 3$  are found. As shown in Fig. 5(a), the traits scores are grouped in three main clades (colored in green, red and azure). Then, in order to assign a semantic meaning to the discovered clusters, we compute the average score of each personality trait in the different clusters.

In Fig. 5(b), we show the obtained results, which are consistent with the psychological theory proposed by [9], stating that the Big-5 personality traits can be organized in three major types: undercontrolled, overcontrolled and resilient [18]. In Fig. 5(b), resilient personality type (green color) shows high Extraversion score and the lowest score in Neuroticism, the undercontrolled type (red color) scores high in Extraversion as well as Neuroticism, and finally, the overcontrolled type (azure color) has the lowest score in Extraversion and scores high in Neuroticism.

These new classes are used as labels in the personality recognition experiments. Moreover, regarding the classification task, using personality types labels, two testing approaches are carried out: "K-cross validation" where  $K = 10$  and "leave-one-out" (specifically, the data from one subject is left out from the training procedure and used only for testing).

**Ablation Study.** In order to test the contribution of every component in the proposed framework, different combinations are tested. For the Salsa dataset, we tested:  $D_{pers}$ , where only the Person motion stream is used.  $D_{prox} + D_{group}$ , where only the two context streams are used.  $D_{pers} + D_{prox}$ , where person motion (Sec. 3.2) is combined with social proxemics descriptors (Sec. 3.4.1),  $D_{pers} + D_{group}$ , in which the social group motion descriptors are used (Sec. 3.3), and finally,




Fig. 5. (a) Dendrogram of the hierarchical clustering applied on the normalized personality trait scores from the used datasets. Three main clades (coloured in green, red and azure) are found. (b) Average values for each personality trait in the discovered personality types (undercontrolled personality type colored in red, resilient personality type colored in green, and overcontrolled personality type colored in azure).

$D_{pers} + D_{prox} + D_{group}$ , where all the descriptors are fused. As for the nonsocial context, we tested:  $D_{pers}$  stream alone, and  $D_{pers} + D_{prox}$  encoding the person motion with the scene proxemics (Sec. 3.4.2).

**Baseline Methods.** We compare our framework against the three most related works. In particular,  $LSTM_{cl}$ , developed by [18] was initially proposed for personality recognition on the nonsocial dataset. *Clips + MTLN* was proposed by [24], and it uses similar skeleton descriptors as input to a CNN framework called MTLN. Finally, EL-LMKL [7] was proposed for leadership recognition as well as extraversion trait recognition on a leadership Dataset [42]. Please note that all the baselines were implemented by the authors, as to the best of their knowledge, there do not exist personality recognition results on the Salsa dataset.

### 5.3 Personality types recognition

Personality recognition results obtained on the two parts of the Salsa dataset, as well as on the nonsocial dataset, are shown in Table 1 and Table 2 respectively. Overall, the proposed method is able to reach the best accuracy results in both datasets, demonstrating the benefit of fusing Person-Context information. In Table 1 we report the results on the two parts of Salsa dataset (poster and cocktail party sessions) for both testing approaches (Cross-validation and Leave-One-Out). For the results on the poster session part, the ablation study indicates that, although the fusion of all the context descriptors ( $D_{pers} + D_{prox} + D_{group}$ ) obtains the highest accuracy, it does not improve significantly the result obtained using only  $D_{pers} + D_{group}$ . This may be explained by studying the scenario of the dataset, which displays interactions during a scientific poster session. During this type of sessions, individuals have to respect some spatial constraints, for example, if a poster presentation is already crowded, individuals cannot position themselves freely. Hence, in this situation, the proxemics descriptor does not always describe affective behaviors as explained by [15]. Additionally, note that  $D_{pers}$  descriptor alone reaches higher accuracy than  $D_{prox} + D_{group}$ , supporting the explanation that proxemics features in this scenario does not show a clear link to personality patterns. On the other hand, the ablation study on the cocktail party part demonstrates the advantage of considering all the descriptors, where the three proposed descriptors ( $D_{pers} + D_{prox} + D_{group}$ ) obtain the best accuracy, significantly improving all the other descriptors' combinations. As this part of the dataset depicts the participants freely interacting with each other,

Table 1. Experiments on the Salsa dataset for Personality Recognition. CV= Cross-validation, LOO= leave-one-out. The p-values less than 0.001 are summarized with three asterisks, p-values less than 0.01 are summarized with two asterisks, p-values less than 0.05 are summarized with one asterisk, and finally, no asterisks show not significant p-values.

