

# UNSUPERVISED AUDIO-VISUAL SUBSPACE ALIGNMENT FOR HIGH-STAKES DECEPTION DETECTION

*Leena Mathur and Maja J Matarić*

Department of Computer Science  
University of Southern California, Los Angeles, CA

## ABSTRACT

Automated systems that detect deception in *high-stakes* situations can enhance societal well-being across medical, social work, and legal domains. Existing models for detecting high-stakes deception in videos have been *supervised*, but labeled datasets to train models can rarely be collected for most real-world applications. To address this problem, we propose the first multimodal *unsupervised* transfer learning approach that detects real-world, high-stakes deception in videos without using high-stakes labels. Our subspace-alignment (SA) approach adapts audio-visual representations of deception in lab-controlled *low-stakes* scenarios to detect deception in real-world, *high-stakes* situations. Our best unsupervised SA models outperform models without SA, outperform human ability, and perform comparably to a number of existing supervised models. Our research demonstrates the potential for introducing subspace-based transfer learning to model high-stakes deception and other social behaviors in real-world contexts with a scarcity of labeled behavioral data.

**Index Terms**— transfer learning, deception detection

## 1. INTRODUCTION

Advances in human-centered signal processing and multimodal machine learning are enabling the development of automated systems that can detect human social behaviors, including *deception* [1]. Deception involves the intentional communication of false or misleading information [2] and has been categorized as occurring in either *high-stakes* or *low-stakes* social situations [3]. Deceivers in high-stakes contexts face substantial consequences if their deception is discovered, in contrast to deceivers in low-stakes contexts. Healthcare providers, social workers, and legal groups have fostered an interest in detecting *high-stakes* deception for applications that enhance societal well-being (e.g., helping therapists recognize whether clients are masking negative emotions, helping judges assess courtroom testimonies of children coerced to lie) [4, 5]. Human deception detection ability has been determined as close to chance level [6], motivating the development of computational approaches that can help humans in this challenging task.

A fundamental challenge in modeling high-stakes deception is the *scarcity of labeled high-stakes deception data*, due to the difficulty in collecting large amounts of real-world data with verifiable ground truth [3]. Lab-controlled experiments to simulate realistic high-stakes scenarios are not ethical, because they require the use of threats to impose substantial consequences on deceivers. Therefore, most researchers have collected data of participants communicating truthfully and deceptively in lab-controlled, low-stakes situations (e.g., mock crime scenarios) to study behavioral cues that could be indicative of high-stakes deception [7, 3].

Unsupervised domain adaptation (UDA) leverages knowledge from labeled *source* domains to perform tasks in related, unlabeled *target* domains [8, 9]. We addressed the data scarcity problem of high-stakes deception by proposing a novel UDA approach, based on subspace alignment (SA) [9], to detect high-stakes deception without using any high-stakes labels. Grounded in psychology research regarding the generalizability of deceptive cues across contexts [2], we hypothesized that *audio-visual representations of low-stakes deception in lab-controlled situations can be leveraged by SA to detect high-stakes deception in real-world situations*.

We experimented with unimodal and multimodal audio-visual SA models to contribute effective modalities, fusion approaches, and behavioral signals for unsupervised high-stakes deception detection. Our best unsupervised SA models (75% AUC, 74% accuracy) outperform models without SA, outperform human ability, and perform comparably to a number of existing supervised models [10, 11, 12, 13]. Our research demonstrates the potential for introducing unsupervised subspace-based transfer learning approaches to address the scarcity of labeled data when modeling high-stakes deception and other social behaviors in real-world situations.

## 2. BACKGROUND

Existing approaches for detecting high-stakes deception in videos have leveraged *supervised* machine learning models that exploit discriminative patterns in human visual, verbal, vocal, and physiological cues to distinguish deceptive and truthful communication [5]. To address the scarcity of labeled high-stakes deception data for training models [14], prior re-

search has focused on developing models that are *robust* to small numbers of samples [14, 15]. While *supervised* transfer learning models have been developed to detect deception [16, 17], to the best of our knowledge, no existing research has introduced *unsupervised* models to address the data scarcity problem of high-stakes deception detection.

