Emergent Leadership Detection Across Datasets

Philipp Müller
Max Planck Institute for Informatics
Saarland Informatics Campus
pmueller@mpi-inf.mpg.de

Andreas Bulling
University of Stuttgart
Institute for Visualisation and Interactive Systems
andreas.bulling@vis.uni-stuttgart.de

ABSTRACT

Automatic detection of emergent leaders in small groups from nonverbal behaviour is a growing research topic in social signal processing but existing methods were evaluated on single datasets – an unrealistic assumption for real-world applications in which systems are required to also work in settings unseen at training time. It therefore remains unclear whether current methods for emergent leadership detection generalise to similar but new settings and to which extent. To overcome this limitation, we are the first to study a cross-dataset evaluation setting for the emergent leadership detection task. We provide evaluations for within- and cross-dataset prediction using two current datasets (PAVIS and MPI-IGroupInteraction), as well as an investigation on the robustness of commonly used feature channels (visual focus of attention, body pose, facial action units, speaking activity) and online prediction in the cross-dataset setting. Our evaluations show that using pose and eye contact based features, cross-dataset prediction is possible with an accuracy of 0.68, as such providing another important piece of the puzzle towards emergent leadership detection in the real world.

KEYWORDS

social signal processing, emergent leadership detection, nonverbal behaviour

1 INTRODUCTION

Emergent leaders are group members who naturally obtain a leadership position through interaction with the group, and not via a higher authority [26]. Even without formal authority, emergent leaders are important for group performance [11, 17], and as a result automatic identification of emergent leaders in group interactions is potentially beneficial in organisational research, for hiring decisions in the context of assessment centres [15], or for robots and intelligent agents that are supposed to interact with a group naturally. Consequently, the detection of emergent leaders is a growing topic in social signal processing [6, 12, 25]. These studies used nonverbal behaviour to detect emergent leaders in group interactions, which is supported by a large body of work on the connection between emergent leadership and nonverbal behaviour [1, 14, 16].

While existent methods on emergent leadership detection in small groups showed reasonable performance, they all make the assumption that training and testing data come from the same distribution. This assumption is unrealistic for application scenarios in which a system is required to detect emergent leaders in slightly different social situations for which no labelled data is available. Until now, it remains unclear whether such cross-dataset leadership detection is possible with sufficient accuracy.

Specifically, emergent leadership detection in small groups of unaugmented people has only been investigated separately on two

Figure 1: Illustration of the recording setup of the MPI-IGroupInteraction dataset [19]. The selected view and corresponding visible participants are shown in orange.

datasets employing very similar tasks, thereby ignoring the crucial cross-dataset setting. The ELEA dataset [25] consists of meetings of three or four people each, in which participants are given the winter survival task and instructed to come up with a joint solution. Work on ELEA investigated emergent leadership detection from recordings of the meetings, by using audio- and visual or multi-modal features [24, 25], and more recently by using features obtained from a co-occurrence mining procedure [21]. Kindiroglu et al. investigated domain adaptation and multi-task learning in order to predict leadership and extraversion on ELEA using video blogs annotated with personality impressions [18]. Their work is different to the cross-dataset setting described above, as they assumed access to emergent leadership ground truth on ELEA.

The PAVIS dataset [6] consists of meetings of four people each either performing a "winter survival task" or a "desert survival task". Research on the dataset focussed on detecting emergent leaders from nonverbal features only [6], using multiple kernel learning [4], or using body pose based features [7]. Further studies improved emergent leadership detection on the PAVIS dataset by using deep visual activity features [9], or by employing sequential analysis [8]. Apart from emergent leadership detection, the dataset has also been used to predict the leadership style of emergent leaders [5, 9].

Recently, Müller et al. recorded the MPIIGroupInteraction dataset of small group interactions to study low rapport detection [19]. Although emergent leadership ratings were recorded, no approach to leadership detection was proposed. This dataset is particularly interesting for the emergent leadership detection task, as participants engaged in an open-ended discussion, which is in contrast to the rather constrained tasks that were performed on ELEA and PAVIS.

