Modeling Emotional Contagion

Jason Tsai[1], Emma Bowring[2], Stacy Marsella[3], and Milind Tambe[1]

[1] University of Southern California, Los Angeles, CA 90089
`{jasontts, tambe} @usc.edu`
[2] University of the Pacific, Stockton, CA 95211
`ebowring@pacific.edu`
[3] USC Institute for Creative Technologies, Playa Vista, CA 90094
`marsella@ict.usc.edu`

Abstract. In social psychology, emotional contagion describes the widely observed phenomenon of one person's emotions being influenced by surrounding people's emotions. While the overall effect is agreed upon, the underlying mechanism of the spread of emotions has seen little quantification and application to computational agents.

In this paper, we explore computational models of emotional contagion by implementing two models (Bosse et al., Durupinar et al.) and augmenting them to better model real world observations. Our additions include examining the impact of physical proximity and authority figures. We show that these additions provide substantial improvements to the qualitative trends of emotion spreading, more in line with expectations than either of the two previous models. We also evaluate their impact on evacuation safety in an evacuation simulation, ESCAPES, showing substantial differences in predicted safety based on the contagion model.

1 Introduction

Emotional contagion has been shown to arise in a wide range of scenarios in everyday life. Its effects are felt in homes everyday when comedic shows employ laugh tracks to elicit stronger emotional responses from audiences. Less often, but with far more severe implications, it is also felt during the spread of fear and anxiety that surrounds any crowd-based disaster. With the growing interest in emotional modeling in agents, the contagion of these emotions can no longer be marginalized when modeling crowds. Recent work has sought to quantify the qualitative findings of social psychology into useable models with varying degrees of success. Bosse et al. (VU University) introduced one of these in 2009 [3] that used an interaction-based model derived directly from social psychology theories of emotional contagion wherein members of the simulation converged towards a weighted average of each emotional type. Durupinar et al. [6] used an epidemiological-style threshold-based model wherein successive interactions with emotionally 'infected' people raises the chance of infection.

While both of these models showed predictions in line with some qualitative findings in social psychology studies, they are inherently very different models of the same phenomenon. Although this type of detail may not be important for understanding the contagion of joy with the use of laugh tracks, it may offer substantially different predictions on the outcome of emotionally-charged crowd simulations. It is in this context that we explore the modeling of emotional contagion. In particular, we use an evacuation simulation called ESCAPES, described more in Section 3, as the test bed for different models of fear contagion.

Despite their promise, both models come short when applied to an evacuation simulation such as ESCAPES. First, the models do not explore the impact of proximity on the effect of contagion. The VU model provides a parameter (channel strength) that can allow for this manipulation, but never provides guidance on how it should be done. Second, the authors have not introduced guidelines for designing 'special' agents such as authority figures that might have stronger resistance to fear contagion and be trained to reduce fear in other agents. Again the VU model provides parameters that allow for this manipulation (receiver openness, sender expressiveness), but have not explored this in their work thus far. The Durupinar model provides no mechanism for either feature.

We propose to augment the VU and Durupinar models with proximity-based effects and authority figure calming and examine their performance in the context of ESCAPES. Through extensive experimental results, we show that without proximity's effect on contagion, neither model produces qualitatively believable results. After incorporating authority calming effects into each model, we show that the spread of emotions through the population again changes drastically. Finally, as a second-order effect, we also show that the evacuation simulation's outcome predictions vary substantially, motivating the need for an accurate model of contagion and authority effects.

2 Related Work

Seminal works in social psychology first began the discussion around emotional contagion. In particular, Hatfield et al. [7] first codified the observed phenomena that were just beginning to receive researcher attention. Follow-up work by the co-authors as well as in related fields such as Barsade et al. [1] in managerial sciences continued to detail the effects of the phenomenon in new domains. Recently, there have been attempts to begin quantifying emotional contagion and explore cross-cultural variations in attributes that effect emotional contagion [5, 9].

