

# Affect-Aware Deep Belief Network Representations for Multimodal Unsupervised Deception Detection

Leena Mathur and Maja J Matarić

Department of Computer Science, University of Southern California

**Abstract**—Automated systems that detect the social behavior of deception can enhance human well-being across medical, social work, and legal domains. Labeled datasets to train supervised deception detection models can rarely be collected for real-world, high-stakes contexts. To address this challenge, we propose the first unsupervised approach for detecting real-world, high-stakes deception in videos without requiring labels. This paper presents our novel approach for affect-aware unsupervised Deep Belief Networks (DBN) to learn discriminative representations of deceptive and truthful behavior. Drawing on psychology theories that link affect and deception, we experimented with unimodal and multimodal DBN-based approaches trained on facial valence, facial arousal, audio, and visual features. In addition to using facial affect as a *feature* on which DBN models are trained, we also introduce a DBN training procedure that uses facial affect as an *aligner* of audio-visual representations. We conducted classification experiments with unsupervised Gaussian Mixture Model clustering to evaluate our approaches. Our best unsupervised approach (trained on facial valence and visual features) achieved an AUC of 80%, outperforming human ability and performing comparably to fully-supervised models. Our results motivate future work on unsupervised, affect-aware computational approaches for detecting deception and other social behaviors in the wild.

## I. INTRODUCTION

Recent developments in representation learning, affective computing, and social signal processing are advancing the creation of automated systems that can perceive human behaviors [51], [7], [40]. The social behavior of *deception* involves the intentional transmission of false or misleading information during social interactions [17], characterized as either high-stakes or low-stakes. Deceivers in high-stakes contexts face substantial consequences if their deception is discovered, in contrast to those in low-stakes contexts. Human deception detection ability across diverse contexts has been determined as close to chance level [9], motivating the creation of computational systems to help humans with this task. Healthcare providers, social workers, and legal groups can benefit from automated deception detection systems for applications that enhance human well-being (e.g., legal teams assessing courtroom testimonies of children who may be coerced to lie [4]; social workers and therapists recognizing when people are concealing abusive experiences [52]).

A key challenge in automated deception detection is the *scarcity* of labeled, real-world high-stakes deception data available for training models [41]. These data cannot be ethically obtained from lab experiments, because simulating realistic high-stakes scenarios with lab participants requires the use of threats to impose substantial consequences on deceivers. It is also difficult to collect large amounts of

labeled deception data with verifiable ground truth from real-world situations. It is worth noting that scarcity of labeled training data also poses a challenge for scientists and engineers seeking to model other affective and behavioral states; our paper contributes to addressing this broader problem.

We address the scarcity of labeled high-stakes deception data by proposing the first fully-unsupervised approach for detecting high-stakes deception in-the-wild, without requiring labeled data. We experimented with unsupervised *energy-based approaches* using Deep Belief Networks (DBN) [49], [46], to learn discriminative representations of truthful and deceptive behavior when trained on multimodal behavioral features. DBN-based approaches have shown potential for learning useful representations for human-centered tasks [29], [23], [36], [53]. To the best of our knowledge, our paper presents the first unsupervised DBN-based approaches for learning representations of a social behavior.

Prior approaches for deception detection in videos have primarily used features from visual, vocal, verbal, and physiological modalities [10]. There is limited research that includes affect as a modality for deception detection [32], [2]. Affect refers to neurophysiological states that function as components of emotions and moods. Inspired by psychology research on relationships between affect and social behavior [16], specifically relationships between affect and deception [13], [57], we chose to include affect in our models. Affective states are typically computationally represented along two dimensions: how pleasant or unpleasant a state is (valence) and how passive or active a state is (arousal) [44]. Since the face is a key channel through which affect is expressed [3] and a key channel of communication in the dataset that we use, our research was scoped to include *facial valence* and *facial arousal* as affect features. In this paper, we report on our experiments with unimodal and multimodal DBN models trained on features from facial valence, facial arousal, audio, and visual modalities, to identify effective modeling approaches. In addition to using facial affect as a *feature* on which our DBN models are trained, we also introduced a novel DBN training procedure to use facial affect as an *aligner* of audio-visual representations. In this paper, we refer to models that use affect as either a feature or an aligner as *affect-aware* models. Our research is driven by the questions: *To what extent will unsupervised DBN models be effective in learning discriminative representations of deceptive and truthful behavior from videos in-the-wild? How does facial affect contribute to the quality of DBN-based representations when used as a feature or aligner?*

To evaluate the usefulness of our DBN-based approaches for learning discriminative representations of deceptive and truthful behavior, we conducted classification experiments with unsupervised Gaussian Mixture Models (GMM). Our best unsupervised approach (trained on facial valence and visual features) achieved an AUC of 80%, outperforming human ability and performing comparably to fully-supervised models. The highest-performing representations for each DBN architecture included facial affect as either a *feature* or an *aligner*, supporting the inclusion of affect in models for learning discriminative representations of deceptive and truthful behavior. Our findings motivate future work on unsupervised, affect-aware machine learning approaches for detecting high-stakes deception and other social behaviors.

This work makes the following contributions:

- The first unimodal and multimodal DBN-based approaches to learn discriminative representations of a social behavior (deception in high-stakes situations).
- A novel unsupervised approach for high-stakes deception detection without requiring labeled data, serving as a proof-of-concept for future research on unsupervised models of deception, and other social behaviors.
- Introduction and analysis of effective modeling approaches for affect-aware unsupervised deception detection, using facial affect as a feature and as an aligner.