In this paper, we move one step closer to an emergent leadership detection system that can be applied in novel social situations without additional labelling effort. We investigate emergent leadership detection across situations using two recent datasets [6, 19] both featuring small group interactions but differing in participants' tasks, language, and nationality. Our specific contributions are twofold: We are the first to study emergent leadership detection

in a cross-dataset setting, thereby achieving state-of-the-art results on MPIIGroupInteraction [19]. Furthermore, we conduct extensive evaluations providing insights into the usefulness of different features and the feasibility of an online prediction system.

2 DATASETS

To study cross-dataset emergent leadership detection, we utilise the PAVIS [6] and the MPIIGroupInteraction [19] datasets of small group interactions.

2.1 PAVIS

The PAVIS dataset [6] consists of 16 interactions of four Italian speaking unacquainted participants of same gender each. Each group performed either a "winter-" or a "desert survival" task, in which participants had to agree on a ranking of items they assume to be useful in a survival situation. The interactions were recorded by four cameras, one facing each participant, as well as lapel microphones attached to each participant. Interactions lasted from 12 to 30 minutes, resulting in a total corpus length of 393 minutes. All recordings were divided into segments of four to six minutes and subsequently annotated for emergent leadership. In line with previous work [9], we exclude four recordings due to audio problems, resulting in 12 meetings with a total of 48 participants. We use PAVIS as a source dataset, as the segment-based annotation leads to more training data than is available on MPIIGroupInteraction [19].

2.2 MPIIGroupInteraction

MPIIGroupInteraction consists of 22 group interactions in German, each consisting of three- to four unacquainted participants. In contrast to the rather constrained winter- or desert survival task on the PAVIS dataset [6], participants had an open-ended discussion. The meetings were recorded by eight frame-synchronised cameras, two of them placed behind every participants in order to cover all other participants in their field of view (see Figure 1). To record audio, one microphone was placed in front and slightly above participants' heads. Each group was discussing for roughly 20 minutes, resulting in more than 440 minutes of audio-visual recordings in total. After the interaction, each participant rated every other participant on a leadership scale ("PLead" as in [25]). We use the aggregate ratings for each participant to identify the ground truth emergent leader.

3 METHOD

After extracting nonverbal features from gaze, body pose, face and speaking activity we train Support Vector Machines to detect emergent leaders.

3.1 Nonverbal Feature Extraction

3.1.1 VFOA Features. As the first step in computing features based on the visual focus of attention (VFOA), we perform eye contact detection, i.e. detecting at which other persons' face a target person is looking at a given moment in time. To this end, we employ the recently introduced method for unsupervised eye contact detection in small group interactions by Müller et al. [20]. To perform eye contact detection without the need of manual annotation, this method exploits the observation that people usually look at the person who is currently speaking in a weak labelling step. The result

are frame-wise predictions indicating with which other person the target person has eye contact, or whether the target person has no eye contact at all. To arrive at optimal results, we use ground-truth annotated speaker segmentations as input to the method on MPIIGroupInteraction. On PAVIS we resort to speaking activity detection via thesholding facial action units (cf. [20]), as we found the speaker segmentations provided with the dataset to not be perfectly synchronised with the video. Using the eye contact annotations provided by the authors of [20] for evaluation, we obtain an accuracy of 0.7 on MPIIGroupInteraction. To eliminate jitter, we apply a median filter of five frames to the eye contact predictions.

