

# Synthesizing affective virtual reality multicharacter experiences

Angshuman Mazumdar | Christos Mousas 

Department of Computer Graphics  
Technology, Purdue University, West  
Lafayette, Indiana, USA

## Correspondence

Christos Mousas, Department of  
Computer Graphics Technology, Purdue  
University, West Lafayette, IN 47907, USA.  
Email: cmousas@purdue.edu

## Abstract

This article presents a methodology for automatically synthesizing a virtual population (pedestrians placed in a virtual environment) that impacts a user with a specified affective experience. The pipeline began by developing a dataset of behaviors that could be assigned to virtual characters. Next, an annotation phase assigned affective responses of participants to each character's behavior. The design considerations of our affective multicharacter virtual reality experience were then encoded to cost terms and assigned to a total cost function. This method allowed the developer to control the priority and the targets of the cost terms, and given the user inputs, our application could optimize the multicharacter experience using a Markov chain Monte Carlo method known as simulated annealing. A user study was conducted to investigate whether our method could synthesize virtual reality multicharacter experiences that affect participants in an expected way. The results of our study showed that the three different synthesized multicharacter experiences (low, medium, and high negative affect) were perceived as expected by participants; therefore, we argue that we can indeed automatically synthesize virtual reality multicharacter experiences that impact participants' affect levels in an expected way. Limitations and future research directions are discussed.

## KEYWORDS

affect, optimization, virtual crowds, virtual population, virtual reality

## 1 | INTRODUCTION

Human crowds form an important aspect of our daily lives. We are surrounded by other people, and we interact both directly and indirectly (e.g., observe other humans, avoid humans when walking, coordinate our movement, and socially interact) with those surrounding us. Like the experiences we have in the real world, humans commonly interact or cohabit with virtual characters in virtual environments. In such an experience, virtual reality users could either have an active or passive interaction (i.e., either directly interact or just observe the virtual characters).<sup>1</sup> Disregarding the interaction form (active or passive) between humans and virtual characters, prior research has shown that virtual characters can elicit a range of behavioral responses to virtual reality users, ranging from positive to negative.<sup>2-4</sup>

Some studies have addressed human interaction with virtual characters and virtual crowds. Examples include an exploration of the effects of appearance and motion of virtual characters on humans' emotional reactivity,<sup>5</sup> an investigation of the effect that the virtual characters' appearance fidelity has on emotion contagion during interpersonal

interactions,<sup>6</sup> research into the effects of interacting with a crowd of emotional virtual humans,<sup>7</sup> the perception of the interaction with virtual humans through haptic feedback patterns,<sup>8</sup> and more. In such studies, the experiment team developed variations of the stimulus to understand how changes (e.g., small changes in facial expressions or body gestures, changes in appearance of the virtual characters) could affect the behavioral responses of participants. However, the stimulus was prepared manually instead of automatically, which means that the designers needed to put in time and effort.

This article investigates the potential of inverting the stimulus-preparation process for automatically creating an affective multicharacter virtual reality experience. We think that such a method could be used in virtual reality experiences or games, in which developers need to induce a certain level of affect to users, without the need for time-consuming and manual processes. Our method was inspired by previous research on automatically synthesizing virtual reality experiences that fulfill certain design goals;<sup>9–11</sup> therefore, we developed a pipeline that helped us synthesize affective multicharacter experiences automatically. Our pipeline started by developing a dataset of behaviors that could be assigned to the virtual characters of our virtual reality scene. Then, each character's behavior was annotated by participants using the negative affect items retrieved from the positive and negative affect schedule (PANAS) scale.<sup>12</sup> After the annotation phase, cost terms were developed that encoded the design decisions of our multicharacter experience and assigned them to a total cost function. To automatically synthesize an affective experience based on the behaviors that were assigned to the virtual characters, we employed a Markov chain Monte Carlo optimization method<sup>13</sup> called simulated annealing.<sup>14</sup>

Based on the developed methodology, our goal was to investigate whether we could synthesize multicharacter experiences that impacted participants' responses in an expected way. For this reason, a user study was conducted. In this study, a simple virtual reality scenario was developed, and the behavior of characters in the virtual reality scene was optimized to provide low, medium, and high negative affect experiences. We did so because prior studies indicated that users appreciate experiencing negative emotions.<sup>15</sup> Our results showed that significant differences were noted by the participants across the three affective experiences, indicating that the proposed pipeline could be used to automatically synthesize affective multicharacter virtual reality experiences.

The remainder of this article is arranged as follows. Section 2 introduces the related work. Section 3 describes the preliminary steps that were followed to create the annotated dataset. The problem formulation and the optimization process are described in Section 4. The conducted experimental study, the results, and the discussion are presented in Section 5. Finally, the conclusions and future research directions are discussed in Section 6.

## 2 | RELATED WORK

Most real-world interactions with populations are based on the concept of interpersonal communication.<sup>16</sup> The term *interaction* can refer to anything from a user engaging actively with the characters (pushing them away, talking to them, and so forth) or a passive interaction, where the presence of the crowd affects the experience of the user (such as producing a specific emotion in the player's mind). The individual relationships between the members of a virtual population encapsulate many nonverbal communication cues, including eye gaze and proxemics.<sup>17</sup> For example, Hall et al.<sup>18</sup> talked about the importance of personal space and how it is deeply ingrained in our understanding of patterns in the physical world. The five-factor model of personality, more commonly referred to as the OCEAN model,<sup>19</sup> provides insight into how a person can evaluate another person in a collective, using the traits of "openness, conscientiousness, extroversion, agreeableness and neuroticism."<sup>20</sup>

Previous studies have investigated the impact of various characteristics that could be assigned to virtual populations. It has been reported that the size and the number of virtual characters assigned to a virtual crowd positively impact the realism of the crowd,<sup>21–23</sup> and the interaction scenario and the behaviors assigned to the virtual characters impact the realism of the synthesized virtual population.<sup>24</sup> Another study<sup>25</sup> found that besides the size of the virtual crowd and the number of virtual characters, it is also important to consider the orientation of virtual characters within a crowd as well as the formation of the crowd; the study reported that rule-based formations can enhance realism compared with a randomly formatted crowd. Finally, studies have also investigated the egocentric behavior of virtual populations,<sup>26,27</sup> indicating that the absence of collision avoidance between virtual characters and virtual reality users results in lower levels of presence and comfort.