## II. RELATED WORKS

### A. Multimodal Deception Detection

Prior approaches for high-stakes deception detection in videos have largely relied on *supervised* machine learning models that exploit discriminative patterns in visual, vocal, verbal, and physiological cues to distinguish deceptive and truthful communication [10], [39], [5], [24], [43], [11], [55], [32]. Recent approaches have found affect to be a promising modality in deception detection [32], [2]. Approaches that include affect as a *feature* draw on psychology theories that deceivers in high-stakes situations will exhibit affective states with lower valence and higher arousal, compared to truthful speakers, expressed through patterns in nonverbal cues [13], [57]. We build on this research to treat affect not only as a *feature* on which models are trained, but also as an *aligner* of audio-visual representations of deception. Social psychology theories have proposed infusing affect in models of social behavior [16]; our research is driven by these insights. Our paper presents the first attempt at considering affect as an *aligner* of audio-visual feature representations for detecting a social behavior, including deception.

Existing high-stakes deception detection models have been trained on the only publicly-available, real-world high-stakes deception dataset that is the current benchmark: 121 videos of people speaking in courtroom situations (60 truthful and 61 deceptive videos, 28 seconds per video)[39]. Benchmark accuracies for this dataset [39] range from 60-75% with decision-trees trained on manually-annotated hand gestures, head movements, and facial features, in addition to automatically-extracted transcript n-grams. Prior approaches

have developed supervised deep classification models [11]; however use of deep classifiers has been considered inadvisable, due to the dataset’s small size [55]. In this paper, we use this benchmark dataset and compare our unsupervised GMM classifiers to approaches that use automated supervised classical machine learning models trained on interpretable features from human communication, with this dataset (**Table I**). Prior classifiers have included Decision Trees (DT), Support Vector Machines (SVM), Adaptive Boosting (AB), and Random Forests (RF).

TABLE I  
HIGHEST-PERFORMING AUTOMATED DECEPTION DETECTION  
APPROACHES USING SUPERVISED CLASSICAL MACHINE LEARNING  
MODELS TRAINED ON THE BENCHMARK DATASET [39].

Modality	Model	AUC	ACC
Visual, Verbal [39] (benchmark)	DT	—	75%
Visual [5]	SVM	—	77%
Visual, Vocal [43]	SVM	70%	—
Visual, Vocal, Verbal [55]	SVM	88%	—
Visual, Vocal, Verbal [24]	SVM	—	79%
Affect, Visual, Vocal [32]	AB	91%	84%
Visual, Vocal, Verbal [48]	RF	—	68%

To address the scarcity of labeled deception data for training models, researchers have developed supervised models that are more robust to small numbers of samples by discriminative feature subspaces [27], adversarial meta-learning [11], and data augmentation [11]. These approaches are constrained by dependency on labeled data to train models. Recently, [33] has demonstrated the use of domain adaptation to develop models trained on labeled low-stakes deception datasets that detect high-stakes deception; this, too, is constrained by dependence on labeled low-stakes deception data to train models. To the best of our knowledge, we introduce the first unsupervised approaches that detect high-stakes deception *without requiring any labels for training models*.

### B. Multimodal Deep Belief Networks

Deep learning models have demonstrated potential to learn useful representations of visual, verbal, and vocal communication [7], [8]. Models based on DBNs have emerged as promising approaches for multimodal representation learning. DBNs are probabilistic graphical models that are composed of layers of stochastic, latent variables [46]. As generative models [45], DBNs learn complex, non-linear dependencies across high-dimensional input feature spaces *without requiring labeled data*. DBNs can, therefore, be trained with fully-unsupervised algorithms. DBNs are formed by stacking Restricted Boltzmann Machines (RBMs), which are shallow 2-layer undirected probabilistic graphical models that contain a layer of visible units, a layer of hidden units, a weight matrix of connections between each visible unit and hidden unit, and biases for visible and hidden layers. RBMs assign an *energy* to each configuration of visible and hidden states and are trained to minimize this energy. In a standard DBN, the output of each RBM serves as the input for the next layer’s RBM; DBNs can be trained layer-wise with an efficient greedy algorithm [21], [47].

Multimodal DBNs are trained to learn joint representations across modalities that capture complex nonlinear dependencies in input data [49]. Prior research has used multimodal DBNs to combine audio-visual information across modalities through early fusion and late fusion. In *early fusion*, DBNs are trained on inputs concatenated across modalities. In *late fusion*, separate, unimodal DBN layers are trained to learn modality-specific representations before combining these to form a joint input for a multimodal DBN [36], [29], [23], [53]. These DBN approaches have shown the ability to learn useful representations for human-centered modeling tasks, including emotion recognition [29], pose estimation [36], and gesture recognition [53]. This paper presents the first multimodal DBN approaches for learning representations of a social behavior, specifically high-stakes deception.

A key challenge in multimodal representation learning is developing models that *align* information across modalities to learn richer representations [7]. The Kabsch algorithm provides a fundamental approach for aligning representations [26]; it estimates the best rotation matrix to align representations by minimizing their root mean squared deviation. Approaches based on the Kabsch algorithm have been developed to align representations for tasks across diverse fields, including graphics [50], speech [28], and natural language processing [18]. We contribute the first approach for affect-aligned DBN representation learning, adapting the Kabsch algorithm to learn audio, visual, and audio-visual DBN representations of a social behavior (deception) that are aligned with affect representations.