Based on these eye contact detections, we extract VFOA features as described in [6]. As the original implementation is not available from the authors, we implement the following features ourselves using the description in [6]: **totWatcher:** total time a person is watched by others, **totME:** total time a person has mutual eye contact (MEC) with others, **totWatcherNoME:** total time a person is being watched by others without having MEC, **totNoLook:** total time a person is not looking at any other person, **lookSomeOne:** total time a person looks at other people, **totInitiatorME:** proportion of MECs of a person that are initiated by her, **stdInitiatorME:** the standard deviation of lengths of MECs that are initiated by the person, **totInterCurrME:** average time between intiation of a MEC and the start of the MEC, **stdInterCurrME:** standard deviation of totInterCurrME, **totWatchNoME:** total time a person is looking at others without MEC, **maxTwoWatcherWME:** maximum time a person is looked at by two others, **minTwoWatcherWME:** minimum time a person is looked at by two others, **maxTwoWatcherNoME:** maximum time a person is looked at by two others without having MEC with them, **minTwoWatcherNoME:** minimum time a person is looked at by two others without having MEC with them, **ratioWatcherLookSOne:** ratio between totWatcher and lookSomeOne. Note that while the features we compute on top on eye contact detections are the same as in [6], in the work by Beyan et al. they are based on VFOA detections using head pose.

3.1.2 Body Pose Features. We estimate body poses of participants using OpenPose [10] and follow the approach taken in [7] for pose feature computation. In detail, we first detect frames of significant activity by a two-step thresholding approach on the difference images of subsequent greyscale frames: In the first step a pixel is classified as moving if its value exceeds the threshold $T_1 = 30$ in the difference image. The second step is classifying a frame as having significant activity if the number of moving pixels in it exceeds a threshold T_2. We set T_2 such that we obtain the same proportion of frames with significant activity as described in [7] (roughly 8.1%). For MPIIGroupInteraction we set T_2 for each interaction separately to not leak information between interactions at test time.

Subsequently, we compute the 80-dimensional featureset described in [7] on frames with significant activity. These features consist of statistical measures extracted from the angles between vectors that are defined by 2D joint positions. We use code provided to us by the authors of [7].

3.1.3 Facial Features. We use OpenFace [2, 3] to extract presence and intensity of facial action units (AUs) following the approach described in [19] for low rapport detection. We specifically extract
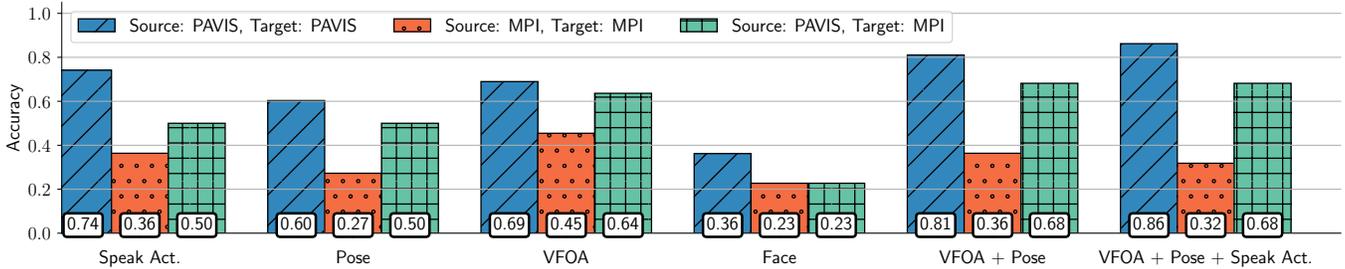
Figure 2: Performance of different featuresets when either training and testing on the same dataset, or training on PAVIS and testing on MPIIGroupInteraction. Random baseline for PAVIS as target is 0.25, for MPIIGroupInteraction as target 0.29.
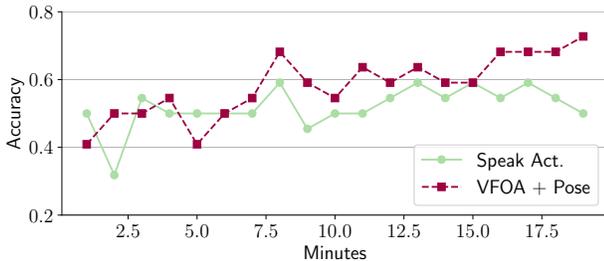
Figure 3: Performance of different featuresets when training on PAVIS and testing on MPIIGroupInteraction, depending on the size of the time window that is used for analysis (starting from the beginning). Random baseline is at 0.29.

the means of AU activations and intensities and the mean and standard deviation of a "facial positivity indicator".