Methods	Salsa Poster CV	Salsa Party CV	Avg. CV	Salsa Poster LOO	Salsa Party LOO	Avg. LOO
<b>Baseline</b>						
$LSTM_{cl}$ [18]	65.7%	67.85%	66.77%	65.8%	53.38%	59.59%
EL-LMKL [7]	75%	74.75%	74.87%	73.6%	48.9%	61.25%
Clips+MTLN [24]	78.7%	74.5%	76.6%	78.1%	58.86%	68.48%
<b>Proposed</b>						
$D_{pers}$	76%	71.3%	73.65%	76.4%	59.66%	68.03%
$D_{prox} + D_{group}$	73.6%	75.83%	74.71%	75%	59.13%	67.06%
$D_{pers} + D_{prox}$	78.2%	75.61%	76.9%	78%	62.4%	70.2%
$D_{pers} + D_{group}$	79.6%	74%	76.8%	79%	64.93%	71.96%
$D_{pers} + D_{prox} + D_{group}$	<b>79.8%</b>	<b>77.25%*</b>	<b>78.52%*</b>	<b>79.2%</b>	<b>66.8%***</b>	<b>73%**</b>

social behaviors are more natural and less constrained by social roles (i.e. poster presenter vs. audience). While natural behaviors may be more representative of a certain personality type, they are more complex to model, and as consequence, the overall recognition accuracy is lower. The complexity in this scenario affects all the parts of the framework, as free interaction and group formation create occlusions and clutter, challenging the tracking and motion modeling components. Finally, the results show that the fusion of all the descriptors is more robust to real-world cluttered scenarios than individual descriptors.

Furthermore, the statistical significance of the results over the two parts of Salsa dataset is computed in respect to the baselines (indicated by the number of asterisks in Table 1). As the poster session presents more static social groups that are not formed by free interactions but imposed by the rules of a poster session event, the combination of Context and Person information does not add a statistical improvement (p-value > 0.05) in comparison to single motion information computed using [24]. On the other hand, the results obtained on the second part of the dataset (i.e. cocktail party) show significance difference (p-value < 0.05) demonstrating the benefit of fusing Person-Context information during social interactions. Finally, our experiments show that there exist many factors interfering with finding the connection between human behaviors and personality labels, such as within-person variability [53] in different situations (i.e. poster session versus cocktail party session). For this reason, we also show the personality recognition results by averaging the two sessions (78.52% for cross-validation and 73% for leave-one-out testing), which are significantly different than the baselines. Concerning the nonsocial dataset (Table 2), our method obtains the highest results, proving the benefit of combining Person-Context information also in a nonsocial environment. Overall, our results pave the way for future studies exploring the connection between contextual information and nonverbal behaviors for personality recognition from a psychological as well as computational point of view.

Table 2. Experiments on the Nonsocial dataset for Personality Recognition. CV= Cross-validation, LOO= leave-one-out. P-values less than 0.05 are summarized with one asterisk.

Methods	Nonsocial CV	Nonsocial LOO
<b>Baseline</b>		
$LSTM_{cl}$ [18]	66.4%	55.3%
Clips+MTLN [24]	70.9%	70.7%
<b>Proposed</b>		
$D_{pers}$	69.4%	69.2%
$D_{pers} + D_{prox}$	<b>72.4%*</b>	<b>72.6%*</b>

Table 3. Personality Recognition using Traits scores, E= Extraversion, A= Agreeableness, C= Conscientiousness, N= Neuroticism, OE= Open to Experience

Personality Traits	Salsa Poster	Salsa Party	Nonsocial
E	86.6%	80.51%	68.6%
A	86.2%	75.7%	77.2%
C	82.8%	75.25%	77%
N	82.5%	80.23%	74.8%
OE	91%	78.15%	77%

#### 5.4 Personality traits

The Big-Five personality Traits [33] is the most popular model for personality recognition, hence, it is of great importance that the proposed framework is evaluated using the big five traits labels. To the best of the authors' knowledge, this is the first study presenting the recognition of personality traits on the two chosen datasets, providing new insights on personality patterns in the depicted scenarios.