From a computational perspective, the previously mentioned work of Bosse et al. (VU model) and Durupinar et al. are two of the most recent models of emotional contagion upon which a few follow-up works have been based [2, 8].

3 ESCAPES

Although not the focal point, the ESCAPES evacuation simulation [10] serves as the test bed for our models of emotional contagion, so we describe it briefly here. ESCAPES focuses on the features identified by experts that particularly effect airport evacuations, including first time visitors' incomplete knowledge of the area, the presence of families, and the presence and effects of authority figures [4]. We also model fear and model its impact on behavior by increasing the speed of more fearful agents.

Although the second-order impacts of changing the emotional contagion model, such as evacuation rates and safety, are dependent upon the specific simulation implementation, we use ESCAPES as a test bed to illustrate obvious deficiencies in the base models that would occur in any simulation with a spatial component. We detail these in Sections 8 and 9 and also highlight second-order impacts that the different models have upon the evacuation as a whole.

For all the experiments discussed in Section 8 and 9, the same scenario was used (spatial layout can be seen in Figure 3) and 30 trials were run for each setting. It features 2 large spaces, each with an exit, connected by hallways which are lined with smaller spaces that represent shops. 15 seconds into the simulation, an event occurs at the center of the scenario, inciting fear (0.75 for nearby agents, 0.1 for further agents) and a need to evacuate that is communicated by authority figures to pedestrians. The scenario features 100 normal pedestrians, including 10 families of 4 each, as well as 10 authority figures that patrol the scenario.

4 VU Model

Introduced in 2009 by Bosse et al. [3] and built upon in [2, 8], this model is an independent interaction-based model. The initial version that we use here moves people towards a weighted-average of the group's emotional levels. Since subsequent works do not address the needs of our evacuation simulation, we begin with a discussion of the original model and its attributes. In Sections 6 and 7 we will further explore this base model to examine the impact of proximity and a partcular model of authority figures.

We briefly mention the primary components of interest in the VU model here. In particular, emotional contagion is modeled using 5 parameters for every pair of people that may interact: level of sender's emotion q_S, level of receiver's emotion q_R, sender's expressiveness ϵ_S, receiver's openness δ_R, and the channel strength between S and R α_{SR}. All values are numbers in the interval $[0, 1]$. The parameters are derived from the theory put forth in [1], giving the model a theoretical foundation.

At each time step, each agent calculates the average emotional transfer from all relevant agents. Specifically, from a sender S to a receiver R, the strength of the emotion received would be $\gamma_{SR} = \epsilon_S \cdot \alpha_{SR} \cdot \delta_R$. Logically, stronger channel, stronger sender expressiveness, and stronger receiver openness all lead to stronger emotional transfer to the receiver. [3] details the mathematical formulation, but, qualitatively, the fear level of an agent converges towards a weighted average of the group's fear level. The speed at which this convergence occurs as well as the weighting depend on the parameter settings for the channel strength, expressiveness, and openness for each agent.

5 Durupinar Model

As opposed to independent interaction models, Durupinar et al. used a threshold model based on epidemiological models of disease contagion. While many types of epidemiological models exist, Durupinar implemented a simple version with only *susceptible* and *infected* states (as opposed to recovered, innoculated, etc. states). The model's applicability to emotional contagion was not discussed in the initial publication, but its use assumes similarity between disease spread and emotion spread.

Each agent begins with a randomized threshold drawn from a pre-determined log-normal distribution. At each time step, for each agent, a random agent is chosen from the relevant population group and if the agent is infected, will generate a random dose drawn from a pre-determined log-normal distribution and pass it to the original agent. If the agent is not infected, then a dose of 0.0 is generated. Each agent maintains a running history of the last K doses received. If the cumulative total of all doses in the agent's history exceeds his threshold, the agent enters the infected state. This causes the emotion level to be set to 1.0 with an exponential decay towards 0.0, at which point the agent re-enters the susceptible state. The random dose and threshold are generated from log-normal distributions with user-specified averages and standard deviations and K is a static global variable.