3.1.4 Speaking Activity Features. To evaluate the importance of speaking activity, we implement features used in previous work [24], specifically the total speaking time of a participant (SPL), the number of speaking turns of a participant (SPT), the total number of times a participant interrupts other participants (SPI), and the average duration of a participants' speaking turns (ASP). We normalise SPL, SPT and SPI with the length of the time interval from which we extract the feature. On both datasets, we extract speaking activity features from ground truth speaker segmentations.

3.2 Classification

For classification, we use Support Vector Machines (SVMs) with radial basis function kernels. To obtain a single predicted leader for each interaction during test time, we obtain probability estimates using Platt scaling [23] and select the participant with the highest probability as the predicted emergent leader. We choose the regularisation parameter C of the SVM via cross-validation on the source dataset (PAVIS) and use the default value $1/n_{feats}$ for γ. Dedicated domain adaptation methods including Transfer Component Analysis [22], Correlation Alignment [27], Random Walk Adaptation [28] as well as transductive methods like label propagation [29] could not consistently improve over the plain SVM approach in our experiments.

The standard way to normalise both train and test data is via mean and standard deviation computed on the training data [13]. This prevents information leakage from the test set at training

time (e.g. when normalising train and test data jointly), and also leakage from "future" test samples at test time (when normalising the whole test set at once). However, in our case training and testing data distributions differ and our data is structured by interactions made up of three to four individual participants. As a consequence, while normalising the training data as usual, we normalise each test interaction separately (i.e. independently from the training data as well as other test interactions). In this way, no information "from the future" is leaked while testing and we comply to the fact of different training and testing distributions. In preliminary experiments, we found this way of normalising to be crucial. The common way of normalising by applying mean and standard deviation of the training data on the test data resulted in much worse performance.

When employing several featuresets for classification, we always use late fusion, i.e. averaging scores of classifiers applied independently on the respective featuresets. In preliminary experiments this produced more reliable results than early fusion.

4 EXPERIMENTAL RESULTS

All our evaluations are based on per-interaction accuracy of emergent leadership predictions as in [24, 25]. Specifically, an interaction is counted as correct, if and only if the predicted emergent leader is the same as the ground truth emergent leader.

4.1 Offline Prediction

To evaluate the extent to which classifiers trained on a source dataset are able to achieve high performance on a target dataset, we train on PAVIS and test on MPIIGroupInteraction. At test time we assume to have access to a full test recording, i.e. we are predicting emergent leadership after an interaction took place ("offline" setting). In order to ensure using the same length for each of the approximately 20 minute long interactions on MPIIGroupInteraction we always use the first 19 minutes for feature extraction.

Figure 2 shows the obtained results for different feature sets and source- and target dataset combinations. The highest performance in the cross-dataset setting ("Source: PAVIS, Target: MPI") is achieved by a combination of VFOA and pose features with an accuracy of 0.68, slightly outperforming VFOA features only at 0.64 accuracy. Combining other features with VFOA and pose did not improve results. For applications where video recordings are not available or not desired, an accuracy of 0.5 can be achieved by relying on speaking activity features only. Both results are clearly above

the random baseline of 0.29, showing the feasibility of cross-dataset emergent leadership prediction.

Comparing cross-dataset to within-dataset results reveals that cross-dataset accuracies are consistently lower than within-dataset accuracies on PAVIS. More surprisingly, by training on PAVIS, we achieve higher accuracies on MPIIGroupInteraction compared to training on MPIIGroupInteraction directly. This is most likely an effect of the limited training data available on MPIIGroupInteraction. In total there are only 78 samples (one per participant), whereas on PAVIS we have 232 samples due to the segment based annotations.

Concerning within-dataset results, we achieve the best accuracy for PAVIS with a combination of speaking activity, VFOA and pose features (0.86). The best result for the emergent leadership detection task on PAVIS was published in [7], achieving detection scores of 0.76 for the positive class and 0.93 for the negative class with a combination of pose and VFOA features. Later work by the same authors adopted a different evaluation setting, and thus can not serve as a comparison [8, 9]. The detection scores for our predictions on PAVIS based on VFOA, pose and speaking activity features, reach 0.86 for the positive class and 0.95 for the negative class, exceeding the previously published results. Within-dataset results on MPIIGroupInteraction are much lower, which is most probably due to the limited number of training examples. Here, the best performance is reached by VFOA features at 0.45 accuracy.