### III. METHODOLOGY

We experimented with affect-aware unimodal and multimodal DBN-based approaches for unsupervised high-stakes deception detection; models were trained on facial valence, facial arousal, audio, and visual features. In addition to using facial affect as a *feature* on which models are trained, we also developed a multimodal DBN training procedure that uses facial affect as an *aligner* of audio-visual representations.

#### A. High-Stakes Deception Dataset

We used the only publicly-available, real-world, high-stakes deception dataset of people speaking in 121 courtroom trial situations, discussed in Section II.A; additional details are in [39]. The creators of this dataset labeled each video as *truthful* or *deceptive*, based on information verified by police testimony and trials. We used their labels as ground truth to evaluate the effectiveness of our deception detection models. Videos were collected in unconstrained situations with a variety of illuminations, camera angles, camera movements, background settings, poses, and face occlusions. A usable subset, based on criteria defined by prior research with this dataset [55], [32], included 108 videos (53 truthful and 55 deceptive videos; 47 speakers of diverse race and gender).

#### B. Multimodal Feature Extraction

1) *Affect Features*: Similar to prior methods [32], we implemented an AffWildNet model [30] in TensorFlow [1]

trained on the Aff-Wild database to extract continuous representations of facial valence and facial arousal from speakers in the videos [56], [31]. The Aff-Wild data are labeled for continuous facial valence and facial arousal (between -1 and 1) at each facial frame, by 8 trained annotators. This data from situations in-the-wild included a variety of environments, poses, illuminations, and occlusions, and diverse speaker ethnicities. AffWildNet employs a CNN and RNN spatio-temporal architecture to predict valence and arousal at each facial frame. Tracking the faces of primary speakers, we extracted images of the facial bounding box at each video frame, resized images (96x96x3), and normalized pixel intensities between -1 and 1. Model weights came from the AffWildNet developers. We prepared Tensors with sequences of 80 consecutive images, per specified AffWildNet hyperparameters, and performed inference on these with AffWildNet to extract affect features. We extracted 2 facial affect features (valence and arousal) at each video frame.

2) *Audio Features*: We used the OpenSMILE toolkit [15] (version 2.0) to extract audio features from the eGeMAPs [14] and MFCC feature sets, guided by prior deception detection approaches [43], [55], [24], [32]. These features capture cepstral, spectral, prosodic, energy, and voice quality information. OpenSMILE’s eGeMAPs and MFCC feature sets overlap on MFCC features; we excluded duplicates. We extracted 58 audio features at each audio frame.

3) *Visual Features*: We used the OpenFace toolkit [6] (version 2.2.0) to extract visual features that captured eye gaze, facial action units, and head pose information, guided by prior deception detection approaches [43], [55], [24], [32], [5]. We extracted 31 visual features at each video frame.

4) *Temporal Representations*: Similar to prior methods [43], [32], we chose to represent each affect, audio, and visual feature as a fixed-length vector of time-series attributes to capture their temporal behavior along each variable-length video. We computed 17 time-series attributes for each feature: mean, median, standard deviation, minimum, maximum, kurtosis, mean and median autocorrelation across 1-second time lags, and changes in feature values across quantiles. A total of 1547 multimodal features were computed to represent each video (17 facial valence, 17 facial arousal, 527 visual, and 986 audio features).

#### C. Multimodal DBN Representation Learning

We extracted low-dimensional representations of each video by training DBNs with 4 modeling approaches: unimodal networks, multimodal early fusion networks, multimodal late fusion networks, and multimodal affect-aligned networks. Conceptual diagrams visualizing our approaches are in **Figure 1** and described in this section. The goal of our DBN approaches is to capture complex, non-linear dependencies in visible, behavioral input data by learning lower dimensional, hidden representations.

1) *DBN Training*: Each of our DBN approaches follow the same underlying DBN training procedure. DBNs have stacks of RBMs that each have one visible layer  $\mathbf{v}$  and one hidden layer  $\mathbf{h}$ . The energy [22] of a joint configuration  $(\mathbf{v},$

$\mathbf{h}$ ), given the weight matrix and bias parameters  $\theta = (\mathbf{W}, \mathbf{a}, \mathbf{b})$ , can be expressed with the following equation [20], [21]:

$$E(v, h; \theta) = - \sum_{i \in \text{visible}} a_i v_i - \sum_{j \in \text{hidden}} b_j h_j - \sum_{i,j} v_i h_j W_{ij} \quad (1)$$

This energy function can be used to express the joint probability of  $\mathbf{v}$  and  $\mathbf{h}$  as follows:

$$P(v, h; \theta) = \frac{1}{Z(\theta)} e^{-E(v, h; \theta)} \quad (2)$$

$$Z(\theta) = \sum_{v, h} e^{-E(v, h; \theta)} \quad (3)$$

During training, we draw on [20] to implement the contrastive divergence (CD) algorithm [21], approximating the gradient through Gibbs sampling. We use the stochastic gradient descent optimization algorithm to update the model's parameters. We train each DBN with a greedy, layerwise approach, detailed in [21], [47]. To allow comparisons of our DBN approaches, we kept the random seed (0) and hyperparameters constant when training models (learning rate 0.01, 10 CD iterations, 200 epochs, batch size 32).

2) *DBN Architectures*: To examine trends across modalities, we report results from 2 and 4-dimensional embeddings for multiple DBN architectures. For unimodal, multimodal early fusion, and affect-aligned representations, we trained DBNs with 8 different architectures: 2 RBMs were trained to learn hidden representations of size 2 and size 4; 6 DBNs were trained with stacked RBM architectures of 128-64-2, 256-128-2, 512-256-2, 128-64-4, 256-128-4, and 512-256-4. For multimodal late fusion representations, we trained 6 DBNs with the stacked RBM architectures, listed above.

