

EVA: Generating Emotional Behavior of Virtual Agents using Expressive Features of Gait and Gaze

Tanmay Randhavan
University of North Carolina
Chapel Hill, USA
tanmay@cs.unc.edu

Rahul Sheth
Snap Inc.
Santa Monica, USA
rahul@snap.com

Aniket Bera
University of North Carolina
Chapel Hill, USA
ab@cs.unc.edu

Kurt Gray
University of North Carolina
Chapel Hill, USA
kurtgray@unc.edu

Kyra Kapsaskis
University of North Carolina
Chapel Hill, USA
kyrakaps@gmail.com

Dinesh Manocha
University of Maryland
College Park, USA
dm@cs.umd.edu

ABSTRACT

We present a novel, real-time algorithm, *EVA*, for generating virtual agents with various perceived emotions. Our approach is based on using *Expressive Features* of gaze and gait to convey emotions corresponding to happy, sad, angry, or neutral. We precompute a data-driven mapping between gaits and their perceived emotions. *EVA* uses this gait emotion association at runtime to generate appropriate walking styles in terms of gaits and gaze. Using the *EVA* algorithm, we can simulate gaits and gazing behaviors of hundreds of virtual agents in real-time with known emotional characteristics. We have evaluated the benefits in different multi-agent VR simulation environments. Our studies suggest that the use of expressive features corresponding to gait and gaze can considerably increase the sense of presence in scenarios with multiple virtual agents.

ACM Reference Format:

Tanmay Randhavan, Aniket Bera, Kyra Kapsaskis, Rahul Sheth, Kurt Gray, and Dinesh Manocha. 2019. *EVA: Generating Emotional Behavior of Virtual Agents using Expressive Features of Gait and Gaze*. In *Proceedings of ACM Conference (Symposium on Applied Perception)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

In recent years, interest in developing intelligent virtual agents has increased because of the applications in many areas including gaming and social VR [19], evacuation simulations [34], therapy and rehabilitation environments [56], training, etc. These applications require simulation of virtual agents that have personalities, moods, emotions and exhibit human-like behaviors. Social robots, too, require human-like appearance and behaviors to elicit more empathy from humans interacting with them [54]. In the physical world, humans communicate emotions and moods using speech, gestures, and facial expressions. Non-verbal expressions, including body posture and movements, are also important and used for social interactions. Virtual agents and robots that can communicate using both verbal and non-verbal channels are, therefore, important for human-computer and human-robot interactions.

Emotions influence our interactions with the environment and other agents [6]. The perception of someone's emotional state helps

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Symposium on Applied Perception, September 2019, Barcelona, Spain

© 2019 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>




Figure 1: We present a novel, real-time algorithm, *EVA*, for generating virtual agents with various emotions. We compute a novel data-driven gait emotion association and combine them with gazing behaviors to create virtual agents with different emotions. To simulate gazing, we use expressive features such as direct (Happy, Angry) vs. averted gazing (Sad) and neck flexion (Sad) vs. extension (Happy). To simulate gaits, we use expressive features such as body expansion (Happy), compact posture (Sad), rapid and increased movements (Angry, Happy), and slow movements (Sad).

to define how we view their personality and how we react to their actions [46]. Therefore, it is crucial to incorporate the emotional context in the simulation of intelligent virtual agents. In this paper, we mainly focus on conveying emotions of virtual agents based on gait and gaze.

Psychology literature refers to the features that are used to communicate emotional states through nonverbal movements as *Expressive Features* [55]. These features include facial expressions, body postures, and movements [27]. During walking, an observer uses the style of walking (i.e., gait) to assist the perception of emotions [57]. Gait features such as the extent of arm swings, the foot landing force, length of the strides, and erectness of the posture have been shown to assist in the identification of emotions such as sadness, anger, happiness, and pride [57]. For example, faster walking combined with high arm swings is perceived as happy, and a collapsed upper body with less movement activity is considered as sad [27].

In addition to gaits, gazing behaviors also affect the perception and expression of emotions [20]. As one of the critical components of non-verbal communication [5], gaze increases the plausibility of embodied conversational agents (ECA) [49]. A direct versus an averted gaze also changes the perceptions of emotions; direct gaze is associated with anger, whereas averted gaze is associated with sadness [3]. In this paper, we restrict ourselves to expressive features of gait and gaze features.

Main Results: We present two novel contributions:

1) A data-driven metric between gaits and perceived emotions based on a user study. We chose *happy*, *angry*, *sad*, and *neutral* emotions because they are known to persist for a prolonged interval of time and are more apparent during walking [38]. Using the continuous space representation of emotions [18, 28, 40, 41], we can combine these four basic emotions to represent all other emotions of virtual agents [42, 43]. This gait emotion association can also be used to predict the perceived emotion of any new input gait with an accuracy of 70.04%.

2) **EVA:** A novel, data-driven approach for generating emotions in virtual agents, *EVA*, using the expressive features of gait and gaze (Figure 1). *EVA* algorithm uses the gait emotion association to generate gaits corresponding to different emotions. We combine these gaits with gazing behaviors and create virtual agents that are perceived as experiencing *happy*, *angry*, *sad*, or *neutral* emotions.¹

We validate our algorithms using standard metrics of co-presence [21]. The questionnaire measures the level of presence achieved by a real user immersed in a virtual environment, based on the emotional state of virtual agents generated using *EVA*. Our results indicate that in scenarios with a high number of virtual agents, expressive features can increase the sense of presence felt by the users in the virtual environment. Our results also indicate that both gait and gazing features contribute to the perceptions of emotions in virtual agents. We observe that *EVA* simulates virtual agents with desired emotions with an accuracy of 70.83%, and achieves 100% accuracy in terms of simulating emotion corresponding to sadness.

This paper is organized as follows. In Section 2, we present related work on expressive features and emotions of virtual agents. In Section 3, we give a brief overview of our virtual agent generation algorithm and present the details of the gait emotion association. In Section 4, we present the details of our *EVA* algorithm, which uses the expressive features to generate emotions in virtual agents. We present the details of our user evaluation in Section 5.

2 RELATED WORK

In this section, we give a brief overview of related work on expressive features and emotions of virtual agents.

2.1 Expressive Features of Virtual Agents

Emotional states can be communicated through expressive features such as nonverbal movements and gestures, including facial expressions [55]. Ferstl et al. [19] studied the effect of the manipulation of facial features for trait portrayal in virtual agents. Previous research has used trajectories to convey dominance [52], approachability [53], group entitativity [8], and personalities [9]. Biometric data can also be used for predicting the emotions of individuals [66]. In this paper, we focus on the expressive features of gait and gaze movements. Different methods have been proposed to simulate the gait and gaze of virtual agents to convey a variety of behaviors in virtual agents. Narang et al. [44] studied self-recognition in virtual agents with walking gaits and avatars of real humans. Gaits have been used to convey friendliness [50] and dominance [51] of virtual agents. Gazing has been a focus of attention in many VR studies too [45, 63]. Loth et al. [37] studied how users perceived the gaze of virtual agents. Bailenson et al. [5] investigated the relationship between gazing and personal space. In this paper, we combine both gait and gaze behaviors to convey emotions in virtual agents.