4.2 Online Prediction

Some applications scenarios require information about emergent leaders already during the course of an interaction. To evaluate performance in this setting, we restrict the time interval from which to extract features from the target interactions at test time. Figure 3 shows accuracies for classifiers that only observe data from a limited number of minutes at the beginning of the interaction. Both our best performing featureset (VFOA and pose) and speaking activity features tend to achieve higher accuracies after longer observation time. This tendency is more pronounced for the VFOA and pose featureset, which stays between 0.4 and 0.6 accuracy during the first minutes of an interaction, and clearly above 0.6 accuracy after observing more than 15 minutes of the interaction. Thus, while prediction above chance is possible early on, for optimal precision a significant portion of the interaction has to be observed.

4.3 Feature Analysis

VFOA features were the best performing individual featureset in our evaluation. To better understand which VFOA features generalise best across datasets, we investigate how well each individual feature discriminates the ground truth classes on MPIIGroupInteraction and PAVIS. That is, for each interaction, we construct an unlearned classifier from a single feature by selecting the person with either the maximum or the minimum value on that feature as the emergent leader. The choice of selection via maximum or minimum is based on achieved accuracy when comparing to ground truth. We refer to features of which we take the maximum/minimum as having positive/negative orientation respectively. It is important to note that this is not a valid classification approach, as we do not employ cross-validation. Instead, it is a post-hoc analysis on the connection between individual features and ground truth. The

Feature	MPI		PAVIS	
	Acc.	Ori.	Acc.	Ori.
totWatcherNoME	0.59	+	0.66	+
ratioWatcherLookSOne	0.59	+	0.62	+
totWatcher	0.55	+	0.76	+
totWatchNoME	0.55	−	0.43	−
totInitiatorME	0.45	−	0.40	−
lookSomeOne	0.45	−	0.34	−
stdInitiatorME	0.45	+	0.34	+
totNoLook	0.45	+	0.34	+
stdInterCurrME	0.45	−	0.41	−
maxTwoWatcherNoME	0.45	+	0.21	+
minTwoWatcherWME	0.45	−	0.14	+
maxTwoWatcherWME	0.41	+	0.36	+
minTwoWatcherNoME	0.41	−	0.14	−
totInterCurrME	0.41	−	0.43	−
totME	0.36	+	0.60	+

Table 1: Accuracies for single feature based classification using VFOA features on PAVIS and MPIIGroupInteraction. "Ori." indicates whether the maximum or the minimum of the feature was used for prediction.

results are summarised in Table 1. The usefulness of VFOA features for cross-dataset prediction is illustrated by the fact that all features except one share the same orientation. The features with the highest accuracies on both datasets are *totWatcherNoME*, *ratioWatcherLookSOne* and *totWatcher*. This indicates that being looked at by others is a central property of leaders that is robust across datasets. In contrast, the low performance of *totME* on MPIIGroupInteraction in comparison to PAVIS indicates that mutual eye contact is less robustly associated with leadership across the two datasets.

5 CONCLUSION

In this paper, we are first to investigate a cross-dataset evaluation setting for the emergent leadership detection task. We showed that it is possible to predict emergent leadership from nonverbal features on a new dataset that was not observed at test time. We found that a combination of VFOA and pose features achieved best performance in the cross-dataset evaluation. Furthermore, we analysed the feasibility of online prediction and the usefulness of single VFOA features. All in all, our initial study on cross-dataset emergent leadership prediction opens the way to studying this important task in more realistic settings.