3) *Unimodal DBN*: We trained unimodal DBN models on the respective input features of each modality to explore the effectiveness of learning representations of individual modalities for unsupervised deception detection.

4) *Multimodal Early Fusion DBN*: We experimented with early fusion DBNs which initially concatenate features from different modalities to form a single input on which DBNs

are trained. We trained early fusion DBNs on each of the 11 combinations of the 4 modalities. Early fusion experiments were used to explore the effectiveness of exploiting low-level interactions of multimodal features to learn representations for unsupervised deception detection.

5) *Multimodal Late Fusion DBN*: We experimented with late fusion DBNs which train a separate unimodal RBM layer before concatenating the unimodal representations into a new input for joint DBN layers. Late fusion experiments were used to explore the effectiveness of exploiting higher-level interactions of multimodal feature spaces to learn representations for unsupervised deception detection.

6) *Multimodal Affect-Aligned DBN*: We formulated an approach for affect-aligned DBN representation learning by adapting the Kabsch algorithm [26] to learn audio, visual, and audio-visual DBN representations that are aligned with affect. For each affect-aligned approach, we concurrently train 2 DBN models:  $DBN_{AV}$  is trained on audio, visual, or audio-visual input features and  $DBN_{AFF}$  is trained on affect input features. During DBN training, we successively align the hidden representations of  $DBN_{AV}$  with the corresponding layer of hidden representations in  $DBN_{AFF}$ . At each hidden layer  $h_1 \dots h_n$  of both DBNs,  $DBN_{AV}$  will have a  $D$ -dimensional hidden representation  $X = (x^i)_{i=1}^D$ , and  $DBN_{AFF}$  will also have a  $D$ -dimensional hidden representation  $A = (a^i)_{i=1}^D$ . We compute an optimal alignment matrix  $\phi$  to align  $X$  with  $A$  per the following:

$$\phi = \underset{\phi}{\operatorname{argmin}} \|\phi X - A\| \quad (4)$$

where  $\|\cdot\|$  denotes the Frobenius norm. We constrained this problem to consider rotation matrices as solutions, allowing us to use the Kabsch algorithm [26] to solve the equation. The covariance matrix  $H$  of  $X$  and  $A$  is computed as follows:

$$H = [X - \operatorname{centroid}(X)]^T [A - \operatorname{centroid}(A)] \quad (5)$$

The singular value decomposition of  $H$  is then computed to express  $H$  as a product of matrices  $U$ ,  $S$ , and  $V^T$ . The




Fig. 1. Visualization of Unimodal and Multimodal DBN Approaches.

optimal rotation matrix is then computed with the following:

$$\phi = V \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \text{sign}(\det(VU^T)) \end{bmatrix} U^T \quad (6)$$

We use  $\phi$  to align  $X$  with  $A$  during DBN training, resulting in DBN audio, visual, and audio-visual hidden representations that are successively aligned with affect at each DBN layer. Affect-alignment was efficiently implemented with NumPy and PyTorch [37]. We trained 9 affect-aligned DBNs that used (1) facial valence, (2) facial arousal, and (3) both facial valence and facial arousal to align audio, visual, and audio-visual modalities. These experiments were used to explore the effectiveness of affect as an *aligner* of audio-visual representations for unsupervised deception detection.

#### D. Unsupervised Classification Experiments

To explore the extent to which our DBN representation learning approaches can be used for deception detection, we experimented with unsupervised GMMs [19] implemented in scikit-learn [38] to find clusters in representations. We compared the cluster assignments with the ground truth *deceptive* and *truthful* video labels to evaluate the performance of our models. Our process is visualized in **Figure 2**.

GMMs have been widely used to evaluate representations during human-centered modeling tasks for speech [42] and emotions [34], among other domains. Our clustering approach assumes that samples are generated through a mixture of two Gaussian distributions that are estimated with the Expectation-Maximization (EM) algorithm [35]. We evaluate the quality of our representations by their effectiveness when clustering deceptive and truthful behavior; high-quality representations have a *naturally clusterable* property [8] (samples that come from the same class should be contained within distinguishable clusters in the representation space). We kept the random seed (1) and hyperparameters constant when training all GMM clustering models (diagonal covariance matrices, maximum of 100 EM iterations). Our methodology is similar to [54], which experimented with unsupervised clustering algorithms for stress detection.

Consistent with prior research [11], [32], [43], [55], we ensured that the same person was not in both the training and testing sets of any experiment. All experiments were conducted with 5-fold stratified cross-validation, split across

the 47 speaker identities, and repeated 10 times (50 cross-validation fold experiments). For feature normalization, training sets and testing sets for each cross-validation fold were scaled between 0 and 1 by the distribution of the training set, to prevent information from leaking into the testing sets.

#### E. Evaluation Metrics and Baselines

Aligned with prior research [32], [55], [43], we use *AUC* as the primary metric to evaluate unsupervised deception detection classifiers: the probability of the classifier ranking a randomly chosen deceptive sample higher than a randomly chosen truthful one. We computed this across 50 cross-validation folds; results are in **Table II**. Significance values of differences in model performances across representations were computed with McNemar’s chi-squared test with continuity correction ( $\alpha=0.05$ ) [12]. We also compute *ACC*, the classification accuracy across the videos, and *precision*, the proportion of deceptive videos among the videos classified as deceptive. Due to space limitation, we report accuracy and precision for the highest-performing approach.