Scenarios with virtual populations or crowds were mainly associated with evacuation-based studies<sup>28,29</sup> and disaster-management protocol testing.<sup>30,31</sup> Previous findings shed light on how crowd movement in multihazard scenarios could affect the decisions of participants.<sup>28,32</sup> Kyriakou et al.<sup>33</sup> examined how characteristics of a virtual population could affect a user's experience in a virtual environment. They found that when the virtual crowd was more interactive,

the participants perceived it as more realistic. Volonte et al.<sup>7</sup> examined the effects of various properties—such as eye gaze, facial expression, and body language—in an interactive virtual reality application. They found that during the negative conditions, the users generally had the least interaction with the virtual population. Marschner et al.<sup>34</sup> explored the role of gaze mechanics and eye contact using virtual agents. They found that mutual eye contact and body orientation were important sources for effective communication between virtual agents and users. The exploratory study conducted by Moustafa and Steed<sup>35</sup> found that group dynamics get carried over into the virtual world, with the only drawback being proper avatar implementation.<sup>36</sup>

This article explores the use of an optimization-based method to optimize affective multicharacter virtual reality experiences. In the past, such techniques have been used for matters of adaptability or replayability in games,<sup>37,38</sup> since such techniques provide the ability to design games that adapt to a variety of constraints and parameters, both during the initialization process and before the game starts.<sup>10,39</sup> Optimization-based methods can even alter a game dynamically in response to events in the game.<sup>40-43</sup> We therefore think that such methods can be adapted and used to synthesize multicharacter experiences that fulfill certain design goals.

This study considered previously published work on human interaction with virtual characters as well as prior work on optimizing virtual reality experiences and combined such knowledge to explore whether it is possible to automatically synthesize virtual reality multicharacter experiences that elicit certain emotional responses of participants. Thus, to the best of our knowledge, this article contributes toward (1) proposing a pipeline for automatically synthesizing virtual reality multicharacter experiences based on several design decisions and (2) validating through a user study that the optimized multicharacter experiences could elicit a certain emotional state in our participants. We think our method could be useful for the automatic synthesis of unique virtual reality experiences that eliminate the need for manual and time-consuming stimulus preparation processes.

### 3 | PRELIMINARIES

In this section, we describe the steps we followed for synthesizing the behaviors of our virtual characters as well as the process that took place to annotate those behaviors.

#### 3.1 | Character behavior

The first step of the study was to develop a dataset of different behaviors that could be assigned to a virtual character, and its connotation was determined to be between neutral and negative affect. The dataset consists of both motion sequences (walk, point, yell, sidestep, and idle) and scripted behaviors (e.g., Look At participant's position in the virtual environment). We also considered triggers that could prompt the behaviors to be displayed to the participants. For this part, we considered the interpersonal space, also known as the proxemic model.<sup>18,44</sup> Based on this model, a character behavior is generated once the position of the user is inside either the intimate, personal, or social space of a virtual character. This allowed us to add another dimension to the behavior display of the virtual population by introducing the concept of *invasion of space*.

The animations to elicit the behaviors were downloaded from Mixamo,<sup>1</sup> which has free animation resources. These animations were then put into our animator controllers in the Unity game engine so that the characters could display them during the simulation. Some animations were used directly without any changes needing to be made, while others were a combination of two or more animations. This was achieved using custom scripts, layering, and transitions. Finally, the virtual characters used in our project were also downloaded from Mixamo. Figure 1 illustrates a virtual character assigned with different behaviors.

#### 3.2 | Behavior annotation

The second step of the preliminary part of our pipeline was the annotation of the developed behaviors of the virtual characters. To do so, we recruited 10 students from our department: five males (age:  $M = 24.40$ ,  $SD = 2.96$ ) and five females (age:  $M = 25.00$ ,  $SD = 1.41$ ).

<sup>1</sup><https://www.mixamo.com/>



**FIGURE 1** Example behaviors that could be assigned to a virtual character and that were used in our project. From left to right: Idle, Point, and Walk. The top row has no Look At functionality, and the bottom row has Look At