REFERENCES

[1] John E Baird Jr. 1977. Some nonverbal elements of leadership emergence. *Southern Speech Communication Journal* 42, 4 (1977), 352–361. https://doi.org/10.1080/10417947709372361

[2] Tadas Baltrušaitis, Marwa Mahmoud, and Peter Robinson. 2015. Cross-dataset learning and person-specific normalisation for automatic action unit detection. In *Proc. of the IEEE International Conference and Workshops on Automatic Face and Gesture Recognition*, Vol. 6. 1–6. https://doi.org/10.1109/FG.2015.7284869

[3] Tadas Baltrusaitis, Amir Zadeh, Yao Chong Lim, and Louis-Philippe Morency. 2018. Openface 2.0: Facial behavior analysis toolkit. In *Proc. of the IEEE International Conference on Automatic Face & Gesture Recognition*. 59–66. https://doi.org/10.1109/FG.2018.00019

[4] Cigdem Beyan, Francesca Capozzi, Cristina Becchio, and Vittorio Murino. 2016. Identification of emergent leaders in a meeting scenario using multiple kernel learning. In *Proc. of the Workshop on Advancements in Social Signal Processing for Multimodal Interaction*. 3–10. https://doi.org/10.1145/3005467.3005469

[5] Cigdem Beyan, Francesca Capozzi, Cristina Becchio, and Vittorio Murino. 2018. Prediction of the Leadership Style of an Emergent Leader Using Audio and Visual Nonverbal Features. *IEEE Transactions on Multimedia* 20, 2 (2018), 441–456. https://doi.org/10.1109/TMM.2017.2740062

[6] Cigdem Beyan, Nicolò Carissimi, Francesca Capozzi, Sebastiano Vascon, Matteo Bustreo, Antonio Pierro, Cristina Becchio, and Vittorio Murino. 2016. Detecting emergent leader in a meeting environment using nonverbal visual features only. In *Proc. of the ACM International Conference on Multimodal Interaction*. 317–324. https://doi.org/10.1145/2993148.2993175

[7] Cigdem Beyan, Vasiliki-Maria Katsageorgiou, and Vittorio Murino. 2017. Moving as a Leader: Detecting Emergent Leadership in Small Groups using Body Pose. In *Proc. of the ACM Multimedia Conference*. 1425–1433. https://doi.org/10.1145/3123266.3123404

[8] Cigdem Beyan, Vasiliki-Maria Katsageorgiou, and Vittorio Murino. 2019. A Sequential Data Analysis Approach to Detect Emergent Leaders in Small Groups. *IEEE Transactions on Multimedia* (2019). https://doi.org/10.1109/TMM.2019.2895505

[9] Cigdem Beyan, Muhammad Shahid, and Vittorio Murino. 2018. Investigation of Small Group Social Interactions Using Deep Visual Activity-Based Nonverbal Features. In *Proc. of the ACM Multimedia Conference*. 311–319. https://doi.org/10.1145/3240508.3240685

[10] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2018. OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields. In *arXiv preprint arXiv:1812.08008*.

[11] Vanessa Urch Druskat and Anthony T Pescosolido. 2006. The impact of emergent leader's emotionally competent behavior on team trust, communication, engagement, and effectiveness. *Research on Emotions in Organizations* 2 (2006), 25–55.

[12] Sebastian Feese, Amir Muaremi, Bert Arnrich, Gerhard Troster, Bertolt Meyer, and Klaus Jonas. 2011. Discriminating Individually Considerate and Authoritarian Leaders by Speech Activity Cues. In *Proc. of the IEEE International Conference on Privacy, Security, Risk and Trust and IEEE International Conference on Social Computing*. 1460–1465. https://doi.org/10.1109/PASSAT/SocialCom.2011.209

[13] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. 2001. *The elements of statistical learning*. Vol. 1. Springer Series in Statistics, New York.