To compare the effectiveness of our DBN approaches to a simpler dimensionality-reduction approach, we defined a *principal component analysis (PCA) baseline*, per prior methodology [25]. We learned and clustered 2 and 4-dimensional PCA embeddings for each unimodal and multimodal combination of features, following the same experimental setup and cross-validation procedures.

To compare our models to human deception detection ability at chance level, we defined a *human performance baseline* as a classifier that always predicts *deceptive* for 55 deceptive videos out of 108 videos (51% accuracy) [32], [33]. It is worth noting that human accuracy on this dataset, averaged across audio, visual, and audio-visual modalities, has been determined as 52.78% [39] from 3 human annotators and 70.07% from a separate study with 3 different human annotators [48]. The difference in these two measurements demonstrates the challenge of human deception detection.

## IV. RESULTS AND DISCUSSION

AUC results from the classification experiments are presented in **Table II**. The highest-performing representation for each DBN architecture (column) is highlighted in orange. *Across all DBN architectures, the highest-performing representations always included facial affect, supporting the inclusion of affect as a modality when learning discriminative DBN representations for unsupervised high-stakes deception detection.* Our best representations were learned by DBNs




Fig. 2. Process for Unsupervised Deception Detection.

TABLE II  
CLASSIFICATION RESULTS (AUC)

Modality	Hidden Representation Size: 2				Hidden Representation Size: 4			
	2 (RBM)	128_64	256_128	512_256	4 (RBM)	128_64	256_128	512_256
<b>Unimodal Representations</b>								
Arousal	0.47	0.47	0.50	0.54	0.51	0.48	0.53	0.46
Audio	0.46	0.43	0.45	0.51	0.48	0.47	0.41	0.48
Valence	0.68	0.65	0.67	0.60	0.56	0.69	0.65	0.62
Visual	0.65	0.59	0.65	0.59	0.70	0.53	0.61	0.61
<b>Multimodal Early Fusion Representations</b>								
Arousal + Audio	0.50	0.50	0.45	0.49	0.48	0.48	0.49	0.41
Arousal + Valence	0.50	0.69	0.61	0.74	0.56	0.73	0.62	0.70
Arousal + Visual	0.56	0.64	0.70	0.65	0.59	0.57	0.52	0.54
Audio + Valence	0.48	0.46	0.52	0.43	0.48	0.50	0.46	0.45
Audio + Visual	0.52	0.50	0.47	0.54	0.40	0.50	0.46	0.49
Valence + Visual	0.62	0.61	0.63	0.47	0.71	0.55	0.52	0.60
Arousal + Audio + Valence	0.46	0.48	0.39	0.47	0.48	0.49	0.53	0.45
Arousal + Audio + Visual	0.51	0.48	0.58	0.46	0.44	0.42	0.51	0.49
Arousal + Valence + Visual	0.64	0.43	0.71	0.71	0.69	0.51	0.61	0.62
Audio + Valence + Visual	0.52	0.47	0.52	0.46	0.46	0.44	0.45	0.47
All Modalities	0.49	0.51	0.49	0.53	0.40	0.53	0.48	0.55
<b>Multimodal Late Fusion Representations</b>								
Arousal + Audio	-	0.46	0.52	0.50	-	0.38	0.44	0.47
Arousal + Valence	-	0.58	0.73	0.49	-	0.50	0.67	0.56
Arousal + Visual	-	0.56	0.71	0.76	-	0.54	0.73	0.67
Audio + Valence	-	0.45	0.52	0.58	-	0.48	0.45	0.50
Audio + Visual	-	0.47	0.47	0.59	-	0.49	0.51	0.53
Valence + Visual	-	0.56	0.64	0.80*	-	0.52	0.75	0.70
Arousal + Audio + Valence	-	0.50	0.49	0.50	-	0.46	0.47	0.46
Arousal + Audio + Visual	-	0.45	0.63	0.63	-	0.59	0.43	0.55
Arousal + Valence + Visual	-	0.62	0.79	0.62	-	0.62	0.69	0.77
Audio + Valence + Visual	-	0.58	0.52	0.63	-	0.48	0.58	0.53
All Modalities	-	0.50	0.49	0.50	-	0.46	0.47	0.46
<b>Multimodal Affect-Aligned Representations</b>								
Arousal + Audio	0.50	0.50	0.54	0.52	0.48	0.54	0.51	0.51
Arousal + Visual	0.69	0.43	0.56	0.51	0.70	0.50	0.38	0.52
Arousal + Audio-Visual	0.44	0.60	0.50	0.49	0.56	0.47	0.52	0.52
Valence + Audio	0.50	0.50	0.56	0.55	0.48	0.46	0.44	0.54
Valence + Visual	0.62	0.48	0.48	0.35	0.66	0.51	0.75	0.45
Valence + Audio-Visual	0.43	0.50	0.50	0.60	0.47	0.51	0.56	0.57
Arousal & Valence + Audio	0.45	0.47	0.46	0.55	0.48	0.44	0.48	0.44
Arousal & Valence + Visual	0.67	0.62	0.54	0.56	0.59	0.64	0.53	0.55
Arousal & Valence + Audio-Visual	0.44	0.59	0.62	0.62	0.41	0.47	0.45	0.44

with a 512–256–2 architecture trained with multimodal late fusion on facial valence and visual features. Unsupervised GMM classifiers trained on these representations achieved an AUC of 80% (accuracy of 70% and precision of 88%), outperforming the 51% human performance baseline and the corresponding PCA baseline of 54% ( $p < 0.01$ ). Our results demonstrate the potential for DBN-based models to effectively learn representations of deceptive and truthful behavior for unsupervised deception detection.