2.2 Emotions of Virtual Agents

Affective computing has been applied to improve the simulation of many human-like behaviors. Jaques et al. [25] investigated the design of intelligent agents to promote bonding effectively. Liebold et al. [32] discuss how humans understand and process emotions in virtual environments. Lee et al. [31] present an emotion model for virtual agents that deals with the resolution to ambivalence, in which two emotions conflict. Pelczar et al. [48] present a method to evaluate the accuracy of the recognition of expressions of emotions in virtual agents.

Many approaches have been proposed to simulate emotions in virtual agents using verbal communication [11, 58], facial expressions [4, 7, 11, 19, 35, 60], gaze [30], gaits [39], and gestures [47]. McHugh et al. [39] studied the role of dynamic body postures on the perception of emotion in crowded scenes. Clavel et al. [12] evaluated how both face and posture modalities affect the perceptions of emotions in virtual characters. Lance and Marsella [30] used a generative model of expressive gaze for emotional expression in virtual agents. Many generalized virtual agent modeling frameworks have been developed (e.g., SmartBody [59]) which model the animation, navigation, and behavior of virtual agents. Vinayagamoorthy et al. [64] describe models of individual characters' emotions and personalities, models of interpersonal behaviors, and methods for generating emotions using facial expressions. Liebold et al. [33] observe that multimodal (e.g., a combination of verbal cues and facial expression) expressions of emotions yield the highest recognition rates. Inspired by these approaches, we use a perception study (for gait features) and psychological characterization (for gaze features) to simulate emotions of virtual agents.

3 GAIT EMOTION ASSOCIATION METRIC

In this section, we provide an overview of our algorithm for generating emotions of virtual agents. We also describe the gait emotion metric. This metric is used at runtime to generate gaits corresponding to different emotions.

3.1 Overview

Our *EVA* algorithm (Figure 2) consists of two parts: (1) offline computation of an association between gaits and emotions, and (2) real-time generation of virtual agents with desired emotional characteristics using expressive features. We obtain emotion labels

¹Please refer to the accompanying video for visual results.

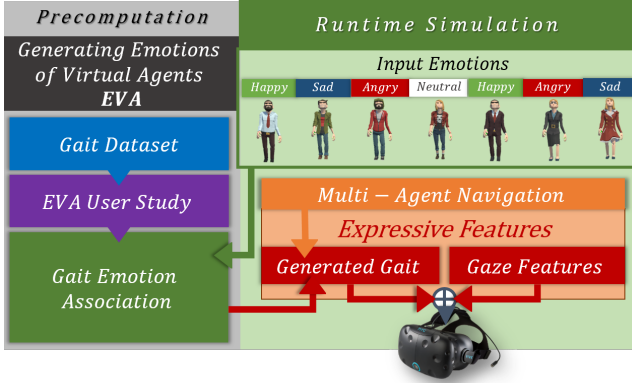


Figure 2: Overview: Our virtual agent generation algorithm consists of two parts: (1) offline computation of gait emotion association and (2) real-time generation of virtual agents (EVA) with desired emotional characteristics using expressive features. We obtain emotion labels for a set of motion-captured gaits with a user study and model the gait emotion association metric using these labels. At runtime, we take desired emotions for each virtual agent as input and generate expressive features corresponding to the desired emotions.

for a set of motion-captured gaits with a user study. We use these labels and formulate a data-driven association between the gaits and emotions. We can also use this association to predict the perceived emotion of any input gaits. At runtime, we take desired emotions for each virtual agent as inputs and use the gait emotion association to generate gaits corresponding to the desired emotions. We also combine our method with a multi-agent navigation algorithm (e.g., RVO [61]) for collision-free navigation of the agents in the virtual environment. We combine gaze features according to the desired emotion for each virtual agent.

3.2 Gait Dataset

We used a combination of publicly available datasets of walking gaits to formulate our gait emotion association metric:

- **BML** [38]: This motion-captured dataset contains 120 different gaits with four different walking styles 30 subjects (15 male and 15 female).
- **CMU** [1]: This motion-captured dataset contains 49 gaits obtained from subjects walking with different styles.
- **EWalk**: This dataset contains gaits extracted from 94 RGB videos using state-of-the-art 3D pose estimation [15].
- **Human3.6M** [24]: This motion-captured dataset contains 14 gaits acquired by recording the performance of five female and six male subjects walking.
- **ICT** [44]: This motion-captured dataset contains 24 gaits obtained from videos of 24 subjects.
- **SIG** [65]: This is a synthetic dataset of 41 gaits generated using local mixtures of autoregressive models.

We visualized each gait using a skeleton mesh (Figure 3). We rendered the gait visualizations from the viewpoint of a camera situated in front of the skeleton mesh and generated a total of 342 gait visualizations.

3.3 EVA User Study

To compute our gait emotion association metric, we showed the 342 gait visualizations to participants in a web-based user study. Based on the participant responses, we obtained emotion labels for each video.



Figure 3: Gait Visualization: We presented 342 gaits visualized from the viewpoint of a camera situated in front of the mesh.

	Happy	Angry	Sad	Neutral
Happy	1.000	-0.268	-0.775	-0.175
Angry	-0.268	1.000	-0.086	-0.058
Sad	-0.775	-0.086	1.000	-0.036
Neutral	-0.175	-0.058	-0.036	1.000

Table 1: Correlation Between Emotion Responses: We present the correlation between participants’ responses to questions relating to the four emotions.

3.3.1 Experiment Details. We recruited 688 participants (279 female, 406 male, $age = 34.8$) from Amazon Mechanical Turk. We presented 10 randomly chosen videos to each participant. After watching each video, the participant answered whether he/she perceived the gait in the video as happy, angry, sad, or neutral on a 5-point Likert scale ranging from *strongly disagree* to *strongly agree*. For each video, we obtained a minimum of 10 participant responses.

3.3.2 Results. We converted the participant responses to an integer scale from [1, 5], where 1 corresponds to *strongly disagree*, and 5 corresponds to *strongly agree*. For each emotion e , we obtained the mean of all participant responses ($r_{i,j}^e$) for each gait G_i in the datasets:

$$r_i^e = \frac{\sum_{j=1}^{n_p} r_{i,j}^e}{n_p}, \quad (1)$$

where j is the participant id and n_p is the number of participant responses collected.

We computed the correlation between participants’ responses to the questions relating to the four emotions (Table 1). A correlation value closer to 1 indicates that the two variables are positively correlated and a correlation value closer to -1 indicates that the two variables are negatively correlated. A correlation value closer to 0 indicates that two variables are uncorrelated. As expected, *happy* and *sad* are negatively correlated and *neutral* is uncorrelated with the other emotions.

3.4 Gait Emotion Association

We used the participant responses to obtain emotion labels for the gaits in the gait datasets. We obtained the emotion label e_i for a gait G_i as $e_i = e \mid r_i^e > \theta$, where $\theta = 3.5$ is an experimentally determined threshold for emotion perception.

If $r_i^e < \theta$ for all 4 emotions, then the corresponding gait G_i is not emotionally expressive and we do not use it for metric computation (21 gaits). If there are multiple emotions with average participant responses greater than $r_i^e > \theta$, we choose the emotion with the highest average participant response as the label e_i for the gait G_i .