Psychology studies have found that humans can rely on audio-visual behavioral cues (e.g., facial expressions, voice pitch) to detect deception across different people and contexts in lab-controlled experiments that attempt to simulate high-stakes situations (e.g., mock crime scenarios) [2, 6]. These findings motivated our decision to leverage representations of audio-visual cues during low-stakes deception to detect high-stakes deception across different people and contexts. We note that subspace-based transfer learning has been successfully used to detect emotion [18, 19], but has not been previously developed to detect social behaviors. To the best of our knowledge, our unsupervised subspace-based transfer learning approach is the first to detect a social behavior, leveraging audio-visual cues to detect high-stakes deception.

### 3. DATASETS

This section describes the video datasets used for real-world *high-stakes* deception and lab-collected *low-stakes* deception.

#### 3.1. High-stakes deception dataset

For high-stakes deception, we used a publicly-available video dataset of people in 121 real-world courtroom trial situations ( $\sim 28$  sec per video); this dataset is the current benchmark for high-stakes deception detection in videos [11]. Each video was labeled as *truthful* or *deceptive* per police testimony and trial information; we used these labels as ground truth. Per criteria used by prior research with this dataset [20, 21], we identified a usable subset of 108 videos (53 truthful videos, 55 deceptive videos; 47 speakers of diverse race and gender). All videos were collected in unconstrained situations with a variety of illuminations, camera angles, and face occlusions.

#### 3.2. Low-stakes deception dataset

For low-stakes deception, we used the UR Lying Dataset [7], the only publicly-available video dataset that had participants voluntarily communicating truthfully and deceptively in lab-controlled game scenarios ( $\sim 23$  min per video). Participants chose to respond either truthfully or deceptively to questions about images (e.g., picture of a flower). They won \$10 if they were believed while communicating truthfully and \$20 if they were believed while communicating deceptively. We used 107 videos (44 truthful videos, 63 deceptive videos; 29 speakers of diverse race and gender). All videos were collected in situations with consistent illuminations and camera angles, as well as minimal face occlusions.

## 4. METHODOLOGY

This section describes our approach (**Fig. 1**) for detecting *high-stakes* deception through unsupervised audio-visual SA.

### 4.1. Feature Extraction

We extracted the same set of 89 audio-visual behavioral cues from speakers in both datasets. The OpenSMILE [22] toolkit was used to extract 58 audio features from the eGeMAPs [23] and MFCC feature sets, to capture spectral, cepstral, prosodic, and voice quality information at each audio frame. The eGeMAPs and MFCC features sets in OpenSMILE overlap on MFCC features, and we do not include any duplicates. The OpenFace [24] toolkit was used to extract 31 visual features to capture facial action units (FAU), head pose, and eye gaze information from speakers at each visual frame.


The 89 audio-visual features were extracted frame-by-frame from videos. To prepare these features for a binary video classification task, we represented each as a fixed-length vector of time-series attributes during variable-length videos, similar to prior methods [10, 21]. For each feature, the TsFresh toolkit [25] computed 12 attributes: mean, standard deviation, aggregate autocorrelation across different time lags, and changes in feature values across quantiles. These functions are documented in the EfficientFCParameters class of TsFresh. A fixed-length feature vector of 1068 audio-visual time-series features ( $89 \times 12$ ) was computed to represent each video.

### 4.2. Unsupervised Subspace Alignment (SA)

We formulated high-stakes deception detection as an unsupervised subspace alignment (SA) transfer learning problem [9]. Given a set  $S$  of  $Z$  labeled low-stakes deception samples (source domain) and a set  $T$  of  $Y$  unlabeled high-stakes deception samples (target domain), the modeling goal is to train a classifier on  $S$  that predicts *deceptive* and *truthful* labels of  $T$ . Both  $S$  and  $T$  are in  $M$ -dimensional feature spaces and are drawn according to different marginal distributions. PCA computes the first  $D$  principal components of each domain to create  $D$ -dimensional subspace embeddings  $C_S$  (source components) and  $C_T$  (target components). The optimal linear transformation matrix  $\phi$  is computed by minimizing the difference between  $C_S$  and  $C_T$  with the following equation:

$$\begin{aligned} \phi &= \operatorname{argmin}_{\phi} \|C_S \phi - C_T\|_F^2 \\ &= \operatorname{argmin}_{\phi} \|C_S^T C_S \phi - C_S^T C_T\|_F^2 \\ &= \operatorname{argmin}_{\phi} \|\phi - C_S^T C_T\|_F^2 \\ &= C_S^T C_T \end{aligned} \tag{1}$$

where  $\|\cdot\|$  denotes the Frobenius norm. This matrix  $\phi$  is used to transform low-stakes deception embeddings  $C_S$  to align with high-stakes deception. Classifiers are then trained