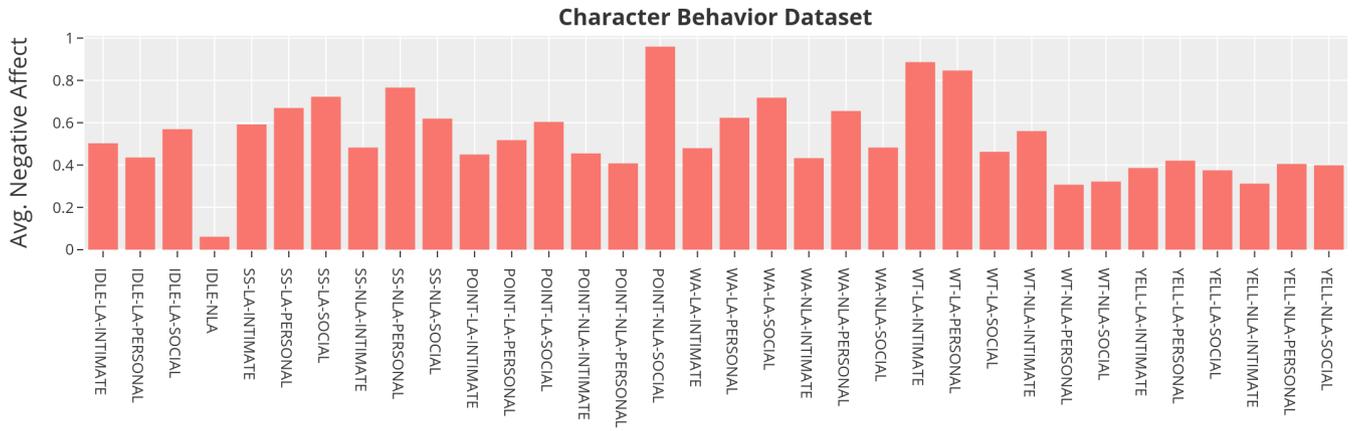
**FIGURE 2** Example scenes that were used for the annotation phase. From left to right: Look At Idle, Look At Point, Look At Yell

The annotation stage was critical to our method because it helped us establish a baseline of how people perceived the different behaviors that were developed. For the annotation phase, 34 scenes were developed (one scene for each behavior, with only one No Look At scene since such behavior was not impacted by the proxemic zones) in Unity. Each scene comprised 10 randomly chosen virtual characters exhibiting a single behavior. It should be noted that each participant experienced each scene with a different combination of the virtual characters. This helped us eliminate the effects that the virtual characters' appearances had on the participants' responses; therefore, we were able to collect data related to the behavior of the virtual characters instead of their appearances. Figure 2 shows examples of the scenes developed for the annotation phase.

For the annotation phase, after the participants observed each scene, they were asked to respond to four negative affect items (upset, alert, nervous, distressed) that were retrieved from the PANAS scale<sup>12</sup> using a 100-point visual scale (1 denoted "Not at all," and 100 denoted "Extremely") in a computer-based environment. Before assigning the collected data to each of the developed behaviors, we screened the reliability of the measure using Cronbach's alpha coefficient. Due to sufficient correlation ( $.77 < \alpha < .91$ ), we used a cumulative score of all four items (upset, nervous, alert, and distressed) as the final result and treated it as a continuous scale. It should be noted that removal of items would not enhance the reliability measures. Once all data was collected, a table was developed in which each behavior had a negative affect value attached to it, which would be fed to the optimization algorithm. Figure 3 shows the results of the annotated study in the form of a bar graph, with values between 0 (low negative affect) and 1 (high negative affect) for each behavior.

## 4 | PROBLEM FORMULATION AND OPTIMIZATION

We wanted to automatically synthesize an affective experience  $E$  that consisted of a number of virtual characters  $c_i$  (each  $c_i$  with a specific annotated behavior) assembled in a sequential order such that  $E = [c_1, c_2, c_3, \dots, c_n]$ . We decided to use



**FIGURE 3** Average negative affect of each behavior derived from the annotation phase (LA, Look At; NLA, No Look At; IDLE, Idle Behavior; SS, Sidestep; POINT, Point Behavior; WA, Walk Across Behavior; WT, Walk Toward; YELL, Yell Behavior; INTIMATE, PERSONAL, and SOCIAL, Proxemics Zones)

three design decisions—the affect cost ( $C_A$ ), variance cost ( $C_V$ ), and duplicate behavior cost ( $C_D$ )—which we encoded to our total cost function  $C_{\text{Total}}(E)$ :

$$C_{\text{Total}}(E) = w_A C_A + w_V C_V + w_D C_D, \quad (1)$$

where  $w_A$ ,  $w_V$ , and  $w_D$  are associated weights for each of the cost terms to control their priority in the total cost function. It should be noted that besides the three proposed cost terms, various other cost terms can be implemented, depending on the characteristics and the design decisions of the developers. The following sections detail each of the cost terms individually.

#### 4.1 | Affect cost

The affect cost term encodes the mean target affect that should be exhibited by the multicharacter experience. The affect cost is defined as:

$$C_A(E) = \frac{1}{|E|} \sum_{c_i} A(c_i) - \sigma_A, \quad (2)$$

where  $|E|$  is the total number of characters,  $\sigma_A$  is the target mean affect value of the synthesized experience (e.g., when  $\sigma_A = .25$  means that the average negative affect of all characters should be close to .25), and  $A(c_i)$  is the negative affect value for the  $c_i$  virtual character of the synthesized scene.

#### 4.2 | Variance cost

We implemented a variance cost term to constrain the variance of affect that should be included in the synthesized scene, therefore overcoming synthesizing scenes with great variations across characters. The variance cost for each synthesized experience  $E$  is defined as:

$$C_V(E) = \frac{1}{|E|} \sum_{c_i} \left( A(c_i) - \bar{A} \right)^2 - \sigma_V, \quad (3)$$

where  $\sigma_V$  denotes the target affect variation and  $\bar{A}$  denotes the mean negative affect of the characters in the scene. Note that when a higher value is assigned to  $\sigma_V$ , characters with higher affective variations will be included in the scene.

### 4.3 | Duplicate behavior cost

Our last cost term was developed to prevent our scene from including multiple characters that exhibit the same behavior. The duplicate behavior cost is defined as:

$$C_D(E) = \frac{1}{\frac{|E|!}{(2!(|E|-2)!)}} \sum_{c_i, c_j} \Gamma(c_i, c_j), \quad (4)$$

where  $\frac{|E|!}{(2!(|E|-2)!)}$  returns the total number of combinations between  $c_i$  and  $c_j$  that are a pair of characters currently in the scene, and  $\Gamma$  returns 1 if the characters have the same behavioral characteristic or 0 if they have different behavior based on the following condition:

$$\Gamma(c_i, c_j) = \begin{cases} 1 & \text{if } B(c_i) == B(c_j) \\ 0 & \text{otherwise,} \end{cases}$$

where  $B(c_i)$  and  $B(c_j)$  represent the behavior of characters  $c_i$  and  $c_j$ , respectively.