In order to use the traits scores as classification labels, following the approach used by [60], the score of each trait is transformed into a binary value (HIGH or LOW), based on the median value computed from the population. Note that every trait is evaluated independently from the scores of the other traits. As for the proposed architecture, we employ the descriptors that obtained the best results in the previous experiment. Hence, the  $D_{pers} + D_{group} + D_{prox}$  framework is used for the Salsa dataset, and the  $D_{pers} + D_{prox}$  framework is used for the Nonsocial dataset.

Results of the K-fold cross validation experiment are reported in Table 3. For the Salsa dataset, like in the previous experiment, results are presented for the two parts of the dataset independently. In the poster session part, the highest accuracy is obtained for the *OE* trait, followed by *E* and *A* traits. Considering the scenario in which the dataset was recorded (university environment), behavioral attributes related to these traits, such as curiosity or being talkative, can be easily matched with researchers' behaviors depicted in the video data. On the other hand, in the cocktail party part, the highest accuracy is obtained for the *E*, as well as *N* traits. As in this scenario the participants are allowed to freely interact with each other, attributes like talkative and sociable belonging to




Fig. 6. Example of correctly predicted sequence, belonging to high  $A$  trait. The graph contains the descriptors values (depicted in red, green and purple) on the y-axis, for each frame in the sequence on the x-axis. An alternation of high values is noticeable between the individual motion descriptor  $D_{pers}$  and the social group motion  $D_{group}$ , indicating that the two parts are taking turns in the conversation. Correspondingly, the image at frame 3 shows high motion of the skeleton colored in red, and the image at frame 9 shows high motion of the skeleton colored in green.

the  $E$  trait, are more evident than in the other scenario. For the nonsocial dataset, which contains problem solving tasks data for ADL applications, the highest accuracy is obtained for the  $A$ ,  $C$  and  $OE$  traits. When it comes to engagement with the scene, searching for objects, deliberate scanning strategies or curiosity are attributes related to  $C$  and  $OE$  traits, which can be matched to behaviors observed in the nonsocial dataset. Moreover, given the nonsocial scenario, as expected, the  $E$  trait is the hardest to predict.

## 6 QUALITATIVE RESULTS

As explained in section 5, all the proposed descriptors contribute to the understanding of certain behavioral patterns, bearing to an improved personality recognition performance. In this section, aiming to further investigate the learning process, we disentangle the behavior of each descriptor at test time. Specifically, we select two frame sequences  $t_1, t_2$  belonging to a subject with high Agreeableness trait. One sequence, ( $t_1$ ) depicted in Figure 6 was classified correctly and the network showed high confidence that the sequence belonged to the right label, whereas the other sequence ( $t_2$ ) shown in Figure 7 was misclassified. The graphs display the three descriptor values before being fed to the CNN learning framework. On the y-axis, we show the individual motion values ( $D_{pers}$  red line), social group motion ( $D_{group}$  green line), and proxemics ( $D_{prox}$  purple line) for each frame of the sequence, on the x-axis. In the correctly classified sample (Fig. 6), we can notice an alternation of




Fig. 7. Example of misclassified sequence belonging to high  $A$  trait. The graph contains the descriptors values (depicted in red, green and purple) on the y-axis, for each frame in the sequence on the x-axis. As the interaction between the two individuals is limited, and the descriptors values are quite low, the system did not associate this behavioral pattern with high  $A$  trait.

high values between the individual motion descriptor  $D_{pers}$ , and the social group motion descriptor  $D_{group}$ . This variation of high motion may denote that the two parts are conversing, each taking turns in a discussion. The displayed behavioral pattern, identified by our framework as belonging to the Agreeableness trait, is in line with previous literature studies, which correlate this trait with aspects like cooperation and empathy [56].

In the misclassified sample (Fig. 7), the interaction is more limited, and the descriptors values display low variation, while, the conversation occurs between only two individuals making the  $D_{prox}$  values low as well. This pattern violates previous literature findings, which correlate high Agreeableness trait with large and highly interactive conversational groups [3]. Hence, to further improve the results, our framework should be able to better model situations in which the motion is rather limited, or the behavior of the individuals is close to being “static”, by including additional sources of information, such as facial or audio data.