Figure 4: Emotion Distribution: We present the distribution of the emotion labels obtained from participant responses. Our data is equitably distributed among the four emotion categories.

Using these gaits G_i and emotion labels e_i , we define a function $GEA(e)$ that returns a set of gaits G_i for which the emotion label e_i matches with the emotion e as $GEA(e) = \{G_i | e_i = e\}$. We refer to this function as our novel gait emotion association metric.

We present the distribution of the gaits and their associated emotions in Figure 4. The smallest emotion category, *Neutral*, contains 16.35% of the gaits, whereas the largest emotion category, *Happy*, contains 32.07%, indicating that our data is equitably distributed among the four emotion categories.

3.5 Application: Prediction of Perceived Emotion

We present an application of our gait emotion association to predict the perceived emotion given any new gait. This allows us to extend our gait emotion association to include any new gait by predicting its perceived emotion value without having to obtain ratings from participants.

We provide an overview of the application framework in Figure 5. We use the gait emotion association to train an SVM. Given a gait, we compute numerical values for its expressive features using psychological characteristics [13, 26] corresponding to posture and movement of joints. We use the SVM to classify our computed gait features into one of the four basic emotions. We extract features corresponding to both posture and movement modalities because both are important for the accurate prediction of an individual's emotional state [27]. We assume a skeleton with 16 joints. Before extracting features, we extract frames corresponding to a single stride from the video, i.e., frames corresponding to consecutive foot strikes of the same foot (walk cycle). We compute the posture and movement features at each frame and aggregate them over the walk cycle.

3.5.1 Posture Features. We use the following posture features:

- **Volume:** According to Crenn et al. [13], body expansion conveys positive emotions, whereas a person has a more compact posture during negative expressions. We model this by a single scalar value corresponding to the volume occupied by the bounding box around the human.
- **Area:** We model body expansion by two scalar values corresponding to the areas of triangles between the hands and the neck and between the feet and the root joint.
- **Distance:** Distances of the feet and the hands from the root joint

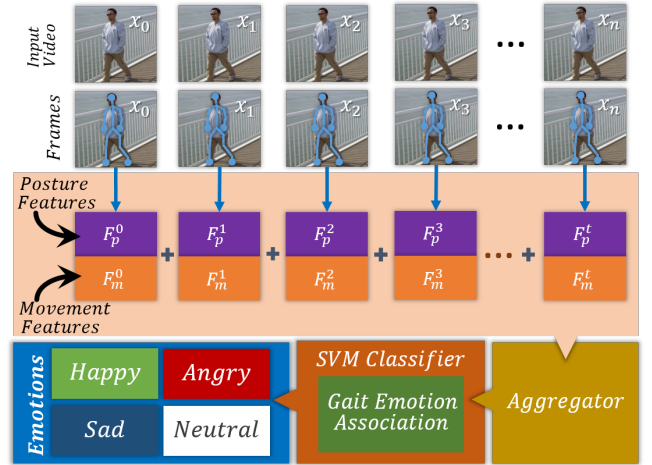


Figure 5: Prediction of Perceived Emotion: We can use gait emotion association to predict perceived emotions of any new gait. Given a new gait, we compute gait features using psychological characterization [13]. We use an SVM classifier trained using the gait emotion association to classify these features into one of the four basic emotions with 70.04% accuracy.

can be used to model body expansion (four scalar values of distances).

- **Angle:** Head tilt is used to distinguish between happy and sad emotions [13, 26]. We model this using five scalar values corresponding to the angles: (1) at the neck by the shoulder joints, (2) at the shoulder joints by the neck and the other shoulder, (3) at the neck by the vertical direction and the back, and (4) at the neck by the head and the back.

We also include stride length as a posture feature. Longer stride lengths convey anger and happiness, and shorter stride lengths convey sadness and neutrality [26]. Suppose we represent the positions of the left foot joint j_{lFoot} and the right foot joint j_{rFoot} in frame t as $\vec{p}(j_{lFoot}, t)$ and $\vec{p}(j_{rFoot}, t)$, respectively. Then the stride length $s \in \mathbb{R}$ is computed as:

$$s = \max_{t \in 1 \dots \tau} \|\vec{p}(j_{lFoot}, t) - \vec{p}(j_{rFoot}, t)\|. \quad (2)$$

We represent the posture features at frame $t = \{1, 2, \dots, \tau\}$ by a 13-dimensional vector $F_{p,t} \in \mathbb{R}^{13}$, where τ is the number of frames in the walk cycle. Then, we can compute the posture features of the gait F_p as the average of the posture features at each frame t combined with the stride length:

$$F_p = \frac{\sum_t F_{p,t}}{\tau} \cup s. \quad (3)$$

3.5.2 Movement Features. High arousal emotions such as anger and happiness are associated with rapid and increased movements, whereas low arousal emotions such as sadness are associated with slow movements [10]. We consider the magnitude of the velocity, acceleration, and movement jerk of the hand, foot, and head joints and compute the movement features $F_{m,t} \in \mathbb{R}^{15}$ at frame t . For each of these five joints $j_i, i = 1, \dots, 5$, we compute the magnitude of the first, second, and third derivatives of the position vector $\vec{p}(j_i, t)$ at frame t .

Since faster gaits are perceived as happy or angry and slower gaits are considered as sad [26], we also include the time taken for one walk cycle ($gt \in \mathbb{R}$) as a movement feature. We define the

Method	Accuracy
Karg et al. [26]	39.58%
Venture et al. [62]	30.83%
Crenn et al. [13]	66.22%
Crenn et al. [14]	40.63%
Daoudi et al. [16]	42.52%
<i>Our Method (PEP)</i>	70.04%

Table 2: Accuracy: Our method using gait features, including posture and movement features, and SVM classifier achieves an accuracy of 70.04% for emotion classification. We obtain considerable improvement over prior methods.

movement features $F_m \in \mathbb{R}^{16}$ as the average of $F_{m,t}$, $t = \{1, 2, \dots, \tau\}$:

$$F_m = \frac{\sum_t F_{m,t}}{\tau} \cup gt. \quad (4)$$

We combine posture and movement features and compute the gait features F as $F = F_m \cup F_p$. In the end, we normalize the feature values to $[-1, 1]$ with -1 and 1 representing the minimum and maximum values of the feature in the dataset, respectively.

3.5.3 Emotion Classification Accuracy. For each gait G_i in the dataset, we have computed the gait features F_i and the associated perceived emotion label e_i . Using this gait emotion association, we train an SVM classifier with an RBF kernel with a one-versus-rest decision function of shape. For any new gait, we compute the values of the gait features and use the trained SVM to predict its perceived emotion. We obtain the accuracy results by performing 10-fold cross-validation on all six datasets (Section 3.2). We compare the accuracy (Table 2) of our classification algorithm with other state-of-the-art methods that predict emotions based on gaits:

- Karg et al. [26]: This feature-based classification method uses PCA to classify gait features related to shoulder, neck, and thorax angles, stride length, and velocity. This method only models the posture features for the joints and doesn't model the movement features.
- Venture et al. [62]: This method uses the auto-correlation matrix of the joint angles at each frame and uses similarity indices for classification. Their algorithm achieves good intra-subject accuracy but performs poorly for inter-subject databases.
- Crenn et al. [13]: This feature-based method uses both posture and movement features and classifies these features using SVMs. This method considers only the upper body and does not model features related to feet joints for classification.
- Daoudi et al. [16]: This method uses a manifold of symmetric positive definite matrices to represent body movement and classifies them using the nearest neighbors method.
- Crenn et al. [14]: This method is based on the synthesis of neutral motion from an input motion and uses the residual between the input and neutral motion for classification.