#### A. Analysis of Unimodal Representations

Unimodal classification results revealed that models trained on facial valence and visual representations significantly outperformed models trained on facial arousal and audio representations across DBN architectures ( $p < 0.001$ ). Differences in performance are visualized in **Figure 3**. The highest AUCs achieved by classifiers trained on unimodal visual, facial valence, facial arousal, and audio representations were 70%, 69%, 54%, and 51%, respectively (highlighted in blue in **Table II**). Visualizations of these classifier performances compared to their corresponding PCA baselines

are also in **Figure 3**. Classifiers trained on the highest-performing visual, facial valence, and facial arousal representations significantly outperformed their corresponding PCA baselines of 49%, 49%, and 43%, respectively ( $p < 0.001$ ,  $p < 0.001$ ,  $p < 0.05$ ). Our results suggest that unimodal DBN models can effectively learn discriminative representations of deceptive and truthful behavior. Results indicate that facial affect contributes to the quality of DBN representations when used as a feature for high-stakes deception detection.

Classifiers trained on the highest-performing audio representations did not significantly outperform the corresponding PCA baseline of 50%, indicating that audio representations were not as discriminative for detecting high-stakes deception in this dataset’s context (courtroom trials). It is worth noting that audio cues have been useful for detecting deception in other contexts [55], [32], [43], [24].

Classifiers trained on facial valence representations outperformed classifiers trained on visual representations across all DBN architectures, with the exception of the 4-dimensional visual RBM representation. Comparable performances of facial valence and visual representations suggest that they have




Fig. 3. Distribution of AUC performances achieved by unimodal representations across all RBM/DBM architectures.

comparable discriminative potential for use in unsupervised deception detection.

*Our findings from these unimodal classification experiments support the inclusion of affect, specifically facial valence, as a feature to be used in DBN-based representation learning for high-stakes deception detection in the wild.*

### B. Analysis of Multimodal Representations

Across the 8 DBN architectures, multimodal representations outperformed all unimodal representations, demonstrating that multimodal models are more capable of learning useful representations of human behavior than unimodal ones [7]. Four modality combinations, highlighted in orange, emerged as most effective: (valence + visual), (arousal + valence + visual), (arousal + valence), (arousal + visual). The highest-performing multimodal representations significantly outperformed their corresponding PCA baselines ( $p < 0.01$ ), visualized in **Figure 4**. In all of the highest-performing representations across each DBN architecture, affect was either a *feature* or an *aligner* within the modality combination.




Fig. 4. AUC performance achieved by highest-performing multimodal DBN representations relative to their baselines.

We examined the performance of our models from the 3 multimodal representation learning approaches: early fusion, late fusion, and affect-aligned representations. Across the 8 DBN architectures, early fusion yielded the highest performance 3 times, late fusion yielded the highest performance 4 times, and affect-aligned representations yielded the highest performance 2 times (there was one tie).

*Our findings provide support for the inclusion of affect as a feature or aligner in multimodal machine learning approaches to learn useful representations of human behavior.*

Comparing our results to those obtained from prior fully-supervised high-stakes deception detection approaches (Table I), we demonstrate the potential for modeling high-stakes deception with unsupervised, affect-aware DBN representation learning. While some prior fully-supervised models outperform our unsupervised models, our highest-performing AUC of 80% falls within the range achieved by the highest-performing prior results (70%–91%). All these prior models required labeled datasets for training; *our approach has the advantage that it does not require any labeled data.*

## V. CONCLUSION

Our paper addresses the scarcity of labeled, real-world high-stakes deception data for training models, by developing the first unsupervised approaches that detect high-stakes deception without requiring labels. We experimented with unimodal and multimodal DBN-based approaches that were trained on facial valence, facial arousal, audio, and visual features to identify effective modeling approaches for unsupervised deception detection. In addition to using facial affect as a *feature* on which DBN models are trained, we also introduced a multimodal DBN training procedure that used facial affect as an *aligner* of audio-visual representations. We demonstrated that unsupervised DBN approaches can learn discriminative, low-dimensional representations of deceptive and truthful behavior. Our best unsupervised approach achieved an AUC of 80%, accuracy of 70%, and precision of 88% (trained on facial valence and visual features), outperforming human ability and performing comparably to a number of existing fully-supervised high-stakes deception detection models. To the best of our knowledge, our research also represents the first attempt at training DBNs to learn discriminative representations of a social behavior.

In all of the highest-performing approaches across each DBN architecture, affect emerged as a contributor, functioning either as a feature or as an aligner, supporting the inclusion of facial affect as a modality when learning discriminative DBN representations for unsupervised high-stakes deception detection. While this paper used facial affect, rich affect features may also be gained from the audio modality, motivating future work in this direction.

Our findings serve as a proof-of-concept and motivation for future research in developing unsupervised, affect-aware machine learning approaches for representing and detecting deception and other social behaviors during interaction contexts in real-world situations. We are also encouraged by the findings to proceed beyond our paper’s research context to examine the role of affect as a feature and aligner in models of human behavior, beyond deception, in-the-wild.

## VI. ACKNOWLEDGMENTS

USC Provost’s Undergraduate Research Fellowship.