[14] Fabiola H Gerpott, Nale Lehmann-Willenbrock, Jeroen D Silvis, and Mark Van Vugt. 2018. In the eye of the beholder? An eye-tracking experiment on emergent leadership in team interactions. *The Leadership Quarterly* 29, 4 (2018), 523–532. https://doi.org/10.1016/j.leaqua.2017.11.003

[15] Leonard D Goodstein and Richard I Lanyon. 1999. Applications of Personality Assessment to the Workplace: A Review. *Journal of Business and Psychology* 13, 3 (1999), 291–322. https://doi.org/10.1023/A:1022941331649

[16] Akko Kalma. 1992. Gazing in triads: A powerful signal in floor apportionment. *British Journal of Social Psychology* 31, 1 (1992), 21–39. https://doi.org/10.1111/j.2044-8309.1992.tb00953.x

[17] Jill Kickul and George Neuman. 2000. Emergent Leadership Behaviors: The Function of Personality and Cognitive Ability in Determining Teamwork Performance and KSAs. *Journal of Business and Psychology* 15, 1 (2000), 27–51. https://doi.org/10.1023/A:1007714801558

[18] Ahmet Alp Kindiroglu, Lale Akarun, and Oya Aran. 2017. Multi-domain and multi-task prediction of extraversion and leadership from meeting videos. *EURASIP Journal on Image and Video Processing* 2017, 1 (2017), 77. https://doi.org/10.1186/s13640-017-0224-z

[19] Philipp Müller, Michael Xuelin Huang, and Andreas Bulling. 2018. Detecting Low Rapport During Natural Interactions in Small Groups from Non-Verbal Behavior. In *Proc. of the ACM International Conference on Intelligent User Interfaces*. https://doi.org/10.1145/3172944.3172969

[20] Philipp Müller, Michael Xuelin Huang, Xucong Zhang, and Andreas Bulling. 2018. Robust Eye Contact Detection in Natural Multi-Person Interactions Using Gaze and Speaking Behaviour. In *Proc. of the International Symposium on Eye Tracking Research and Applications*. 31:1–31:10. https://doi.org/10.1145/3204493.3204549

[21] Shogo Okada, Laurent Son Nguyen, Oya Aran, and Daniel Gatica-Perez. 2019. Modeling Dyadic and Group Impressions with Intermodal and Interperson Features. *ACM Transactions on Multimedia Computing, Communications, and Applications* 15, 1s (2019), 13. https://doi.org/10.1145/3265754

[22] Sinno Jialin Pan, Ivor W Tsang, James T Kwok, and Qiang Yang. 2011. Domain Adaptation via Transfer Component Analysis. *IEEE Transactions on Neural Networks* 22, 2 (2011), 199–210. https://doi.org/10.1109/TNN.2010.2091281

[23] John Platt et al. 1999. Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods. *Advances in Large Margin Classifiers* 10, 3 (1999), 61–74.

[24] Dairazalia Sanchez-Cortes, Oya Aran, Dinesh Babu Jayagopi, Marianne Schmid Mast, and Daniel Gatica-Perez. 2013. Emergent leaders through looking and speaking: from audio-visual data to multimodal recognition. *Journal on Multimodal User Interfaces* 7, 1-2 (2013), 39–53. https://doi.org/10.1007/s12193-012-0101-0

[25] Dairazalia Sanchez-Cortes, Oya Aran, Marianne Schmid Mast, and Daniel Gatica-Perez. 2012. A Nonverbal Behavior Approach to Identify Emergent Leaders in Small Groups. *IEEE Transactions on Multimedia* 14, 3 (2012), 816–832. https://doi.org/10.1109/TMM.2011.2181941

[26] R Timothy Stein and Tamar Heller. 1979. An empirical analysis of the correlations between leadership status and participation rates reported in the literature. *Journal of Personality and Social Psychology* 37, 11 (1979), 1993–2002. https://doi.org/10.1037/0022-3514.37.11.1993

[27] Baochen Sun, Jiashi Feng, and Kate Saenko. 2016. Return of frustratingly easy domain adaptation.. In *Proc. of the AAAI Conference on Artificial Intelligence*. 2058–2065.

[28] Twan van Laarhoven and Elena Marchiori. 2017. Unsupervised Domain Adaptation with Random Walks on Target Labelings. *arXiv preprint arXiv:1706.05335* (2017).

[29] Xiaojin Zhu and Zoubin Ghahramani. 2002. *Learning from labeled and unlabeled data with label propagation*. Technical Report.