4 EMOTIONS IN VIRTUAL AGENTS (EVA)

In this section, we provide the details of our EVA algorithm, which uses the gait emotion association to generate emotions in virtual agents.

4.1 Notation

We define a gait G as a set of 3D poses P_1, P_2, \dots, P_τ where τ is the number of frames. The set of 3D positions of each joint J_i is referred to as the pose of an agent $P \in \mathbb{R}^{3*n}$ where n is the number of joints.

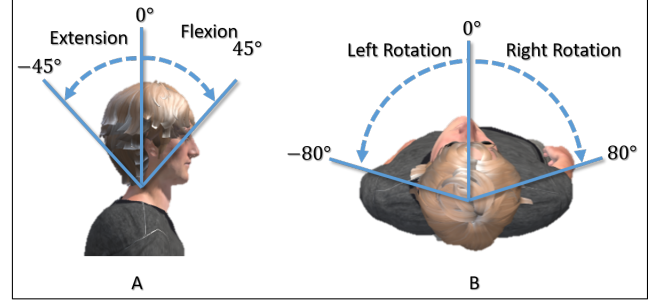


Figure 6: Gaze Control: We control virtual agents' gaze by controlling the rotations of the neck joint in the skeleton mesh. Specifically, we control the two degrees of freedoms associated with (1) neck flexion and extension, and (2) neck left and right rotation.

4.2 Gait Generation

At runtime, for each agent, we take the desired emotion e_{des} as input. We use the gait emotion association metric and obtain a set of gaits $GEA(e_{des})$ that correspond to the desired emotion. We choose one of the gaits G_{des} from $GEA(e_{des})$ and update the joint positions of the agent in the virtual world using the joint positions from G_{des} . The selection of G_{des} from $GEA(e_{des})$ can be made according to many criteria (such as personality or preferred walking speed). In our current implementation, we choose the gait randomly.

4.3 Local Navigation using Expressive Features

For navigation, we represent each agent on the 2D ground plane and compute its collision-free velocity. We use a reciprocal velocity obstacle-based method, RVO, to generate smooth, stable, collision-free velocities [61]. RVO is an agent-based approach that computes a collision-free 2D velocity of an agent given its preferred velocity, time horizon (t_{max}), and current positions and velocities of the all virtual agents in the environment. In other words, it computes a velocity that can generate a collision-free trajectory at time t_{max} .

Based on the chosen gait G_{des} for an agent i , we set the preferred speed of the agent as the speed of the gait G_{des} . The speed of the gait v_{des} is computed as follows:

$$v_{des} = \|p(J_{root}, \tau) - p(J_{root}, 0)\| / \tau, \quad (5)$$

where $p(J_{root}, \tau)$ and $p(J_{root}, 0)$ represent the positions of the root joint in the first P_0 and last pose P_τ , respectively.

4.4 Gaze Behaviors

According to psychology literature, the direction of the gaze affects the perception of emotions [3]. Direct gazes are associated with approach-oriented emotions such as anger and averted gazes are associated with avoidance-oriented emotions such as sadness [3]. Using these observations, we present an emotion-based gaze formulation.

We control a virtual agent's gaze by controlling the rotations of the neck joint in the skeleton mesh. Specifically, we control the angles corresponding to two degrees-of-freedom: (1) neck flexion and extension and (2) neck left and right rotation. Given the desired emotion e_{des} for a virtual agent i , our goal is to compute the two angles θ_{flex} and θ_{rot} . We present a visualization of our gaze control angles in Figure 6. Here, the subscripts i and u represent the virtual agent and the user respectively, and (x, y, z) represent the 3D coordinates.

• **Happy:** According to prior literature, direct gaze and neck extension are associated with happiness [3]. Therefore, for $e_{des} = happy$, the agent's neck orientation is computed such that the virtual agent gazes directly at the user with an upward head tilt. The neck angles corresponding to $e_{des} = happy$ are computed as follows:

$$\theta_{flex} = \theta_{happy}, \quad (6)$$

$$\theta_{rot} = \text{angle}(\vec{p}_u - \vec{p}_i, \vec{f}), \quad (7)$$

where θ_{happy} is a negative constant value indicating neck extension and $\text{angle}(*, *)$ represents the angle between two vectors. We use $\theta_{happy} = -5^\circ$ for our validation study.

• **Angry:** Direct gaze is associated with approach-based emotions such as anger [3]. Therefore, for $e_{des} = angry$, the agent's neck orientation is computed such that the agent gazes directly at the user agent. The neck angles for $e_{des} = angry$ are computed as follows:

$$\theta_{flex} = \arcsin\left(\frac{y_i - y_u}{\sqrt{(x_i - x_u)^2 + (z_i - z_u)^2}}\right), \quad (8)$$

$$\theta_{rot} = \text{angle}(\vec{p}_u - \vec{p}_i, \vec{f}), \quad (9)$$

where $\text{angle}(*, *)$ represents the angle between two vectors.

• **Sad:** Averted gaze and neck flexion are associated with sadness [3]. The neck angles for $e_{des} = sad$ are computed as follows:

$$\theta_{flex} = \theta_{sad}, \quad (10)$$

$$\theta_{rot} = 0^\circ, \quad (11)$$

where θ_{sad} is a negative constant value indicating neck extension. We use $\theta_{sad} = 10^\circ$ for our validation study.

• **Neutral:** For the neutral emotion, we use the orientations of the neck angles θ_{flex} and θ_{rot} computed from the positions of the head and neck joints in the gait G_{des} .

4.5 EVA Algorithm

During runtime, we first compute desired gait G_{des} using gait generation method described in Section 4.2. Next, we compute collision-free trajectories and update the position of the agents in the ground plane. We retarget the generated gait G_{des} to the avatar of the virtual agent using Unity's skeleton retargeting [2] and update the positions of the agent's joints in the virtual world using the joint positions from G_{des} . We use the emotion-based gazing behaviors to update the neck joint's flexion and rotation angles. The resultant virtual agents are then rendered to an HMD.

5 USER EVALUATION

In this section, we present the details of our user evaluation conducted to evaluate the benefits of our virtual agent movement generation algorithm (EVA) with and without gazing behaviors. We compare its performance with a baseline method.

5.1 Experiment Goals and Expectations

We performed the user evaluation to determine whether or not significant improvement in the sense of presence can be observed by a real user immersed in a virtual environment based on the emotional state of virtual agents computed using EVA. We chose the sense of presence metric because it has been widely used to evaluate the quality of a VR experience, and a higher sense of presence is desirable for most VR applications [21, 45]. We wanted to demonstrate that our algorithm improves the sense of presence, thus improving the overall quality of the simulations. We compared the following three virtual agent generation algorithms:

• **Baseline:** All virtual agents were simulated with a single neutral gait. Virtual agents did not perform any gazing behaviors.