## REFERENCES

- [1] M. Abadi et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] S. Amiriparian, J. Pohjalainen, E. Marchi, S. Pugachevskiy, and B. Schuller. Is deception emotional? an emotion-driven predictive approach. In *Interspeech 2016*, pages 2011–2015, 2016.

- [3] B. App, D. McIntosh, C. Reed, and M. J. Hertenstein. Nonverbal channel use in communication of emotion: how may depend on why. *Emotion*, 11 3:603–17, 2011.
- [4] V. Ardulov, Z. Durante, S. Williams, T. Lyon, and S. Narayanan. Identifying truthful language in child interviews. In *IEEE Int. Conf. on Acoust., Speech and Signal Process.*, 2020.
- [5] D. Avola et al. Automatic deception detection in rgb videos using facial action units. In *ACM Int. Conf. on Distrib. Smart Cameras*, 2019.
- [6] T. Baltrušaitis, A. Zadeh, Y. C. Lim, and L.-P. Morency. Openface 2.0: Facial behavior analysis toolkit. In *IEEE Int. Conf. on Autom. Face & Gesture Recognition*, pages 59–66. IEEE, 2018.
- [7] T. Baltrušaitis, C. Ahuja, and L.-P. Morency. Multimodal machine learning: A survey and taxonomy, 2017.
- [8] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE Trans. on Pattern Anal. and Mach. Intell.*, 35(8):1798–1828, 2013.
- [9] C. Bond and B. Depaulo. Accuracy of deception judgments. *Personality and Social Psychol. Rev.*, 10(3), 214–34, 2006.
- [10] M. Burzo, M. Abouelenien, V. Perez-Rosas, and R. Mihalcea. *Multimodal Deception Detection*, page 419–453. 2018.
- [11] M. Ding, A. Zhao, Z. Lu, T. Xiang, and J. Wen. Face-focused cross-stream network for deception detection in videos. In *IEEE Conf. on Comput. Vision and Pattern Recognition*, pages 7794–7803, 2019.
- [12] W. Dupont and W. Plummer. Power and sample size calculations. a rev. and computer program. *Controlled Clin. Trials*, 11:116–28, 05 1990.
- [13] P. Ekman and W. V. Friesen. Nonverbal leakage and clues to deception. *Psychiatry*, 32(1):88–106, 1969.
- [14] F. Eyben et al. The geneva minimalistic acoustic parameter set (gemaps) for voice research and affective computing. *IEEE Trans. on Affect. Comput.*, 7(2):190–202, 2015.
- [15] F. Eyben, M. Wöllmer, and B. Schuller. Opensmile: The munich versatile and fast open-source audio feature extractor. In *ACM Int. Conf. on Multimedia*, MM '10, page 1459–1462, 2010.
- [16] J. P. Forgas. Mood and judgment: the affect infusion model (aim). *Psychological Bull.*, 117(1):39, 1995.
- [17] M. Frank and P. Ekman. The ability to detect deceit generalizes across different types of high-stake lies. *Personality and Social Psychol.*, 72:1429–39, 07 1997.
- [18] G. Glavaš and I. Vulić. Non-linear instance-based cross-lingual mapping for non-isomorphic embedding spaces. In *Annu. Meeting of the Assoc. for Comput. Linguistics*, pages 7548–7555, 2020.
- [19] V. Hasselblad. Estimation of parameters for a mixture of normal distributions. *Technometrics*, 8(3):431–444, 1966.
- [20] G. E. Hinton. A practical guide to training restricted boltzmann machines. In *Neural networks: Tricks of the trade*, pages 599–619. Springer, 2012.
- [21] G. E. Hinton, S. Osindero, and Y.-W. Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7):1527–1554, 2006.
- [22] J. J. Hopfield. Neural networks and physical symp. with emergent collective computational abilities. *Proc. of the Nat. Acad. of Sciences*, 79(8):2554–2558, 1982.
- [23] J. Huang and B. Kingsbury. Audio-visual deep learning for noise robust speech recognition. In *IEEE Int. Conf. on Acoust., Speech and Signal Process.*, pages 7596–7599, 2013.
- [24] M. Jaiswal, S. Tabibu, and R. Bajpai. The truth and nothing but the truth: Multimodal analysis for deception detection. In *IEEE Int. Conf. on Data Mining Workshops*, pages 938–943, 2016.
- [25] N. Jaques, S. Taylor, A. Sano, and R. Picard. Multimodal autoencoder: A deep learning approach to filling in missing sensor data and enabling better mood prediction. In *Int. Conf. on Affect. Comput. and Intell. Interaction*, pages 202–208. IEEE, 2017.
- [26] W. Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Sect. A: Crystal Physics, Diffraction, Theor. and General Crystallography*, 32(5):922–923, 1976.
- [27] H. Karimi, J. Tang, and Y. Li. Toward end-to-end deception detection in videos. In *IEEE Int. Conf. on Big Data*, pages 1278–1283, 2018.
- [28] W. B. Kheder, D. Matrouf, M. Ajili, and J.-F. Bonastre. Iterative bayesian and mmse-based noise compensation techniques for speaker recognition in the i-vector space. In *Odyssey 2016*, pages 60–67, 2016.
- [29] Y. Kim, H. Lee, and E. M. Provost. Deep learning for robust feature generation in audiovisual emotion recognition. In *IEEE Int. Conf. on Acoust., Speech and Signal Process.*, pages 3687–3691, 2013.