### 4.4 | Optimization

The optimization problem was solved using a Markov chain Monte Carlo method<sup>13</sup> called simulated annealing<sup>14</sup> with a Metropolis–Hastings state-searching step.<sup>45</sup> The optimization process starts by initializing a scene with different behaviors assigned to the virtual characters, and the system computes the current total cost  $C_{\text{Total}}(E)$  of that set of characters. In the next iteration, our system proposes a new configuration  $E'$  and computes the proposed total cost  $C_{\text{Total}}(E')$ . The new configuration of our scene is achieved by randomly choosing one of the characters  $c_i$  in the  $E$  configuration and assigning a randomly chosen behavior to that character. The optimizer accepts a proposed configuration either when the proposed cost is lower than the current cost or based on the temperature parameter of the simulated annealing process. In our optimizer, we set a temperature  $t = 1.00$  at the beginning of the optimization, and it is reduced by .10 every 200 iterations. As the temperature parameter decreases, the optimizer becomes greedier in finding the optimal solutions; therefore, fewer high total cost configurations are accepted. The optimization process is completed when the total cost change is less than 2% in the past 500 iterations.

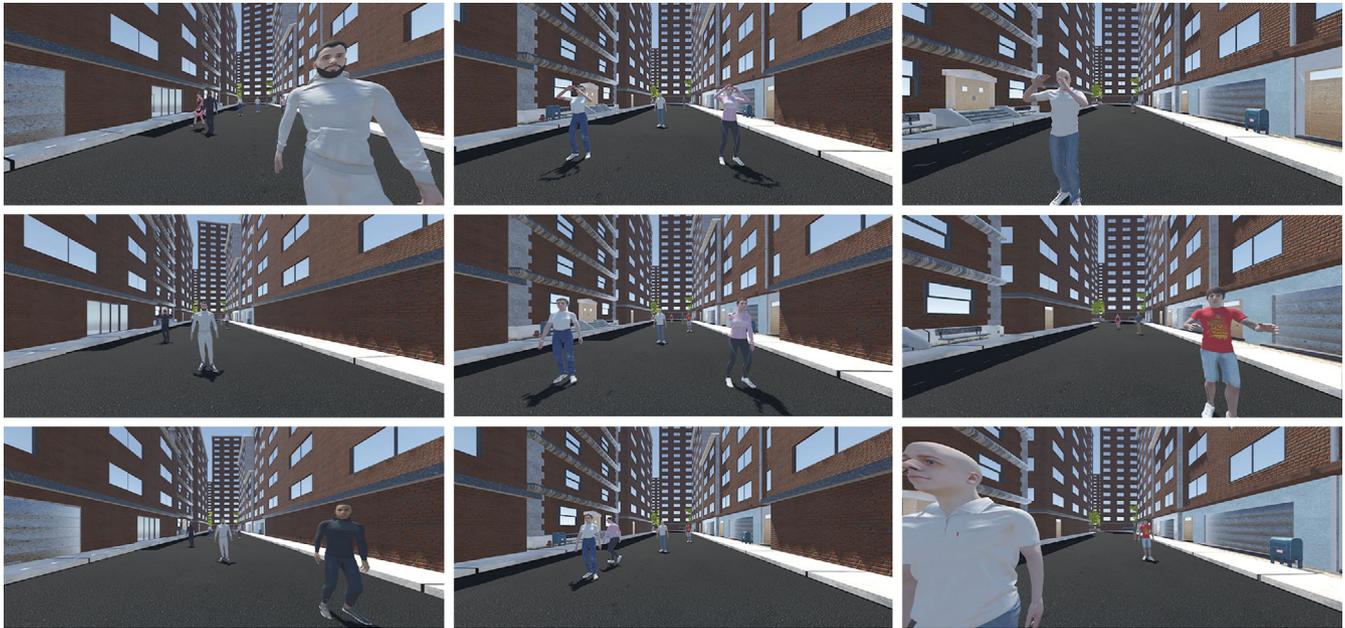
Unless otherwise specified, the target variance was set as  $\sigma_V = .50$ , and the weights were set as  $w_A = 1.00$ ,  $w_V = .10$ , and  $w_D = .10$ . The total number of characters was set to 10, so the optimization algorithm had to consider the affect values for 10 behaviors at a time. The weights assigned to each cost term gave a certain level of control to the designer by allowing them to either prioritize certain terms or to give equal priority to all costs, to create scenarios with different levels of affect, variance, and acceptance of duplicate behaviors. Figure 4 illustrates example scenes synthesized with different target values for negative affect.

## 5 | USER STUDY

The sections below present details about the study that was conducted to evaluate the ability of our method to synthesize affective multicharacter experiences.

### 5.1 | Participants

In our study, 57 university students participated (42 male, 14 females, and one other) within the age range of 18–29 years old ( $M = 19.44$ ,  $SD = 2.28$ ). All participants were recruited through emails, posters, and class announcements. Before participating in our study, the participants signed a consent form that was approved by the Institutional Review Board (IRB) of our university. No compensation was provided for their participation.