• **EVA-O:** Our novel virtual agent generation algorithm (EVA) without any gazing behaviors. We generate gaits corresponding to the four emotions: *happy*, *angry*, *sad*, or *neutral*.

• **EVA-G:** Our novel virtual agent generation algorithm (EVA) with emotion-specific gaits along with the gazing behaviors.

Additionally, we also performed an emotion identification task, where we wanted to estimate the accuracy of our EVA algorithm in generating characters with the desired perceived emotions.

5.2 Experimental Design

We conducted the study using a within-subjects design. We showed each participant four scenarios with virtual agents walking in different environments. For each scenario, participants performed three trials corresponding to the three different virtual agent generation algorithms. The order of the scenarios and the trials was counterbalanced. In each of the scenarios, we used a set of male and female virtual agents. For EVA-O and EVA-G, the virtual agents were simulated with desired emotions chosen randomly out of the four categories: *Happy*, *Angry*, *Sad*, and *Neutral*. For the baseline, all the agents were simulated using the gaits corresponding to the *Neutral* category.

After the four scenarios were completed, the participants performed the emotion identification task. For this task, we generated eight virtual agents: two agents corresponding to each desired emotion from *Happy*, *Angry*, *Sad*, and *Neutral*. We assigned male and female virtual agents to each emotion in this task as well. We displayed these virtual agents one-by-one to the participants. Participants were asked to classify these agents as *Happy*, *Angry*, *Sad*, and *Neutral*.

5.2.1 Procedure. We welcomed the participants and informed them about the overall process and purpose of the study. We introduced them to an HTC Vive HMD and informed that it might cause nausea and slight discomfort. The experiment was approved by Institutional Review Boards and the Office of Human Research Ethics. Before beginning the experiment, the participants were invited to read and agree to the approved consent form. They were also informed that they could opt out of the experiment at any time. During the study, the participants could walk in the tracking space and look around in the virtual environment by rotating their heads. Participants provided optional demographic information about their gender and age. The study required approximately 20 minutes, and participants were compensated with a gift card worth \$5.

5.2.2 Scenarios. We evaluated the virtual agent generation algorithms in 4 different scenarios (Figure 7):

- **Crossing:** This scenario had 8 virtual agents crossing a busy intersection in a city environment.
- **Garden:** This scenario had 4 virtual agents walking in a small leisure garden.
- **Parking:** This scenario consisted of 8 virtual agents walking in a parking lot near a set of stores.
- **Scifi:** This scenario consisted of 4 virtual agents walking inside a fantastical residential pod.

In each of the scenario, we chose characters with low polygon counts because they provide the capability to simulate a large number of virtual agents while maintaining the necessary frame rate for VR. It also avoids the uncanny valley problem [23].

5.2.3 Participants. A total of 30 participants (20 male, 10 female, $\bar{x}_{age} = 25.4$) from a university took part in the study.



Figure 7: Scenarios: Our user evaluation consists of these 4 scenarios.

5.2.4 Questions. To test the proposed hypotheses, we used a modified version of a well-established questionnaire by Garau et al [21]. These questions attempt to assess the various aspects of social presence [36]. In addition to a subset of the original questions, we asked a question whether the virtual agents seemed to be experiencing different emotions. Participants answered the Agree/Disagree type questions on seven-level Likert items. We ask the following questions:

- Q1, Spatial Presence:** I had a sense of being in the same space as the characters.
- Q2, Awareness:** The characters seemed to be aware of me.
- Q3, Interaction:** I felt that I should talk/nod/respond to the characters.
- Q4, Realism:** The characters seemed to resemble real people.
- Q5, Emotions:** The characters seemed to be experiencing different emotions.

5.3 Results

In this section, we discuss the participants' responses. We computed average participant responses for all scenarios (Figure 8). We also tested the differences between responses for the three algorithms using the *Friedman test*. For this test, the simulation algorithm is the independent variable, and the participant response is the dependent variable. We present the test statistic value χ^2 and the significance level p in Table 3. We observe significant differences between the three compared methods for all questions across all scenarios (except Q1 in Scifi and Garden scenario). We also performed a post hoc analysis using the *Wilcoxon Signed-Rank tests*. We applied Bonferroni correction resulting in a significance level at $p < 0.017$. We present the significance level (p) for pairwise comparisons between the three algorithms in Table 5. We observe significant differences between baseline and EVA-G, indicating that our method of generating emotions in virtual agents using both gait and gaze features performs consistently better than the baseline method. We discuss the results in detail below:

- **Spatial Presence:** According to the higher averaged responses, participants experienced a higher sense of being in the same space as the characters simulated with EVA, as compared to the baseline. For Garden and Scifi scenarios, we do not observe a significant difference between the sense of presence felt by the participants

	Crossing		Garden		Parking		Scifi	
	χ^2	p	χ^2	p	χ^2	p	χ^2	p
Q1	9.75	0.01	4.67	0.10	15.13	0.00	2.46	0.29
Q2	19.58	0.00	26.45	0.00	30.28	0.00	27.30	0.00
Q3	20.26	0.00	27.27	0.00	20.17	0.00	24.90	0.00
Q4	13.77	0.00	12.30	0.00	24.45	0.00	10.53	0.01
Q5	23.36	0.00	41.32	0.00	35.04	0.00	43.38	0.00

Table 3: Results of the Friedman Test: We present the test statistic (χ^2) and the significance level (p) for each scenario for all questions. The values of $p < 0.05$ indicate a significant difference between the responses for the three algorithms.

		Participant Response			
		Happy	Angry	Sad	Neutral
Desired Emotion	Happy	60.41%	4.17 %	4.17 %	31.25%
	Angry	18.75%	52.08%	0.00 %	29.17%
	Sad	0.00 %	0.00 %	100.00%	0.00 %
	Neutral	8.33 %	4.17 %	16.67 %	70.83%

Table 4: Confusion Matrix: We present the confusion matrix for the emotion identification task. Here, rows indicate the desired emotion input provided to the EVA algorithm, and the columns indicate the participants' responses.

across the three algorithms according to the Friedman test. Because of the insignificant differences in the Friedman test, post hoc tests do not apply to these scenarios. For Crossing and Parking scenarios, where there were more agents, we observe a significant difference between the responses for different methods in the Friedman test. The Wilcoxon tests reveal significant differences between the *baseline* and *EVA-G* methods.

- **Awareness:** Participants felt that the virtual agents were more aware of them compared to the baseline method as evidenced by the averaged responses. We observe significant differences between the responses across all scenarios using the Friedman test. Wilcoxon tests reveal significant differences in the *baseline* vs. *EVA-G* and *EVA-O* vs. *EVA-G* comparisons, highlighting the importance of gazing behaviors in our algorithm.