6 Proximity

When used in a simulation that includes physical space, an immediate deficiency arises in both models - the lack of specification of proximity's role in contagion. In any such simulation, proximity must enter the equation in some form. The VU model provides the channel strength parameter, which, if varied properly, can incorporate proximity effects into the contagion. Despite this, the authors did not provide guidance or exploration of possible implementations, thus we explore one in this work. The Durupinar publication did not provide experimental results specifically pertaining to the contagion and do not have a variable parameter such as in the VU model.

In this work, we implement a fixed neighborhood of effect for all agents within both models. Only agents within the specified distance are used in the model's calculations for contagion. In the VU model, this is equivalent to setting all channel strengths to 0.0 for pairs of agents that are too distant from each other and setting the remaining channel strengths to their pre-set levels otherwise. In the Durupinar model, this was a direct augmentation to the contagion effect, where we restrict the population used to neighboring agents only.

7 Authority Figure Effects

The second important modification that we require is specific to our scenario of evacuations, but is an example of the more general need for a contagion model to allow 'special contagion' agents. As noted in recent research [4], the role of authority figures is extremely important in evacuations not only for the information they provide but also for the calming effect they bring to anxious or fearful crowds. Neither model inherently discusses the implementation of agents with unique contagion attributes.

The VU model's individual-specific parameters allows for simple settings that would logically correspond with an authority figure (or to other special agents), but the authors did not explore possible impacts or implementations. While many implementations are possible, we choose to set all authority figures' openness parameters to 0.0, simulating the effect of proper training preventing authority figures from being susceptible to others' influences on their emotions. In combination with the proximity effect, this encourages agents near authority figures to calm their emotions towards 0.0 at each time step, producing the desired effect.

The Durupinar model specifies only population-wide, randomized parameter settings, necessitating model-level augmentations to include unique authority effects. Again, many implementations are possible. We choose to reproduce the resistance of authority figures to fear contagion by removing their contagion module entirely. In addition, to reproduce the calming effect that the VU model can naturally produce with the openness and expressiveness parameters, we introduce two changes to the base Durupinar model. First, we halve the level of fear in the agents surrounding authority figures at each time step. Second, we introduce an inoculated state that agents enter upon contact with an authority figure. They remain in this state for a fixed period of time that is reset as long as they remain in the presence of an authority figure. The second addition prevents the situation where a group of fearful agents at different points in the decay process simply pass fear back and forth to each other despite the presence of an authority figure. With these two augmentations, we are able to reproduce the authority calming effect noted by [4].

8 VU Experiments

We first explore the implications of varying the parameters in the base VU model when applied to the scenario described in Section 3. Then we show results pertaining to the rate and strength of emotion spread under the different versions of the model. Finally, we briefly touch on the implications on the predictions of safety under each of the versions of the model as they appear in the ESCAPES simulation engine.

8.1 Sensitivity Analysis

The parameters of interest in the VU model were the channel strengths, individual expressiveness settings, and individual openness settings. Given that we had a whole population of agents, we elected to use randomly drawn values for each of these based on a normal distribution. We explored variations of the averages and standard deviations used, but surprisingly, none yield substantial changes in the outcome of the simulation from both a contagion perspective (i.e., how the fear spread) and a safety analysis (i.e., how safe the evacuation was). The only exception was, unsurprisingly, when the receiver openness parameter varied tightly around a very low mean, leaving many agents with 0.0 openness. This caused the majority of agents to remain at their initial fear level, sometimes raising all agents' fear levels, which was vastly different from the convergence behavior seen in the other settings.