**FIGURE 4** Stills from the user point of view of different scenes optimized based on our method. Top: low negative affect; Middle: medium negative affect; Bottom: high negative affect



**FIGURE 5** A top view of the virtual road that was used in our study (left), the virtual environment (middle), and a synthesized scene with virtual characters (right)

## 5.2 | Hardware and application

For both the implementation and the user study, an Asus Republic of Gamers Scar III laptop computer was used with an NVIDIA GeForce RTX 2070 graphics card and 16 GB of RAM. The virtual reality part was implemented using an Oculus Quest head-mounted display and the associated SDKs for Unity. The application was developed in the Unity game engine and run on 60 FPS in Oculus Quest. A virtual street scene was built from the ground up, which consisted of 70 different location coordinates in which the characters were to be spawned (see Figure 5). To standardize the experimental condition across participants, in an initialization phase, we let our application randomly pick 10 spawn points and assign each scene to one of the virtual characters we had in our database.

Given the initialized scene, we ran our optimization method that searched to find the optimal behavior that should be assigned to each character. Thus, besides the behavior assigned to each virtual character, the rest of the virtual reality scene would remain the same (i.e., took place on the same stretch of virtual road and had the exact same 10 characters placed in the exact same locations in the virtual environment to not induce any bias due to different visual styles or placement of the characters).

At the beginning, the participant was placed at one end of the street, and Unity's navigation mesh functionality was used to find the path that the virtual reality camera should follow to reach the other end of the road. It should be noted that we set the camera's speed to not exceed 1.2 m/s. This choice was based on the US Manual of Uniform Traffic Control Devices,<sup>46</sup> which states that the normal walking speed of humans has been estimated to be 1.2 m/s. The screen fades in

from black to ease the participant into the environment. They then move through the virtual population (without needing any input), and at the end, the screen fades out to black, signaling the end of the simulation.

### 5.3 | Experimental conditions

Three experimental conditions were developed for the study by using three different target values of negative affect. These were low negative affect, medium negative affect, and high negative affect of the synthesized virtual population. Details of the conditions are given below:

- *Low negative affect:* In this condition, the 10 behaviors that were selected by the optimization algorithm had to achieve an overall low total cost of the current scene's configuration. The chosen target values were  $\sigma_A = .30$  and  $\sigma_V = .50$ .
- *Medium negative affect:* In this condition, the 10 behaviors that were selected by the optimization algorithm had to achieve an overall medium total cost of the current scene's configuration. The chosen values were  $\sigma_A = .50$  and  $\sigma_V = .50$ .
- *High negative affect:* In this condition, the 10 behaviors that were selected by the optimization algorithm had to achieve an overall high total cost of the current scene's configuration. The chosen values were  $\sigma_A = .70$  and  $\sigma_V = .50$ .

### 5.4 | Measurements

Participants' affective ratings were collected using the four items of the PANAS<sup>12</sup> scale that were also used in the annotation phase: upset, alert, nervous, and distressed. The participants were asked to respond using a 100-point visual scale in which 1 denoted "Not at all" and 100 denoted "Extremely." This self-reported data was used to analyze the overall user-reported negative affect intensities for the results, as it gave a direct idea about the user-perceived negative affect of each of the generated multicharacter virtual reality experiences. Finally, a designated space in the survey was used to allow participants to leave their comments regarding the virtual reality application and the conditions they exposed. Figure 6 illustrates the visual scale that was used in our study to collect the negative affect rating of our participants.