## 7 DISCOVERED PERSONALITY PATTERNS

In this section, class activation maps are explored [64] to exploit which discriminative personality patterns our proposed CNN framework learned in the social scenario data. Specifically, we use the  $D_{pers} + D_{group}$  descriptor on the Salsa dataset, to reveal the interaction between individual and group behaviors with different personality traits. The activation of the feature maps extracted by the VGG component, corresponds to spatial information of the original image [24]. We aim to discover the importance (determined by the classifier weights in the softmax layer) of both






Fig. 8. Discovered personality patterns using Person-Context descriptors, given the **High** and **Low** scores of each personality trait: **E**xtroversion, **A**greeableness, **C**onscientiousness, **N**euroticism and **O**penness to **E**xperience. Low CNN weights activation in blue, medium CNN weights activation in green, and high CNN weights activation in red.

individual dynamics as well as the social group dynamics for each personality trait, by applying the class activation maps on the Person-Context descriptors.

$$Act_p(x, y) = \sum_{k=1}^K w_k^c f_k(x, y), k \in 1, \dots, K, \quad (2)$$

$$K = D_{pers} + D_{group}$$

, where  $w$  is the weights vector of the softmax layer,  $f_k(x, y)$  is the activation of the  $k$  unit in the last convolutional layer at spatial location  $(x, y)$  for the personality class  $c$ , while the number of units  $K$  is formed of two components, one corresponding to the Person information ( $D_{pers}$ ) and one to the Context information ( $D_{group}$ ).

We are interested in visualizing meaningful personality patterns for each personality trait. Specifically, we pose the following questions: 1) *Is the relation between Person-Context dynamics reflected in the traits?* and 2) *Do the dynamics correspond to the trait attributes defined by psychologists?* The discovered personality patterns for each trait are displayed in Fig. 8. In particular, we visualize the learned patterns from both Person and Context CNN networks. We quantize the activation of each joint in three groups: low activation (first quartile) depicted in blue color, medium activation (second quartile) depicted in green color, and finally, high activation (third quartile) is depicted in red color. When an individual has a high *E* trait (e.g. talkative, outgoing), the CNN weights show high activation on the features coming from the Person stream, and low activation in the features coming from the Context stream. As a consequence, if an individual has a low *E* trait, features coming from the Context stream are more important. Thus, to answer the first question, the two stream dynamics are learned by the CNN architecture and are mapped to personality behavioral patterns. In order to answer the second question, we highlight that, since individuals with a high *E* trait are described as talkative and outgoing, when it comes to social interaction, they tend to be the *center of attention*, forcing the rest of the social group to be more passive [29]. On the other hand, CNN weights for individuals with high *A* and *C* traits, show high activation on both Person-Context streams, demonstrating that these personality types like to be engaged with the social group, without trying to dominate the situation [32]. Medium/high CNN activation on both Person-Context streams is found also for individuals with low *N* and *OE* traits, showing that

having low scores in attributes, such as being tense or being imaginative, makes individuals more sociable.

## 8 CONCLUSION

In this paper, we have presented a novel CNN-based framework for personality recognition. Our model analyzes the scene at multiple levels of granularity: firstly, we encode spatio-temporal descriptors for each individual in the scene, secondly, we extract spatio-temporal descriptors from social groups, and thirdly, we encode the global proxemics of every individual in the scene. Additionally, we demonstrate that our proxemics features can be applied also in a nonsocial scenario, encoding scene interaction information. Experiments on two personality recognition datasets demonstrate the effectiveness of our approach, showing that modeling together Person-Context information significantly improves the state-of-the-art personality recognition results. Furthermore, we presented CNN class activation maps for each personality trait, providing novel insights on nonverbal behavioral patterns linked with personality attributes defined by theories from behavioral psychology. We believe these findings are of great importance for future interdisciplinary behavioral studies, aiming to combine data-driven approaches with psychological studies.

## ACKNOWLEDGMENTS

This work has been funded by the European Union' Horizon 2020 Research and Innovation Programme under Grant Agreement N° 690090 (ICT4Life project).