Figure 1 plots the percentage of people with low fear (≤ 0.1) on the y-axis and the time step on the x-axis. Figure 1a shows the results for variations in average channel strength whereas Figure 1b shows the same results for variations in average receiver openness. In both cases, the parameter being explored varied from 0.1 to 0.9 in increments of 0.2 while keeping a fixed standard deviation of 0.1 and the other two parameters were fixed with an average of 0.5 and a standard deviation of 0.1. As expected, when an event first occurs, those near it become fearful, hence the initial dip. However, due to the global convergence of fear levels and the fact that the vast majority of agents have 0.0 fear and do not know of the event, fear levels quickly decrease back to ≤ 0.1

levels. The tightness of the lines implies that the trend is robust to variations in the average channel strength. The same trend can be seen when the average openness is varied, with the exception of the previously mentioned situation. Similar tightness of lines was observed in other parameter variations.

We also conducted experiments exploring the second-order effects on safety, as measured by the ESCAPES system. In particular we examined the evacuation rates of pedestrians as well as the number of collisions experienced on average. Neither set of results showed significant variation through the parameter space, indicating the results' robustness to parameter variation.

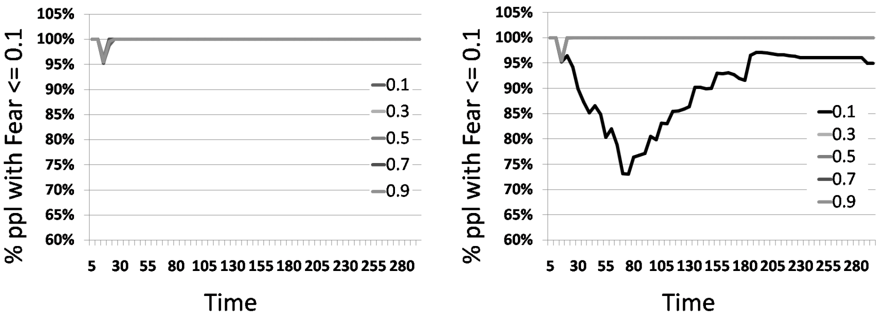

(a) Variations in Channel Strength

(b) Variations in Openness

Fig. 1. Percentage of low-fear agents

8.2 Contagion Analysis

Now we discuss the effect on contagion as we include proximity effects and authority figure calming in the base model . Given the relative indifference of the model to parameter variations, we elect to use median values of 0.5 for the average of all parameters and fix the standard deviations at 0.1 for the results shown in this section.

Figure 2 shows the contagion trends of agents in the simulation under the three different models: original base model with 'worldwide' neighborhood, a model with a limited neighborhood of contagion, and a model with the limited neighborhood in conjunction with the authority figure modification. Each graph shows the percentage of agents remaining in the simulation that possess the labeled level of fear: ≤ 0.1 and ≥ 0.75 on the y-axis and time steps on the x-axis.

As can be seen, the trends are drastically different in each case. In particular, the base model always sees an extremely steep decrease in fear levels as the majority of agents do not know of the event and possess 0.0 fear, lowering the convergence target to near 0.0. When proximity is introduced, the fear levels reduce slowly, as high fear agents must pass close to low fear agents for this to occur. This leaves a large percentage of agents hovering above 0.1 through the entire duration of the simulation (recall that there is no decay of emotion in the VU model). When authority figure calming is introduced, a middle ground between ‘worldwide’ and ‘proximity’ is achieved as authority figures are able to constantly reduce nearby agents’ fear levels despite them not encountering new 0-fear agents.