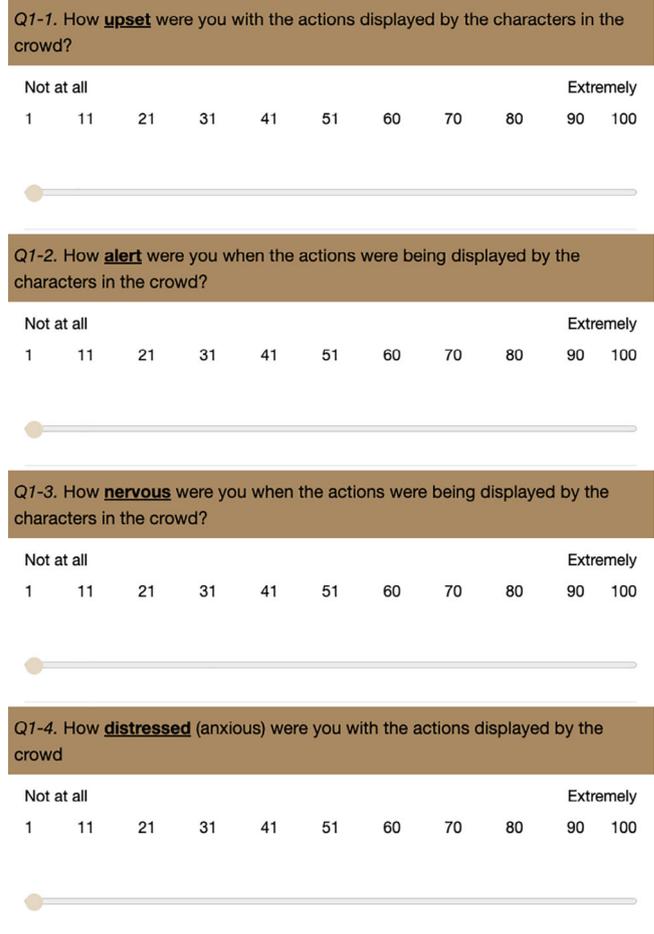
### 5.5 | Procedure

When participants arrived at our lab space, they first sanitized their hands using university-prescribed hand sanitizers to minimize the risk of the spread of the COVID-19 virus. The participants and the researcher wore masks for preventive measures, and appropriate physical-distancing measures were taken whenever necessary. The participants were given a consent form that was approved by the IRB of our university; the researcher advised them to give it a read-through. At the end of the consent form, they were asked whether they were 18 years or older and gave their consent to participate in the study. If they answered that they wanted to take part in the study, a demographics questionnaire (age, gender, whether they had VR experience before, and so forth) appeared. Once they answered that, the researcher briefed them on how the study was going to proceed. First, they were familiarized with the consumer virtual reality headset that would be used (the Oculus Quest) and how to adjust it to their comfort level, and then they were informed that they would be experiencing three scenes with a virtual crowd. At the end of each scene they experienced, they were asked to respond to the questionnaire based on what they had just experienced and rate the negative affect of the virtual population on a scale of 1–100. This was repeated two more times, for a total of three times (for three scenes). Since this was a between-group study, all participants experienced all three conditions mentioned in Section 5.3. Note that the order of the three different conditions was counterbalanced by a Graeco–Latin square<sup>47</sup> for controlling potential carry-over effects. A participant during the study is shown in Figure 7. The total duration of the study did not exceed 30 min.

### 5.6 | Results

The analysis of the results was done using one-way repeated measures analysis of variance using the three conditions as independent variables and the questionnaire responses as dependent variables. The internal validity of

**FIGURE 6** The 100-point visual scales that were developed to capture the negative affective experience of our participants

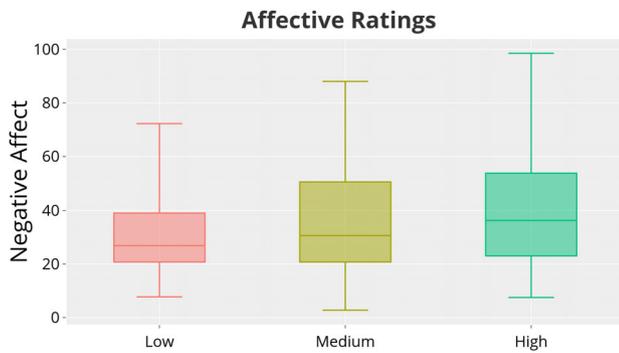


**FIGURE 7** A participant wearing the head-mounted display and observing one of the synthesized scenes



the questionnaire's scale was measured using Cronbach's alpha coefficient. With sufficient scores ( $.75 < \alpha < .81$ ), we used a cumulative score. Removal of items would not enhance the reliability measures. The normality assumption of the ratings was evaluated with Shapiro–Wilk tests at the 5% level and with the Q-Q plots of the residuals. Post hoc comparisons were conducted using Bonferroni corrected estimates. A  $p < .05$  value was deemed statistically significant.

Our analysis revealed significant results [ $\Lambda = .726, F(2, 52) = 9.928, p < .001; \eta_p^2 = .274$ ]. Post hoc comparisons showed that the low negative affect condition ( $M = 31.71, SD = 16.81$ ) was rated lower than that of the medium negative affect condition ( $M = 36.73, SD = 21.13$ ) at the  $p = .041$ , and high negative affect condition ( $M = 40.68, SD = 22.50$ ) at  $p = .001$ . Moreover, the medium negative affect condition was rated lower than that of the high negative effect condition at the  $p = .040$ . Boxplots of our results are shown in Figure 8.



**FIGURE 8** Negative affect ratings of our participants for low, medium, and high negative affect conditions

## 5.7 | Discussion

Our experimental study was conducted to evaluate whether our pipeline could be used to synthesize an affective virtual reality experience that would be perceived by participants in an expected way. Our study evaluated three experimental conditions consisting of low, medium, and high negative affect. Based on the collected self-reported ratings of our participants, we conclude that it is indeed possible to optimize the behavior of the virtual characters and synthesize multicharacter experiences that affect participants in an expected way: low negative affect was rated lower than medium and high negative affect, and medium negative affect was rated lower than the high negative affect multicharacter experience.

Besides the self-reported data that was collected, the participants were also asked to provide feedback about the exposed conditions. Notable responses pointed toward the fact that they could perceive the hostile nature of the virtual population and them being singled out or threatened by the characters. Moreover, a few participants reported that they had a more negative experience when virtual characters invaded their spatial zones, confirming that pairing proxemics with behaviors is an acceptable approach to inducing a negative feeling in the participant. Some notable comments were the following: for a high affect scene: *“This particular event was the most threatening, as the virtual characters made hostile remarks, causing a feeling of being uncomfortable and anxious.”*; for a medium affect scene: *“Yes, I did notice a difference in all three of the crowds. They seemed to react differently every time I walked down the street, but overall, I still felt like an outsider and unwelcome in the environment. I felt that my personal space was not being respected.”*; and for the low affect scene with respect to the medium affect scene: *“Overall, the crowd was not as negative as the simulation before.”* Thus, given both the self-reported ratings and the comments from our participants, it can be said that our method made participants perceive the synthesized scenes differently since the multicharacter experiences impacted their affect levels differently.

## 6 | CONCLUSIONS AND FUTURE WORK

This article presented a method for synthesizing multicharacter experiences that could elicit a certain amount of negative affect on virtual reality users. By conducting a user study, we validated the possibility of automatically synthesizing virtual reality multicharacter experiences that trigger participants’ emotional responses in an expected way. Thus, we conclude that the proposed algorithm is successful in optimizing multicharacter scenarios and generating a set of characters that can induce a specific level of negative affect.

Besides the presented methodology and the interesting findings, a few limitations should be mentioned. First, no audio was involved, which could have added to the aspect of realism and/or added to the factors that would influence the participants’ responses. Second, we think that the dataset itself is limited to a small number of behaviors that could be assigned to virtual characters; hence, a larger behavior dataset should be developed and evaluated. Third, for this study, we only considered annotating the dataset and evaluating our method using self-reported ratings for negative affect, which might not always indicate the true feelings of the participants. We think that the collection of objective data such as collection of arousal levels through galvanic skin response and electroencephalogram could help us understand even more the power of optimizing virtual reality multicharacter affective experiences.