- **Interaction:** Participants felt that they should talk/nod/respond to the agents if their movement was simulated by EVA as compared to the baseline as evidenced by the averaged responses. We observe

	Crossing			Garden			Parking			Scifi		
	Baseline	EVA-O	Baseline	Baseline	EVA-O	Baseline	Baseline	EVA-O	Baseline	Baseline	EVA-O	Baseline
	vs EVA-O	vs EVA-G	vs EVA-G	vs EVA-O	vs EVA-G	vs EVA-G	vs EVA-O	vs EVA-G	vs EVA-G	vs EVA-O	vs EVA-G	vs EVA-G
Q1	0.018	0.323	0.005	NA	NA	NA	0.179	0.019	0.003	NA	NA	NA
Q2	0.037	0.002	0.000	0.095	0.000	0.000	0.540	0.000	0.000	0.179	0.000	0.000
Q3	0.012	0.005	0.000	0.007	0.000	0.000	0.256	0.001	0.000	0.153	0.000	0.000
Q4	0.038	0.276	0.001	0.026	0.160	0.012	0.005	0.035	0.000	0.054	0.041	0.010
Q5	0.000	0.011	0.000	0.000	0.070	0.000	0.001	0.000	0.000	0.000	0.000	0.000

Table 5: Results of the Wilcoxon Signed-Rank Test: We present the significance level (p) for pairwise comparisons between the three algorithms using the post hoc tests. The values of $p < 0.017$ indicate a significant difference between the responses for corresponding pairwise comparisons.

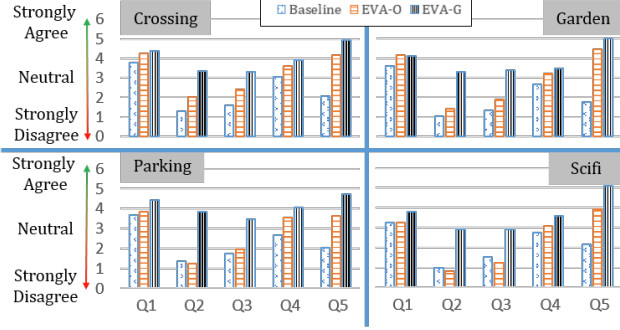


Figure 8: Averaged Responses: We present the average participant responses for all scenarios. These results indicate that the participants consistently preferred our EVA algorithm over the baseline method.

significant differences between the responses across all scenarios using the Friedman test. Wilcoxon tests reveal significant differences for the *baseline* vs. *EVA-G* and *EVA-O* vs. *EVA-G* comparisons highlighting the importance of gazing behaviors in our algorithm. We also observe a significant difference for the *baseline* vs. *EVA-O* comparisons for the Crossing and Garden scenarios, indicating the importance of gaits in eliciting more response from the users.

- **Realism:** When EVA simulated the agents, they seemed to resemble real people more than when they were simulated by the baseline algorithm as indicated by the participant responses. We observe significant differences between the responses across all scenarios using the Friedman test. Wilcoxon tests reveal significant differences in the *baseline* vs. *EVA-G*, highlighting the importance of both gait and gazing behaviors in our algorithm.

- **Emotions:** The agents seemed more likely to be experiencing different emotions if the EVA algorithm simulated them compared to the baseline algorithm according to the averaged responses. We observe significant differences between the responses across all scenarios using the Friedman test. Wilcoxon tests reveal significant differences for all the pairwise comparisons (except for *EVA-G* vs. *EVA* comparison in the Garden scenario).

5.4 Benefits of Expressive Features

These results show that expressive features provide the following benefits:

- for scenarios with a higher number of virtual agents, virtual agents simulated using expressive features improve the sense of presence felt by the user,
- expressive features make the virtual agents appear more aware of the user,
- elicit more response from the users, and

- increase the resemblance of the virtual agents to real people.

5.5 Emotion Identification

For the emotion identification task, the participants identified the emotions of the virtual agents correctly with an accuracy of 70.83%. We present the confusion matrix in Table 4. We obtain 100% accuracy in simulating virtual agents that are perceived as sad and more than 50% accuracy for all the emotions. These results indicate that expressive features provide benefits for the perception of emotions in virtual agents.

6 CONCLUSIONS, LIMITATIONS, AND FUTURE WORK

We present a data-driven approach for generating emotions in virtual agents (*EVA*) using expressive features. We use a novel data-driven metric between gaits and emotions. We combine these gaits with gazing behaviors and generate virtual agents that are perceived as experiencing *happy*, *angry*, *sad*, or *neutral* emotions. We also present an application of our gait emotion association to predict how the real users will perceive the emotions of virtual agents simulated with any new gait and achieve an accuracy of 70.04%. We validated our EVA algorithm with user evaluation. Results of this evaluation indicate that the expressive features can increase the sense of presence felt by the user in the virtual environment and make the virtual agents elicit more response from them.

Our approach has some limitations. Our approach only takes into account expressive features of gait and gaze, and ignore the impact of other agent characteristics related to appearance or facial expressions or other non-verbal cues like appearance, voice, pupil dilation, etc. for the perception of emotion [29]. We conducted the evaluation in a multi-agent VR environment to measure the contribution of our algorithm in a more interactive and immersive setting. In the future, we would like to examine the gait emotion association metric per agent to strengthen the desired emotion as much as possible while dampening the others. Furthermore, we limited this approach to walking activity. In future work, we would like to extend our approach to other activities. Since previous research shows that gender is recognizable from gaits [67], we would like to consider the gender of the virtual agents while generating a gait for them. We would also like to combine our approach with other emotional cues and compare it with different algorithms that simulate other types of motion styles (e.g., personalities [17]). Finally, we would like to examine the impact of the context and culture of the observer on the perception of emotion from bodily expressions using recent psychological findings [22].

REFERENCES

- [1] 2018. CMU Graphics Lab Motion Capture Database. <http://mocap.cs.cmu.edu/>.