**Fig. 1:** We propose *unsupervised, audio-visual subspace alignment (SA)* for detecting high-stakes deception without using any high-stakes labels.  $M$ -dimensional subspaces of low-stakes deception (blue) are aligned with high-stakes deception (orange) through transition matrices  $\phi$  and used in our  $D$ -dimensional SA approach (green) to predict high-stakes deception.

on these aligned low-stakes embeddings and tested on high-stakes embeddings  $C_T$  to predict *deceptive* or *truthful* labels for high-stakes samples. Our approach is visualized in **Fig. 1**; additional details on unsupervised SA are in [9].

### 4.3. Classification Experiments

We experimented with 9 different unimodal and multimodal unsupervised SA approaches, described below, to identify effective modalities, feature sets, and modeling approaches.

#### 4.3.1. Unimodal and Multimodal Classification Models

Unimodal audio SA classifiers were trained on the set of all audio features and separately on MFCC and eGeMAPs feature sets. Unimodal visual SA classifiers were trained on the set of all visual features and separately on FAU, eye gaze, and head pose feature sets. Multimodal SA classifiers were trained with two approaches: (1) *early-fusion* and (2) *late-fusion*. Early-fusion SA classifiers were trained on concatenated feature vectors of all audio-visual features. For late-fusion, separate unimodal SA classifiers were trained on audio and visual features, and a majority vote of their predicted class probabilities determined the final prediction. The classifier for all SA experiments was K Nearest Neighbors (KNN), implemented with scikit-learn [26]. Similar to KNN hyperparameter tuning for SA in [9], all SA experiments were conducted with 3-fold cross-validation in the source subspace to identify optimal values of the KNN nearest-neighbors hyperparameter  $k$  in the discrete range [1, 30]. An optimal subspace dimension  $D$  was determined in the discrete range [1, 10], since smaller values of  $D$  avoid computationally expensive eigendecomposition.

#### 4.3.2. Baseline Models and Chance Performance

For each unimodal and multimodal SA model, we implemented baseline KNN classifiers trained on low-stakes deception data and tested on high-stakes deception data, without alignment, in order to evaluate the effectiveness of SA. For fair comparison, each baseline KNN was implemented with the same hyper-parameter  $k$  as the optimal  $k$  that was automatically computed during 3-fold cross-validation in the corresponding SA model. To compare our model performance to human deception detection ability, which is at chance level [6], we defined the human performance baseline as a classifier that would achieve 51% accuracy (always predicting *deceptive* for 55 deceptive videos out of 108 videos).

#### 4.3.3. Metrics

Aligned with previous research [21, 20, 10], the following metrics were computed to evaluate classifiers: (1) ACC, classification accuracy across the videos; (2) AUC, the probability of the classifier ranking a randomly chosen deceptive sample higher than a randomly chosen truthful one; (3) F1-score, weighted average of precision and recall.

## 5. RESULTS AND DISCUSSION

Modeling results from classification experiments are presented in **Table 1** and visualized in **Fig. 2**. All unimodal and multimodal unsupervised SA models substantially outperformed the baseline models without SA and the human performance chance level, *demonstrating the effectiveness of unsupervised SA for modeling high-stakes deceptive behavior without any high-stakes labels*. Significance values of differences in model performance were computed with McNemar’s test ( $\alpha=0.05$ ) with continuity correction [27].

Model	ACC	AUC	F1
<b>Audio</b>			
Audio (MFCC + eGeMAPs)	0.63	0.65	0.69
MFCC**	0.61	0.60	0.62
eGeMAPs	0.55	0.55	0.48
Audio Baseline	0.52	0.48	0.60
MFCC Baseline	0.42	0.42	0.53
eGeMAPs Baseline	0.49	0.50	0.35
<b>Visual</b>			
Visual (FAU + Gaze + Pose)**	<b>0.74</b>	<b>0.75</b>	<b>0.73</b>
FAU*	0.64	0.59	0.70
Gaze*	0.70	0.71	0.64
Pose*	0.63	0.62	0.61
Visual Baseline	0.62	0.63	0.59
FAU Baseline	0.51	0.47	0.61
Gaze Baseline	0.52	0.52	0.57
Pose Baseline	0.43	0.42	0.52
<b>Audio-Visual</b>			
Audio-Visual Early Fusion (EF)**	0.64	0.60	0.69
Audio-Visual Late Fusion (LF)	0.63	0.65	0.68
Audio-Visual EF Baseline	0.44	0.44	0.53
Audio-Visual LF Baseline	0.52	0.47	0.66