We consider this study a starting point for studies dealing with affect-driven multicharacter scenarios. Besides the mentioned limitations that should be explored, we plan to expand our pipeline in various other directions. Among others, we would like to synthesize more complex scenarios that participants could observe as well as interactive applications such as virtual reality games, in which participants could directly interact with the virtual characters. For both mentioned cases, we would like to explore design decisions (cost terms) that should be included and artificial intelligence methods such as the use of behavior trees<sup>48,49</sup> for authoring event-centric and affect-driven multicharacter narratives.

**ORCID**

Christos Mousas  <https://orcid.org/0000-0003-0955-7959>

**REFERENCES**

1. Gilroy S, Porteous J, Charles F, Cavazza M. Exploring passive user interaction for adaptive narratives. *Proceedings of the ACM International Conference on Intelligent User Interfaces*; 2012. p. 119–28.
2. Kätsyri J, de Gelder B, Takala T. Virtual faces evoke only a weak uncanny valley effect: an empirical investigation with controlled virtual face images. *Perception*. 2019;48(10):968–91.
3. Aranyi G, Pecune F, Charles F, Pelachaud C, Cavazza M. Affective interaction with a virtual character through an fNIRS brain-computer interface. *Front Comput Neurosci*. 2016;10:70.
4. Saberi M, Bernardet U, DiPaola S. An architecture for personality-based, nonverbal behavior in affective virtual humanoid character. *Proc Comput Sci*. 2014;41:204–11.
5. Mousas C, Anastasiou D, Spantidi O. The effects of appearance and motion of virtual characters on emotional reactivity. *Comput Hum Behav*. 2018;86:99–108.
6. Volante M, Babu SV, Chaturvedi H, Newsome N, Ebrahimi E, Roy T, et al. Effects of virtual human appearance fidelity on emotion contagion in affective inter-personal simulations. *IEEE Trans Vis Comput Graph*. 2016;22(4):1326–35.
7. Volante M, Hsu YC, Liu KY, Mazer JP, Wong SK, Babu SV. Effects of interacting with a crowd of emotional virtual humans on users' affective and non-verbal behaviors. *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*; 2020. p. 293–302.
8. Krogmeier C, Mousas C, Whittinghill D. Human–virtual character interaction: toward understanding the influence of haptic feedback. *Comput Animat Vir Worlds*. 2019;30(3-4):e1883.
9. Alghofaili R, Solah MS, Huang H, Sawahata Y, Pomplun M, Yu LF. Optimizing visual element placement via visual attention analysis. *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*; 2019. p. 464–73.
10. Xie B, Zhang Y, Huang H, Ogawa E, You T, Yu LF. Exercise intensity-driven level design. *IEEE Trans Vis Comput Graph*. 2018;24(4):1661–70.
11. Liu H, Wang Z, Mousas C, Kao D. Virtual reality racket sports: virtual drills for exercise and training. *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality*; 2020. p. 566–76.
12. Watson D, Clark LA, Tellegen A. Development and validation of brief measures of positive and negative affect: the PANAS scales. *J Pers Soc Psychol*. 1988;54(6):1063.
13. Brooks S. Markov chain Monte Carlo method and its application. *J Royal Stat Soc Ser D (Stat)*. 1998;47(1):69–100.
14. Kirkpatrick S, Gelatt CD, Vecchi MP. Optimization by simulated annealing. *Science*. 1983;220(4598):671–80.
15. Bopp JA, Mekler ED, Opwis K. Negative emotion, positive experience? emotionally moving moments in digital games. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*; 2016. p. 2996–3006.
16. Knapp ML, Augustine J. *Handbook of interpersonal communication*. Thousand Oaks, California, United States: Sage; 2002.
17. Argyle M. *Bodily communication*. Milton Park, Abingdon-on-Thames, Oxfordshire, England, UK: Routledge; 2013.
18. Hall ET, Birdwhistell RL, Bock B, Bohannon P, Diebold AR Jr, Durbin M, et al. Proxemics [and comments and replies]. *Curr Anthropol*. 1968;9(2/3):83–108.
19. Wiggins JS. *The five-factor model of personality: theoretical perspectives*. New York, United States: Guilford Press; 1996.
20. Allbeck JM, Badler NI. Creating crowd variation with the ocean personality model. *Proceedings of the International Joint Conference on Autonomous agents and Multiagent Systems*; 2008. p. 1217–1220.
21. Peters C, Ennis C. Modeling groups of plausible virtual pedestrians. *IEEE Comput Graph Appl*. 2009;29(4):54–63.
22. McDonnell R, Ennis C, Dobbyn S, O'Sullivan C. Talking bodies: sensitivity to desynchronization of conversations. *ACM Trans Appl Percept*. 2009;6(4):1–8.
23. Ennis C, McDonnell R, O'Sullivan C. Seeing is believing: body motion dominates in multisensory conversations. *ACM Trans Graph*. 2010;29(4):1–9.
24. Ennis C, Peters C, O'Sullivan C. Perceptual effects of scene context and viewpoint for virtual pedestrian crowds. *ACM Trans Appl Percept*. 2011;8(2):1–22.
25. Huerre S, Lee J, Lin M, O'Sullivan C. Simulating believable crowd and group behaviors. *ACM SIGGRAPH ASIA Courses 2010*; New York, United States: 2010. p. 1–92.
26. Stocker C, Allbeck JM, Badler NI. Being a part of the crowd: towards validating VR crowds using presence. *Proceedings of the International Joint conference on Autonomous Agents and Multiagent Systems*; 2008. p. 136–42.
27. Ahn J, Wang N, Thalmann D, Boulic R. Within-crowd immersive evaluation of collision avoidance behaviors. *Proceedings of the ACM SIGGRAPH International Conference on Virtual-Reality Continuum and Its Applications in Industry*; 2012. p. 231–238.
28. Xu M, Xie X, Lv P, Niu J, Wang H, Li C, et al. Crowd behavior simulation with emotional contagion in unexpected multihazard situations. *IEEE Trans Syst Man Cybern Syst*. 2019;51(3):1567–1581.
29. Stamatopoulou I, Sakellariou I, and Kefalas P. Formal agent-based modelling and simulation of crowd behaviour in emergency evacuation plans. *Proceedings of the IEEE International Conference on Tools with Artificial Intelligence*. vol. 1; 2012. p. 1133–8.
30. Xu ML, Jiang H, Jin XG, Deng Z. Crowd simulation and its applications: recent advances. *J Comput Sci Technol*. 2014;29(5):799–811.
31. Helbing D, Farkas I, Vicsek T. Simulating dynamical features of escape panic. *Nature*. 2000;407(6803):487–90.
32. Moussaïd M, Kapadia M, Thrash T, Sumner RW, Gross M, Helbing D, et al. Crowd behaviour during high-stress evacuations in an immersive virtual environment. *J R Soc Interface*. 2016;13(122):20160414.