## REFERENCES

- [1] Henri Achten. 2013. Buildings with an Attitude. In *Stouffs, R. and Sariyildiz, S.(eds.), Computation and Performance—Proceedings of the 31st eCAADe Conference*, Vol. 1. 477–485.
- [2] Xavier Alameda-Pineda, Jacopo Staiano, Ramanathan Subramanian, Ligia Batrinca, Elisa Ricci, Bruno Lepri, Oswald Lanz, and Nicu Sebe. 2016. Salsa: A novel dataset for multimodal group behavior analysis. *IEEE transactions on pattern analysis and machine intelligence* 38, 8 (2016), 1707–1720.
- [3] Xavier Alameda-Pineda, Yan Yan, Elisa Ricci, Oswald Lanz, and Nicu Sebe. 2015. Analyzing free-standing conversational groups: A multimodal approach. In *Proceedings of the 23rd ACM international conference on Multimedia*. ACM, 5–14.
- [4] Sharifa Alghowinem, Roland Goecke, Michael Wagner, Julien Epps, Matthew Hyett, Gordon Parker, and Michael Breakspear. 2016. Multimodal depression detection: fusion analysis of paralinguistic, head pose and eye gaze behaviors. *IEEE Transactions on Affective Computing* (2016).
- [5] Timur M Bagautdinov, Alexandre Alahi, François Fleuret, Pascal Fua, and Silvio Savarese. 2017. Social Scene Understanding: End-to-End Multi-Person Action Localization and Collective Activity Recognition.. In *CVPR*. 3425–3434.
- [6] Jeffrey D Banfield and Adrian E Raftery. 1993. Model-based Gaussian and non-Gaussian clustering. *Biometrics* (1993), 803–821.
- [7] Cigdem Beyan, Muhammad Shahid, and Vittorio Murino. 2018. Investigation of Small Group Social Interactions Using Deep Visual Activity-Based Nonverbal Features. In *2018 ACM Multimedia Conference on Multimedia Conference*. ACM, 311–319.
- [8] Sovan Biswas and Juergen Gall. 2018. Structural Recurrent Neural Network (SRNN) for Group Activity Analysis. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 1625–1632.
- [9] Jack Block and Jeanne H Block. 2014. The role of ego-control and ego-resiliency in the organization of behavior. In *Development of cognition, affect, and social relations*. Psychology Press, 49–112.
- [10] Kevin W Bowyer. 2004. Face recognition technology: security versus privacy. *IEEE Technology and society magazine* 23, 1 (2004), 9–19.
- [11] G. Bradski. 2000. The OpenCV Library. *Dr. Dobb's Journal of Software Tools* (2000).
- [12] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2016. Realtime multi-person 2d pose estimation using part affinity fields. *arXiv preprint arXiv:1611.08050* (2016).
- [13] Oya Celiktutan, Efstratios Skordos, and Hatice Gunes. 2017. Multimodal human-human-robot interactions (mhhri) dataset for studying personality and engagement. *IEEE Transactions on Affective Computing* (2017).
- [14] Florence Corpet. 1988. Multiple sequence alignment with hierarchical clustering. *Nucleic acids research* 16, 22 (1988), 10881–10890.