\* Significant difference ( $p < 0.05$ ) between this SA model and baseline

\*\* Significant difference ( $p < 0.01$ ) between this SA model and baseline

**Table 1:** Classification Results

The unimodal *visual SA* model, trained on all visual features, had the highest performance (hyper-parameters  $D=6$  and  $k=3$ ). This model achieved an ACC of 74%, AUC of 75%, and F1-score of 73%, and significantly outperformed the visual baseline model without SA ( $p<0.01$ ). This model also outperformed the unimodal SA models trained on audio features and the multimodal early fusion and late fusion SA models trained on audio-visual features. Within the visual modality, SA models trained on the *gaze* features outperformed SA models trained on the facial action units and head pose features. These results suggest that representations of visual behavior, in particular eye gaze, demonstrate more transfer potential than audio and audio-visual representations in subspace-based transfer learning approaches for detecting deception across different people in different social contexts.

Our *unsupervised* visual SA model performed comparably to existing *fully-supervised*, automated approaches that used the same dataset (ACC 75% [11], 77% [13], 79% [12]; AUC 70% [10]). While some prior fully-supervised, automated approaches [21, 15, 20] outperformed our unsupervised SA, our findings support the potential for introducing unsupervised SA to address the *data scarcity* problem of modeling high-stakes deception. *Our results support our hypothesis that audio-visual representations of low-stakes deception in lab-controlled situations can be leveraged by SA to detect high-stakes deception in real-world situations.*



**Fig. 2:** Classification accuracies of each unimodal and multi-modal SA model (blue) compared to baseline models (orange) and the chance human performance (dashed line).

To analyze audio-visual behavioral signals that are similar in distribution across low-stakes and high-stakes deception contexts, we conducted a two-tail independent sample Welch’s t-test between each feature’s low-stakes and high-stakes distributions without assuming equal variance [28]. The following 10 behavioral signals exhibited the most similarity in distribution: jitter, left eye gaze, pitch, yaw of head pose, roll of head pose, right eye gaze, the ratio of the first two harmonics of fundamental frequency, and the MFCC coefficients 0, 7, and 8. The similar distributions of these features across different groups of people and different social contexts demonstrate their potential for use as *transferable* cues in models for cross-situational deception detection.

## 6. CONCLUSION

This paper proposes the first *unsupervised* transfer learning approach for detecting real-world, high-stakes deception in videos without using high-stakes labels. Our subspace-alignment models addressed the *data scarcity* problem of modeling high-stakes deception by adapting audio-visual representations of deception in lab-controlled *low-stakes* scenarios to detect deception in real-world, *high-stakes* situations. We contribute effective modalities, modeling approaches, and audio-visual cues, for high-stakes deception detection. Our research demonstrates the potential for introducing subspace-based transfer learning approaches to model high-stakes deception and other social behaviors in real-world contexts, when faced with a scarcity of labeled behavioral data.

## 7. ACKNOWLEDGEMENTS

We thank the Rochester HCI Group for sharing their low-stakes deception dataset. This research was supported by the USC Provost’s Undergraduate Research Fellowship.