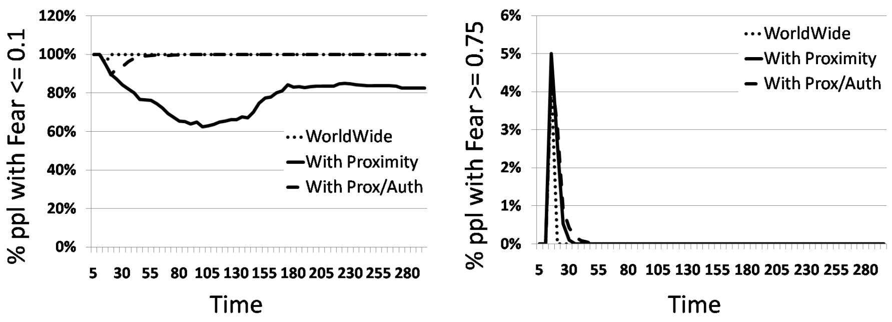

(a) Low fear pedestrians

(b) High fear pedestrians

Fig. 2. Comparison of model with additions

Perhaps the strongest indication of the impact that our proximity addition has on realism comes from a series of illustrative snapshots in time of the locations of fearful agents in the scenario. In Figure 3, agents with fear greater than 0.1 are shown as red dots, agents with less than 0.1 but non-zero fear are shown as white dots. Figure 3b shows the location of fearful agents at time step 16 in the base model without proximity. Figure 3c shows the same snapshot for the model with proximity and without authorities. As can be seen, in the first few seconds following the event, agents throughout the scenario instantly become slightly fearful as they converge towards the fearful agents’ emotional level. When proximity is incorporated, however, a much more realistic spread can be seen with nearby agents becoming fearful.

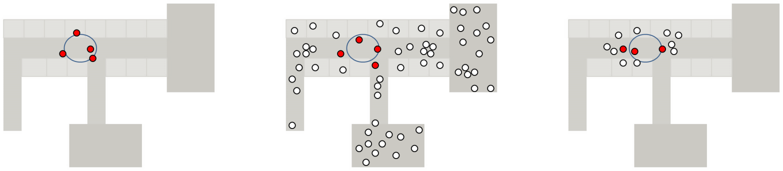

(a) Both models at 15s

(b) Base model at 16s

(e) Proximity at 16s

Fig. 3. Effect of proximity in VU model.

8.3 Safety Analysis

Now we evaluate the impact of the models of contagion on the actual evacuations as measured in the ESCAPES system. In particular we show the evacuation rates and average number of collisions of pedestrians in the simulation. Clearly, faster evacuation rates and lower number of collisions indicate better evacuations.

Figure 4a shows the percentage of pedestrians remaining in the simulation on the y-axis and the time step on the x-axis. Figure 4b shows the average number of collisions accumulated by people remaining in the simulation on the y-axis and the time step on the x-axis. The evacuation rate remains unchanged but there are noticeable differences in the number of collisions. In the ‘worldwide’ model and the model with both proximity and authorities, the number of collisions slopes up substantially slower than the model with proximity only as a result of the slower pace of people. The number of authorities was very high in these simulations, creating a situation similar to the ‘worldwide’ model with very little fear in the population. Thus, although the augmentations do not impact evacuation time, the prediction of safety as measured by the number of collisions is strongly affected.

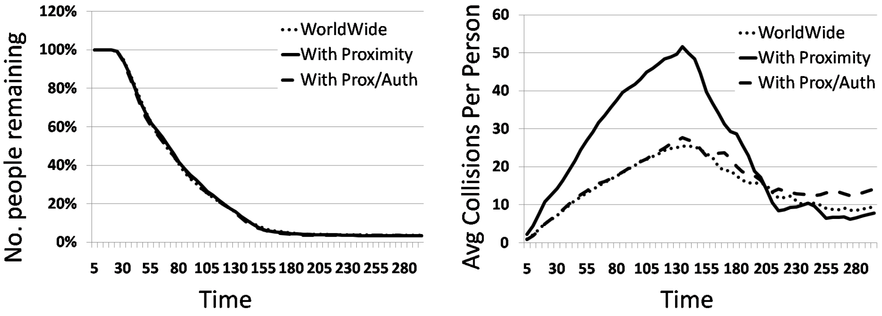

(a) Percent pedestrians unevacuated (b) Avg no. of collisions per person

Fig. 4. Comparison of safety between models

9 Durupinar Experiments

Just as for the VU model, we begin with a sensitivity analysis of the Durupinar model. We then evaluate the implications of the augmentations on the way emotions spread in the simulation. Finally we discuss the implications for safety as they manifest in the ESCAPES simulation.