33. Kyriakou M, Pan X, Chrysanthou Y. Interaction with virtual crowd in immersive and semi-immersive virtual reality systems. *Comput Animat Vir Worlds*. 2017;28(5):e1729.
34. Marschner L, Pannasch S, Schulz J, Graupner ST. Social communication with virtual agents: the effects of body and gaze direction on attention and emotional responding in human observers. *Int J Psychophysiol*. 2015;97(2):85–92.
35. Moustafa F, Steed A. A longitudinal study of small group interaction in social virtual reality. *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*; 2018. p. 1–10.
36. Nelson MG, Mazumdar A, Jamal S, Chen Y, Mousas C. Walking in a crowd full of virtual characters: effects of virtual character appearance on human movement behavior. *Proceedings of the International Symposium on Visual Computing*; 2020. p. 617–29.
37. Hartsook K, Zook A, Das S, Riedl MO. Toward supporting stories with procedurally generated game worlds. *Proceedings of the IEEE Conference on Computational Intelligence and Games*; 2011. p. 297–304.
38. Hullett K, Mateas M. Scenario generation for emergency rescue training games. *Proceedings of the International Conference on Foundations of Digital Games*; 2009. p. 99–106.
39. Zhang Y, Xie B, Huang H, Ogawa E, You T, Yu LF. Pose-guided level design. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*; 2019. p. 1–12.
40. Smith AM, Andersen E, Mateas M, Popović Z. A case study of expressively constrainable level design automation tools for a puzzle game. *Proceedings of the International Conference on the Foundations of Digital Games*; 2012. p. 156–63.
41. Rowe J, Mott B, Lester J. Optimizing player experience in interactive narrative planning: a modular reinforcement learning approach. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*; 2014. p. 160–6.
42. Harrison B, Roberts D. Analytics-driven dynamic game adaption for player retention in a 2-dimensional adventure game. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*; 2014. p. 23–29.
43. Luo L, Yin H, Cai W, Lees M, Zhou S. Interactive scenario generation for mission-based virtual training. *Comput Animat Vir Worlds*. 2013;24(3-4):345–54.
44. Bailenson JN, Blascovich J, Beall AC, Loomis JM. Interpersonal distance in immersive virtual environments. *Personal Soc Psychol Bull*. 2003;29(7):819–33.
45. Chib S, Greenberg E. Understanding the metropolis-hastings algorithm. *Am Stat*. 1995;49(4):327–35.
46. Federal Highway Administration. U.S. manual on uniform traffic control. Revisions 1 and 2 incorporated ed. Washington, DC: U.S. Department of Transportation Federal Highway Administration; 2003.
47. Grant DA. The Latin square principle in the design and analysis of psychological experiments. *Psychol Bull*. 1948;45(5):427.
48. Kapadia M, Shoulson A, Steimer C, Oberholzer S, Sumner RW, Gross M. An event-centric approach to authoring stories in crowds. *Proceedings of the International Conference on Motion in Games*; 2016. p. 15–24.
49. Shoulson A, Gilbert ML, Kapadia M, Badler NI. An event-centric planning approach for dynamic real-time narrative. *Motion on games*; New York, United States: 2013. p. 121–30.

## AUTHOR BIOGRAPHIES



**Angshuman Mazumdar** holds a B.Tech in Electronics and Communications Engineering from VIT Vellore, Vellore (India) and is currently pursuing his MS in Computer Graphics Technology at Purdue University, West Lafayette. His research interests include virtual reality, computer games and game development, computer graphics, and human–computer interaction.



**Christos Mousas** is an Assistant Professor at Purdue University, West Lafayette. His research revolves around virtual reality, computer graphics, computer animation, and human–computer interaction. In the past, he has been affiliated with the Department of Computer Science at Dartmouth College as Postdoctoral Researcher, and with the Department of Cultural Technology and Communication of the University of the Aegean as Research Assistant. He holds a PhD in Informatics and an MSc in Multimedia Applications and Virtual Environments both from the School of Engineering and Informatics of the University of Sussex, and an integrated Masters degree in Audiovisual Science and Art from Ionian University. He is a member of IEEE.

**How to cite this article:** Mazumdar A, Mousas C. Synthesizing affective virtual reality multicharacter experiences. *Comput Anim Virtual Worlds*. 2021;e2004. <https://doi.org/10.1002/cav.2004>