## 8. REFERENCES

- [1] T. Baltrusaitis, C. Ahuja, and L.P. Morency, “Multimodal machine learning: A survey and taxonomy,” *IEEE Trans. PAMI*, vol. 41(2), 423–443, Feb. 2019.
- [2] M. Frank and P. Ekman, “The ability to detect deceit generalizes across different types of high-stake lies,” *J. Pers. Soc. Psychol.*, vol. 72, pp. 1429–39, 1997.
- [3] S. Porter and L. ten Brinke, “The truth about lies: What works in detecting high-stakes deception?,” *Legal and Criminological Psych.*, vol. 15, no. 1, pp. 57–75, 2010.
- [4] V. Ardulov, Z. Durante, S. Williams, T. Lyon, and S. Narayanan, “Identifying truthful language in child interviews,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [5] M. Burzo et al., *Multimodal Deception Detection*, p. 419–453, Association for Computing Machinery, 2018.
- [6] C.F. Bond and B.M. Depaulo, “Accuracy of deception judgments,” *Personality and Social Psych. Rev.*, vol. 10(3), 214–34, 2006.
- [7] T. Sen et al., “Automated dyadic data recorder (addr) framework and analysis of facial cues in deceptive communication,” *ACM Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 1, no. 4, Jan. 2018.
- [8] S. J. Pan and Q. Yang, “A survey on transfer learning,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.
- [9] B. Fernando, A. Habrard, M. Sebban, and T. Tuytelaars, “Unsupervised visual domain adaptation using subspace alignment,” in *2013 IEEE International Conference on Computer Vision (ICCV)*, 2013, pp. 2960–2967.
- [10] R. Rill-García et al., “High-level features for multimodal deception detection in videos,” *2019 IEEE CVPR Workshops*, pp. 1565–1573, 2019.
- [11] V. Pérez-Rosas et al., “Deception detection using real-life trial data,” in *ACM International Conference on Multimodal Interaction (ICMI)*, 2015, p. 59–66.
- [12] M. Jaiswal, S. Tabibu, and R. Bajpai, “The truth and nothing but the truth: Multimodal analysis for deception detection,” in *2016 IEEE 16th International Conference on Data Mining Workshops*, 2016, pp. 938–943.
- [13] D. Avola et al., “Automatic deception detection in rgb videos using facial action units,” in *ACM International Conference on Distributed Smart Cameras*, 2019.
- [14] H. Karimi, J. Tang, and Y. Li, “Toward end-to-end deception detection in videos,” in *IEEE International Conference on Big Data*, 2018, pp. 1278–1283.
- [15] M. Ding, A. Zhao, Z. Lu, T. Xiang, and J. Wen, “Face-focused cross-stream network for deception detection in videos,” in *IEEE CVPR*, 2019, pp. 7794–7803.
- [16] K. Hasan et al., “Facial expression based imagination index and a transfer learning approach to detect deception,” in *International Conference on Affective Computing and Intelligent Interaction (ACII)*, 2019.
- [17] Q. Luo, R. Gupta, and S. Narayanan, “Transfer learning between concepts for human behavior modeling: An application to sincerity and deception prediction,” in *Interspeech*, 2017, pp. 1462–1466.
- [18] Z. Yang, B. Gong, and S. Narayanan, “Weighted geodesic flow kernel for interpersonal mutual influence modeling and emotion recognition in dyadic interactions,” in *International Conference on Affective Computing and Intelligent Interaction (ACII)*, 2017.
- [19] P. Song and W. Zheng, “Feature selection based transfer subspace learning for speech emotion recognition,” *IEEE Trans. Affect. Comput.*, vol. 11(3), 373–82, 2020.
- [20] Z. Wu et al., “Deception detection in videos,” in *Proceedings of AAAI 2018*, pp. 1695–1702.
- [21] L. Mathur and M.J. Matarić, “Introducing representations of facial affect in automated multimodal deception detection,” in *ACM International Conference on Multimodal Interaction (ICMI)*, 2020, p. 305–314.
- [22] Eyben F., Wenginger F., Gross F., and Schuller B., “Recent developments in opensmile, the munich open-source multimedia feature extractor,” in *ACM International Conference on Multimedia*, 2013, p. 835–838.
- [23] F. Eyben et al., “The geneva minimalistic acoustic parameter set (gemaps) for voice research and affective computing,” *IEEE Trans. Affect. Comput.*, vol. 7(2), 190–202, 2016.
- [24] T. Baltrusaitis, A. Zadeh, Y. C. Lim, and L. Morency, “Openface 2.0: Facial behavior analysis toolkit,” in *IEEE International Conference on Automatic Face Gesture Recognition (FG)*, 2018, pp. 59–66.
- [25] M. Christ et al., “Time series feature extraction on basis of scalable hypothesis tests,” *Neurocomputing*, 2018.
- [26] Pedregosa F. et al., “Scikit-learn: Machine learning in Python,” *JMLR*, vol. 12, pp. 2825–2830, 2011.
- [27] W. Dupont and W. Plummer, “Power and sample size calculations. a review and computer program,” *Controlled clinical trials*, vol. 11, pp. 116–28, 05 1990.
- [28] M. Delacre, D. Lakens, and C. Leys, “Why psychologists should by default use welch’s t-test instead of student’s t-test,” *IRSP*, vol. 30(1), 92–101, 2017.