- [2] 2019. Retargeting of Humanoid animations. <https://docs.unity3d.com/Manual/Retargeting>. Accessed: 2019-06-30.
- [3] Reginald B Adams and Robert E Kleck. 2003. Perceived gaze direction and the processing of facial displays of emotion. *Psychological Science* 14, 6 (2003), 644–647.
- [4] Norman Badler, Jan Allbeck, Liwei Zhao, and Meeran Byun. 2002. Representing and parameterizing agent behaviors. In *Computer Animation, 2002. Proceedings of IEEE*, 133–143.
- [5] Jeremy N Bailenson, Andrew C Beall, Jack Loomis, Jim Blascovich, and Matthew Turk. 2005. Transformed social interaction, augmented gaze, and social influence in immersive virtual environments. *Human communication research* 31, 4 (2005), 511–537.
- [6] L. F. Barrett. 2017. *How emotions are made: The secret life of the brain*. Houghton Mifflin Harcourt.
- [7] Jenay M Beer, Arthur D Fisk, and Wendy A Rogers. 2009. Emotion recognition of virtual agents facial expressions: the effects of age and emotion intensity. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 53. SAGE Publications Sage CA: Los Angeles, CA, 131–135.
- [8] Aniket Bera, Tanmay Randhavane, Emily Kubin, Husam Shaik, Kurt Gray, and Dinesh Manocha. 2018. Data-driven modeling of group entitativity in virtual environments. In *Proceedings of the 24th ACM Symposium on Virtual Reality Software and Technology*. ACM, 31.
- [9] Aniket Bera, Tanmay Randhavane, and Dinesh Manocha. 2017. Aggressive, Tense or Shy? Identifying Personality Traits from Crowd Videos. In *IJCAL*. 112–118.
- [10] D Bernhardt and P Robinson. 2007. Detecting affect from non-stylised body motions. In *ACII*.
- [11] Andry Chowanda, Peter Blanchfield, Martin Flintham, and Michel Valstar. 2016. Computational Models of Emotion, Personality, and Social Relationships for Interactions in Games: (Extended Abstract). In *AAMAS*.
- [12] Céline Clavel, Justine Plessier, Jean-Claude Martin, Laurent Ach, and Benoit Morel. 2009. Combining facial and postural expressions of emotions in a virtual character. In *International Workshop on Intelligent Virtual Agents*. Springer, 287–300.
- [13] Arthur Crenn, Rizwan Ahmed Khan, Alexandre Meyer, and Saida Bouakaz. 2016. Body expression recognition from animated 3D skeleton. In *2016 International Conference on 3D Imaging (IC3D)*. IEEE, 1–7.
- [14] Arthur Crenn, Alexandre Meyer, Rizwan Ahmed Khan, Hubert Konik, and Saïda Bouakaz. 2017. Toward an efficient body expression recognition based on the synthesis of a neutral movement. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*. ACM, 15–22.
- [15] Rishabh Dabral, Anurag Mundhada, Uday Kusupati, Safeer Afaque, Abhishek Sharma, and Arjun Jain. 2018. Learning 3d human pose from structure and motion. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 668–683.
- [16] Mohamed Daoudi, Stefano Berretti, Pietro Pala, Yvonne Delevoye, and Alberto Del Bimbo. 2017. Emotion recognition by body movement representation on the manifold of symmetric positive definite matrices. In *International Conference on Image Analysis and Processing*. Springer, 550–560.
- [17] Funda Durupinar, Mubbasi Kapadia, Susan Deutsch, Michael Neff, and Norman I Badler. 2017. Perform: Perceptual approach for adding ocean personality to human motion using laban movement analysis. *ACM Transactions on Graphics (TOG)* 36, 1 (2017), 6.
- [18] Paul Ekman and Wallace V Friesen. 1967. Head and body cues in the judgment of emotion: A reformulation. *Perceptual and motor skills* 24, 3 PT 1 (1967), 711–724.
- [19] Ylva Ferstl and Rachel McDonnell. 2018. A perceptual study on the manipulation of facial features for trait portrayal in virtual agents. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents*. ACM, 281–288.
- [20] Andrew C Gallup, Andrew Chong, Alex Kacelnik, John R Krebs, and Iain D Couzin. 2014. The influence of emotional facial expressions on gaze-following in grouped and solitary pedestrians. *Scientific reports* 4 (2014), 5794.
- [21] Maia Garau, Mel Slater, David-Paul Pertaub, and Sharif Razzaque. 2005. The responses of people to virtual humans in an immersive virtual environment. *Presence: Teleoperators & Virtual Environments* 14, 1 (2005), 104–116.
- [22] Beatrice de Gelder. 2016. *Emotions and the body*. Oxford University Press.
- [23] Tom Geller. 2008. Overcoming the uncanny valley. *IEEE computer graphics and applications* 28, 4 (2008), 11–17.
- [24] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. 2013. Human3. 6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence* 36, 7 (2013), 1325–1339.
- [25] Natasha Jaques, Daniel McDuff, Yoo Lim Kim, and Rosalind Picard. 2016. Understanding and predicting bonding in conversations using thin slices of facial expressions and body language. In *International Conference on Intelligent Virtual Agents*. Springer, 64–74.
- [26] Michelle Karg, Kolja Kuhnlenz, and Martin Buss. 2010. Recognition of affect based on gait patterns. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 40, 4 (2010), 1050–1061.
- [27] A. Kleinsmith and N. Bianchi-Berthouze. 2013. Affective Body Expression Perception and Recognition: A Survey. *IEEE Transactions on Affective Computing* 4, 1 (Jan 2013), 15–33. <https://doi.org/10.1109/T-AFCC.2012.16>
- [28] Andrea Kleinsmith, P Ravindra De Silva, and Nadia Bianchi-Berthouze. 2005. Grounding affective dimensions into posture features. In *International Conference on Affective Computing and Intelligent Interaction*. Springer, 263–270.
- [29] Mariska Esther Kret, Karin Roelofs, Jeroen Stekelenburg, and Beatrice de Gelder. 2013. Emotional signals from faces, bodies and scenes influence observers’ face expressions, fixations and pupil-size. *Frontiers in human neuroscience* 7 (2013), 810.
- [30] Brent J Lance and Stacy C Marsella. 2008. A model of gaze for the purpose of emotional expression in virtual embodied agents. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 199–206.
- [31] Benny Ping-Han Lee, Edward Chao-Chun Kao, and Von-Wun Soo. 2006. Feeling ambivalent: A model of mixed emotions for virtual agents. In *International Workshop on Intelligent Virtual Agents*. Springer, 329–342.
- [32] Benny Liebold. 2016. *Cognitive and Emotional Processing of Virtual Environments: the Significance of Attentional Processes and Mental Models*. Ph.D. Dissertation. Technische Universität Chemnitz.
- [33] Benny Liebold and Peter Ohler. 2013. Multimodal emotion expressions of virtual agents, mimic and vocal emotion expressions and their effects on emotion recognition. In *Affective Computing and Intelligent Interaction (ACII), 2013 Humaine Association Conference on*. IEEE, 405–410.
- [34] Zhen Liu, Tingting Liu, Minhua Ma, Hui-Huang Hsu, Zhongrui Ni, and Yanjie Chai. 2018. A perception-based emotion contagion model in crowd emergent evacuation simulation. *Computer Animation and Virtual Worlds* 29, 3-4 (2018), e1817.
- [35] Zhen Liu and Zhi Geng Pan. 2005. An emotion model of 3d virtual characters in intelligent virtual environment. In *International Conference on Affective Computing and Intelligent Interaction*. Springer, 629–636.
- [36] Matthew Lombard, Theresa B Ditton, and Lisa Weinstein. 2009. Measuring presence: the temple presence inventory. In *Proceedings of the 12th Annual International Workshop on Presence*. 1–15.
- [37] Sebastian Loth, Gernot Horstmann, Corinna Osterbrink, and Stefan Kopp. 2018. Accuracy of Perceiving Precisely Gazing Virtual Agents. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents*. ACM, 263–268.
- [38] Yingliang Ma, Helena M Paterson, and Frank E Pollick. 2006. A motion capture library for the study of identity, gender, and emotion perception from biological motion. *Behavior research methods* 38, 1 (2006), 134–141.
- [39] Joanna Edel McHugh, Rachel McDonnell, Carol O’Á’ Sullivan, and Fiona N Newell. 2010. Perceiving emotion in crowds: the role of dynamic body postures on the perception of emotion in crowded scenes. *Experimental brain research* 204, 3 (2010), 361–372.
- [40] Albert Mehrabian. 1980. Basic dimensions for a general psychological theory implications for personality, social, environmental, and developmental studies. (1980).
- [41] Albert Mehrabian. 1996. Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in temperament. *Current Psychology* 14, 4 (1996), 261–292.
- [42] Joseph A Mikels, Barbara L Fredrickson, Gregory R Larkin, Casey M Lindberg, Sam J Maglio, and Patricia A Reuter-Lorenz. 2005. Emotional category data on images from the International Affective Picture System. *Behavior research methods* 37, 4 (2005), 626–630.
- [43] Jon D Morris. 1995. Observations: SAM: the Self-Assessment Manikin; an efficient cross-cultural measurement of emotional response. *Journal of advertising research* 35, 6 (1995), 63–68.
- [44] Sahil Narang, Andrew Best, Andrew Feng, Sin-hwa Kang, Dinesh Manocha, and Ari Shapiro. 2017. Motion recognition of self and others on realistic 3D avatars. *Computer Animation and Virtual Worlds* 28, 3-4 (2017), e1762.
- [45] Sahil Narang, Andrew Best, Tanmay Randhavane, Ari Shapiro, and Dinesh Manocha. 2016. PedVr: Simulating gaze-based interactions between a real user and virtual crowds. In *Proceedings of the 22nd ACM conference on virtual reality software and technology*. ACM, 91–100.
- [46] R. S. Nickerson. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology* 2, 2 (1998), 175.
- [47] C Pelachaud. 2009. Studies on gesture expressivity for a virtual agent. *Speech Communication* (2009).
- [48] Ildiko Pelczar, Francisco Cabiedes, and Fernando Gamboa. 2007. Expressions of Emotions in Virtual Agents: Empirical Evaluation. In *Virtual Environments, Human-Computer Interfaces and Measurement Systems, 2007. VECIMS 2007. IEEE Symposium on*. IEEE, 31–35.
- [49] Christopher Peters, Catherine Pelachaud, Elisabetta Bevacqua, Maurizio Mancini, and Isabella Poggi. 2005. A model of attention and interest using gaze behavior. In *International Workshop on Intelligent Virtual Agents*. Springer, 229–240.
- [50] Tanmay Randhavane, Aniket Bera, Kyra Kapsaskis, Kurt Gray, and Dinesh Manocha. 2019. FVA: Modeling Perceived Friendliness of Virtual Agents Using Movement Characteristics. *2019 International Symposium on Mixed and Augmented Reality Special Issue of TVCG* (2019).
- [51] Tanmay Randhavane, Aniket Bera, Emily Kubin, Kurt Gray, and Dinesh Manocha. 2019. Modeling Data-Driven Dominance Traits for Virtual Characters using Gait Analysis. *arXiv preprint arXiv:1901.02037* (2019).
- [52] Tanmay Randhavane, Aniket Bera, Emily Kubin, Austin Wang, Kurt Gray, and Dinesh Manocha. 2019. Pedestrian Dominance Modeling for Socially-Aware Robot Navigation. *2019 IEEE International Conference on Robotics and Automation (ICRA)* (2019).
- [53] Tanmay Randhavane, Aniket Bera, and Dinesh Manocha. 2017. F2FCrowds: Planning agent movements to enable face-to-face interactions. *Presence: Teleoperators*