- [15] Marco Cristani, Loris Bazzani, Giulia Paggetti, Andrea Fossati, Diego Tosato, Alessio Del Bue, Gloria Menegaz, and Vittorio Murino. 2011. Social interaction discovery by statistical analysis of F-formations.. In *BMVC*, Vol. 2. 4.
- [16] Marco Cristani, Vittorio Murino, and Alessandro Vinciarelli. 2010. Socially intelligent surveillance and monitoring: Analysing social dimensions of physical space. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2010 IEEE Computer Society Conference on*. IEEE, 51–58.
- [17] Marco Cristani, Giulia Paggetti, Alessandro Vinciarelli, Loris Bazzani, Gloria Menegaz, and Vittorio Murino. 2011. Towards computational proxemics: Inferring social relations from interpersonal distances. In *Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third International Conference on Social Computing (SocialCom), 2011 IEEE Third International Conference on*. IEEE, 290–297.
- [18] Dario Dotti, Mirela Popa, and Stylianos Asteriadis. 2018. Behavior and Personality Analysis in a nonsocial context Dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 2354–2362.
- [19] Nicholas Epley, Adam Waytz, and John T Cacioppo. 2007. On seeing human: a three-factor theory of anthropomorphism. *Psychological review* 114, 4 (2007), 864.
- [20] Yağmur Güçlütürk, Umut Güçlü, Xavier Baro, Hugo Jair Escalante, Isabelle Guyon, Sergio Escalera, Marcel AJ Van Gerven, and Rob Van Lier. 2017. Multimodal first impression analysis with deep residual networks. *IEEE Transactions on Affective Computing* (2017).
- [21] Wilfred J Hansen. 1971. User engineering principles for interactive systems. In *Proceedings of the November 16-18, 1971, fall joint computer conference*. ACM, 523–532.
- [22] Ann Hutchinson. 1954. Labanotation. *Journal of Aesthetics and Art Criticism* 13, 2 (1954), 276–277.
- [23] Oliver P John and Sanjay Srivastava. 1999. The Big Five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook of personality: Theory and research* 2, 1999 (1999), 102–138.
- [24] Qiuhong Ke, Mohammed Bennis, Senjian An, Ferdous Sohel, and Farid Boussaid. 2017. A new representation of skeleton sequences for 3d action recognition. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*. IEEE, 4570–4579.
- [25] Markus Koppensteiner. 2013. Motion cues that make an impression: Predicting perceived personality by minimal motion information. *Journal of experimental social psychology* 49, 6 (2013), 1137–1143.
- [26] Shun Lau and Youyan Nie. 2008. Interplay between personal goals and classroom goal structures in predicting student outcomes: A multilevel analysis of person-context interactions. *Journal of educational Psychology* 100, 1 (2008), 15.
- [27] Yun-Shao Lin and Chi-Chun Lee. 2018. Using Interlocutor-Modulated Attention BLSTM to Predict Personality Traits in Small Group Interaction. In *Proceedings of the 2018 on International Conference on Multimodal Interaction*. ACM, 163–169.
- [28] Jian Liu, Naveed Akhtar, and Ajmal Mian. 2017. Skepxels: Spatio-temporal image representation of human skeleton joints for action recognition. *arXiv preprint arXiv:1711.05941* (2017).
- [29] Yu-En Lu, Sam Roberts, Pietro Lio, Robin Dunbar, and Jon Crowcroft. 2009. Size matters: variation in personal network size, personality and effect on information transmission. In *Computational Science and Engineering, 2009. CSE'09. International Conference on*, Vol. 4. IEEE, 188–193.
- [30] François Mairesse and Marilyn A Walker. 2010. Towards personality-based user adaptation: psychologically informed stylistic language generation. *User Modeling and User-Adapted Interaction* 20, 3 (2010), 227–278.
- [31] Marcin Marszałek, Ivan Laptev, and Cordelia Schmid. 2009. Actions in context. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, 2929–2936.
- [32] Christopher McCarty and HD Green. 2005. Personality and personal networks. *Sunbelt XXV* (2005).
- [33] Robert R McCrae and Oliver P John. 1992. An introduction to the five-factor model and its applications. *Journal of personality* 60, 2 (1992), 175–215.
- [34] Juan Abdon Miranda-Correa, Mojtaba Khomami Abadi, Nicu Sebe, and Ioannis Patras. 2018. AMIGOS: A Dataset for Affect, Personality and Mood Research on Individuals and Groups. *IEEE Transactions on Affective Computing* (2018).
- [35] Hossein Mousavi, Sadegh Mohammadi, Alessandro Perina, Ryad Chellali, and Vittorio Mur. 2015. Analyzing tracklets for the detection of abnormal crowd behavior. In *2015 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 148–155.
- [36] Vinod Nair and Geoffrey E Hinton. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*. 807–814.
- [37] Victor Ponce-López, Baiyu Chen, Marc Oliu, Ciprian Corneanu, Albert Clapés, Isabelle Guyon, Xavier Baró, Hugo Jair Escalante, and Sergio Escalera. 2016. Chalearn lap 2016: First round challenge on first impressions-dataset and results. In *European Conference on Computer Vision*. Springer, 400–418.
- [38] Beatrice Rammstedt and Oliver P John. 2007. Measuring personality in one minute or less: A 10-item short version of the Big Five Inventory in English and German. *Journal of research in Personality* 41, 1 (2007), 203–212.
- [39] Kamrad Khoshhah Roudposhti and Jorge Dias. 2013. Probabilistic human interaction understanding: Exploring relationship between human body motion and the environmental context. *Pattern Recognition Letters* 34, 7 (2013),

820–830.