- [30] D. Kollias et al. Deep affect prediction in-the-wild: Aff-wild database and challenge, deep architectures, and beyond. *Int. J. of Comput. Vision*, pages 1–23, 2019.
- [31] D. Kollias, M. A. Nicolaou, I. Kotsia, G. Zhao, and S. Zafeiriou. Recognition of affect in the wild using deep neural networks. In *IEEE Conf. on Comput. Vision and Pattern Recognition Workshops*, pages 1972–1979. IEEE, 2017.
- [32] L. Mathur and M. J. Mataríć. Introducing representations of facial affect in automated multimodal deception detection. In *Int. Conf. on Multimodal Interaction*, page 305–314, 2020.
- [33] L. Mathur and M. J. Mataríć. Unsupervised audio-visual subspace alignment for high-stakes deception detection. In *IEEE Int. Conf. on Acoust., Speech and Signal Process.*, pages 2255–2259, 2021.
- [34] A. Metallinou, S. Lee, and S. Narayanan. Audio-visual emotion recognition using gaussian mixture models for face and voice. In *IEEE Int. Symp. on Multimedia*, pages 250–257, 2008.
- [35] T. K. Moon. The expectation-maximization algorithm. *IEEE Signal Process. Mag.*, 13(6):47–60, 1996.
- [36] W. Ouyang, X. Chu, and X. Wang. Multi-source deep learning for human pose estimation. In *IEEE Conf. on Comput. Vision and Pattern Recognition*, page 2337–2344, USA, 2014.
- [37] A. Paszke et al. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Inf. Process. Symp.* 32, pages 8024–8035. 2019.
- [38] F. Pedregosa et al. Scikit-learn: Machine learning in Python. *J. of Mach. Learn. Res.*, 12:2825–2830, 2011.
- [39] V. Pérez-Rosas et al. Deception detection using real-life trial data. In *ACM Int. Conf. on Multimodal Interaction*, page 59–66, 2015.
- [40] R. W. Picard. *Affective Computing*. MIT press, 2000.
- [41] S. Porter and L. ten Brinke. The truth about lies: What works in detecting high-stakes deception? *Legal and Criminological Psychol.*, 15(1):57–75, 2010.
- [42] D. Povey et al. Subspace gaussian mixture models for speech recognition. In *IEEE Int. Conf. on Acoust., Speech and Signal Process.*, pages 4330–4333, 2010.
- [43] R. Rill-García et al. High-level features for multimodal deception detection in videos. *IEEE Conf. on Comput. Vision and Pattern Recognition Workshops*, pages 1565–1573, 2019.
- [44] J. A. Russell. A circumplex model of affect. *Personality and Social Psychol.*, 39(6):1161, 1980.
- [45] R. Salakhutdinov. Learning deep generative models. *Annu. Rev. of Statist. and Its Application*, 2:361–385, 2015.
- [46] R. Salakhutdinov and G. Hinton. Deep boltzmann machines. In *Artif. Intell. and Statist.*, pages 448–455. PMLR, 2009.
- [47] R. Salakhutdinov and G. Hinton. An efficient learning procedure for deep boltzmann machines. *Neural Computation*, 24(8):1967–2006, 2012.
- [48] U. M. Sen, V. Perez-Rosas, B. Yanikoglu, M. Abouelenien, M. Burzo, and R. Mihalcea. Multimodal deception detection using real-life trial data. *IEEE Trans. on Affect. Comput.*, pages 1–1, 2020.
- [49] N. Srivastava and R. Salakhutdinov. Multimodal learning with deep boltzmann machines. In *NIPS*, volume 1, page 2. Citeseer, 2012.
- [50] J. Taylor, J. Shotton, T. Sharp, and A. Fitzgibbon. The vitruvian manifold: Inferring dense correspondences for one-shot human pose estimation. In *IEEE Conf. on Comput. Vision and Pattern Recognition*, pages 103–110. IEEE, 2012.
- [51] A. Vinciarelli, M. Pantic, H. Bourlard, and A. Pentland. Social signal processing: State-of-the-art and future perspectives of an emerging domain. In *ACM Int. Conf. on Multimedia*, MM '08, page 1061–1070, 2008.
- [52] A. Vrij, L. Akehurst, and S. Knight. Police officers’, social workers’, teachers’ and the general public’s beliefs about deception in children, adolescents and adults. *Legal and Criminological Psychol.*, 11(2):297–312, 2006.
- [53] D. Wu and L. Shao. Multimodal dynamic networks for gesture recognition. In *ACM Int. Conf. on Multimedia*, MM '14, page 945–948, 2014.
- [54] Y. Wu, M. Daoudi, A. Amad, L. Sparrow, and F. D’Hondt. Unsupervised learning method for exploring students’ mental stress in medical simulation training. In *Companion Publication of the 2020 Int. Conf. on Multimodal Interaction*, page 165–170, 2020.
- [55] Z. Wu et al. Deception detection in videos. In *AAAI Conf. on Artif. Intell.*, pages 1695–1702, 2018.
- [56] S. Zafeiriou, D. Kollias, M. A. Nicolaou, A. Papaioannou, G. Zhao, and I. Kotsia. Aff-wild: Valence and arousal ‘in-the-wild’ challenge. In *IEEE Conf. on Comput. Vision and Pattern Recognition Workshops*, pages 1980–1987. IEEE, 2017.
- [57] M. Zuckerman, B. M. DePaulo, and R. Rosenthal. Verbal and nonverbal communication of deception. In *Advances in Exp. Social Psychol.*, volume 14, pages 1–59. Elsevier, 1981.