- and *Virtual Environments* 26, 2 (2017), 228–246.
- [54] Laurel D Riek, Tal-Chen Rabinowitch, Bhismadev Chakrabarti, and Peter Robinson. 2009. How anthropomorphism affects empathy toward robots. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*. ACM, 245–246.
 - [55] H. Riggio. 2017. Emotional Expressiveness. *Encyclopedia of Personality and Individual Differences* (2017).
 - [56] Jesús J Rivas, Felipe Orihuela-Espina, L Enrique Sucar, Lorena Palafox, Jorge Hernández-Franco, and Nadia Bianchi-Berthouze. 2015. Detecting affective states in virtual rehabilitation. In *Proceedings of the 9th International Conference on Pervasive Computing Technologies for Healthcare*. ICST (Institute for Computer Sciences, Social-Informatics and \ddot{A} Ä, 287–292.
 - [57] J Russell. 1987. The identification of emotions from gait information. *Journal of Nonverbal Beh.* (1987).
 - [58] Samuel S Sohn, Xun Zhang, Fernando Geraci, and Mubbasir Kapadia. 2018. An Emotionally Aware Embodied Conversational Agent. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2250–2252.
 - [59] Marcus Thiebaux, Stacy Marsella, Andrew N Marshall, and Marcelo Kallmann. 2008. Smartbody: Behavior realization for embodied conversational agents. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 151–158.
 - [60] Angela Tinwell, Mark Grimshaw, Debbie Abdel Nabi, and Andrew Williams. 2011. Facial expression of emotion and perception of the Uncanny Valley in virtual characters. *Computers in Human Behavior* 27, 2 (2011), 741–749.
 - [61] Jur Van den Berg, Ming Lin, and Dinesh Manocha. 2008. Reciprocal velocity obstacles for real-time multi-agent navigation. In *2008 IEEE International Conference on Robotics and Automation*. IEEE, 1928–1935.
 - [62] Gentiane Venture, Hideki Kadone, Tianxiang Zhang, Julie Grèzes, Alain Berthoz, and Halim Hicheur. 2014. Recognizing emotions conveyed by human gait. *International Journal of Social Robotics* 6, 4 (2014), 621–632.
 - [63] Maia Garau Anthony Steed Vinayagamoorthy, Vinoba and Mel Slater. 2004. An eye gaze model for dyadic interaction in an immersive virtual environment: Practice and experience. In *Computer Graphics Forum*.
 - [64] V. Vinayagamoorthy, M. Gillies, A. Steed, E. Tanguy, X. Pan, C. Loscos, and M. Slater. 2006. Building Expression into Virtual Characters. In *Eurographics 2006 - State of the Art Reports*, Brian Wyvill and Alexander Wilkie (Eds.). The Eurographics Association. <https://doi.org/10.2312/egst.20061052>
 - [65] Shihong Xia, Congyi Wang, Jinxiang Chai, and Jessica Hodgins. 2015. Realtime Style Transfer for Unlabeled Heterogeneous Human Motion. *ACM Trans. Graph.* 34, 4, Article 119 (July 2015), 10 pages. <https://doi.org/10.1145/2766999>
 - [66] Heath Yates, Brent Chamberlain, Greg Norman, and William H. Hsu. 2017. Arousal Detection for Biometric Data in Built Environments using Machine Learning. In *Proceedings of IJCAI 2017 Workshop on Artificial Intelligence in Affective Computing (Proceedings of Machine Learning Research)*, Neil Lawrence and Mark Reid (Eds.), Vol. 66. PMLR, 58–72. <http://proceedings.mlr.press/v66/yates17a.html>
 - [67] Katja Zibrek, Ludovic Hoyet, Kerstin Ruhland, and Rachel McDonnell. 2015. Exploring the effect of motion type and emotions on the perception of gender in virtual humans. *ACM Transactions on Applied Perception (TAP)* 12, 3 (2015), 11.