- [40] Kamrad Khoshhal Roudposhti, Urbano Nunes, and Jorge Dias. 2016. Probabilistic social behavior analysis by exploring body motion-based patterns. *IEEE transactions on pattern analysis and machine intelligence* 38, 8 (2016), 1679–1691.
- [41] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115, 3 (2015), 211–252.
- [42] Dairazalia Sanchez-Cortes, Oya Aran, Marianne Schmid Mast, and Daniel Gatica-Perez. 2011. A nonverbal behavior approach to identify emergent leaders in small groups. *IEEE Transactions on Multimedia* 14, 3 (2011), 816–832.
- [43] David P Schmitt, Jüri Allik, Robert R McCrae, and Verónica Benet-Martínez. 2007. The geographic distribution of Big Five personality traits: Patterns and profiles of human self-description across 56 nations. *Journal of cross-cultural psychology* 38, 2 (2007), 173–212.
- [44] Karen Simonyan and Andrew Zisserman. 2014. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*. 568–576.
- [45] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [46] Cees GM Snoek, Marcel Worring, and Arnold WM Smeulders. 2005. Early versus late fusion in semantic video analysis. In *Proceedings of the 13th annual ACM international conference on Multimedia*. ACM, 399–402.
- [47] Mark Snyder, Jeffry A Simpson, and Steve Gangestad. 1986. Personality and sexual relations. *Journal of Personality and Social Psychology* 51, 1 (1986), 181.
- [48] Adriana Tapus, Cristian Tapus, and Maja J Mataric. 2007. Hands-off therapist robot behavior adaptation to user personality for post-stroke rehabilitation therapy. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*. IEEE, 1547–1553.
- [49] Alessandro Vinciarelli and Gelareh Mohammadi. 2014. A survey of personality computing. *IEEE Transactions on Affective Computing* 5, 3 (2014), 273–291.
- [50] D. Wang, B. Subagdja, Y. Kang, A. H. Tan, and D. Zhang. 2014. Towards intelligent caring agents for aging-in-place: issues and challenges. In *Proceedings of the 2014 IEEE Symposium on Computational Intelligence for Human-Like Intelligence*. IEEE Computer Society, 1–8.
- [51] Minsi Wang, Bingbing Ni, and Xiaokang Yang. 2017. Recurrent modeling of interaction context for collective activity recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [52] P. Wang, W. Li, P. Ogunbona, J. Wan, and S. Escalera. 2018. RGB-D-based human motion recognition with deep learning: A survey. *Computer Vision and Image Understanding* 171 (2018), 118–139.
- [53] Weichen Wang, Gabriella M Harari, Rui Wang, Sandrine R Müller, Shayan Mirjafari, Kizito Masaba, and Andrew T Campbell. 2018. Sensing behavioral change over time: Using within-person variability features from mobile sensing to predict personality traits. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 3 (2018), 141.
- [54] Xiu-Shen Wei, Chen-Lin Zhang, Hao Zhang, and Jianxin Wu. 2018. Deep bimodal regression of apparent personality traits from short video sequences. *IEEE Transactions on Affective Computing* 9, 3 (2018), 303–315.
- [55] Daniel Weinland, Remi Ronfard, and Edmond Boyer. 2006. Free viewpoint action recognition using motion history volumes. *Computer vision and image understanding* 104, 2-3 (2006), 249–257.
- [56] Yanna J Weisberg, Colin G DeYoung, and Jacob B Hirsh. 2011. Gender differences in personality across the ten aspects of the Big Five. *Frontiers in psychology* 2 (2011), 178.
- [57] Svante Wold, Kim Esbensen, and Paul Geladi. 1987. Principal component analysis. *Chemometrics and intelligent laboratory systems* 2, 1-3 (1987), 37–52.
- [58] Bangpeng Yao and Li Fei-Fei. 2012. Recognizing human-object interactions in still images by modeling the mutual context of objects and human poses. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34, 9 (2012), 1691–1703.
- [59] Shuai Yi, Hongsheng Li, and Xiaogang Wang. 2016. Pedestrian behavior modeling from stationary crowds with applications to intelligent surveillance. *IEEE transactions on image processing* 25, 9 (2016), 4354–4368.
- [60] Gloria Zen, Bruno Lepri, Elisa Ricci, and Oswald Lanz. 2010. Space speaks: towards socially and personality aware visual surveillance. In *Proceedings of the 1st ACM international workshop on Multimodal pervasive video analysis*. ACM, 37–42.
- [61] Dingwen Zhang, Guangyu Guo, Dong Huang, and Junwei Han. 2018. PoseFlow: A Deep Motion Representation for Understanding Human Behaviors in Videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6762–6770.
- [62] Le Zhang, Songyou Peng, and Stefan Winkler. 2018. PersEmoN: A Deep Network for Joint Analysis of Apparent Personality, Emotion and Their Relationship. In *arXiv preprint arXiv:1811.08657*.

- [63] Lei Zhao, Qinghua Hu, and Yucan Zhou. 2015. Heterogeneous features integration via semi-supervised multi-modal deep networks. In *International Conference on Neural Information Processing*. Springer, 11–19.
- [64] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2921–2929.