9.1 Sensitivity Analysis

Sensitivity analysis of the Durupinar model is considerably more complex than the VU model, because although the number of key parameters remain the same, they are interdependent. Lower thresholds, higher dose strengths, or longer dose histories (K) would lead to more agents that are fearful because they would accumulate necessary doses faster. Clearly the relative values are what are important. Thus, we begin with fixed relative values and vary the parameters to identify key sensitivities. In particular, we begin with a baseline of K of 4, dose average of 2, dose standard deviation of 0.5, threshold average of 7, and threshold standard deviation of 2.

Unsurprisingly, altering any one of the parameters' averages OR standard deviations individually drastically alters the magnitude of the contagion effect, but not the overall trends. The exceptions are at extremely low values for K or dose distribution average and at extremely high values for threshold distribution average, when very few agents become fearful at all due to insufficient doses, dose sizes, or extraordinarily high thresholds. Figure 5a shows the percentage of low-fear pedestrians (≤ 0.001) on the y-axis and time steps on the x-axis, with each line representing a different setting of K. Figure 5b shows the same, but with each line showing different settings of the threshold distribution's standard deviation. We use 0.001 instead of 0.1 as before because the decaying aspect of the Durupinar model quickly causes people to fall below 0.1 fear, making 0.001 comparable to 0.1 in the VU experiments. As can be seen, the qualitative trends remain the same over the tested parameter-spaces, with the aforementioned exception. This implies that the model remains robust to parameter changes with respect to the contagion trends that emerge.

We again explored the second-order impacts of parameter variations on the safety of the evacuation by measuring the evacuation rates and average number of collisions of pedestrians in the simulation. As in the VU experiments, we again found no significant variation as the parameters varied across the non-trivial parameter space.

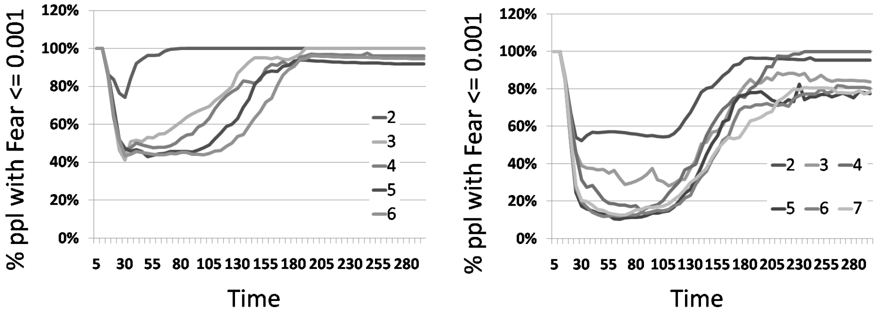

(a) Variations of K

(b) Variations of Threshold SD

Fig. 5. Percent low-fear pedestrians

9.2 Contagion Analysis

Now we examine the effects that the augmentations have on actual contagion in the simulation. Unlike in the VU model, larger populations cause an *increase* in the overall level of fear because of the infection model. Figure 6 shows the percentage of the remaining population with ≤ 0.001 fear and ≥ 0.75 fear on the y-axis and time step on the x-axis. The results use the baseline parameter settings mentioned in Section 9.1, but the tightness of trends is consistent through the non-trivial parameter space.

As expected, in the 'worldwide' model, a larger percentage of the population becomes fearful than in the other two cases, as shown by the lower point reached by the line. This is due to the fact that the entire population can potentially be infected. The model with proximity has a similar dip, although less pronounced since the susceptible population available consists only of neighboring agents. As shown by the less steep increase towards the tail, the number of fearful people tapers off more slowly in the proximity case than in the worldwide case because new susceptible people are encountered over time. The high fear graph, Figure 6b, shows no real surprises with the exception of the abrupt spikes, which is due to the fact that so few people have high fear at any given point in time and fear levels are set immediately to 1.0 upon infection.

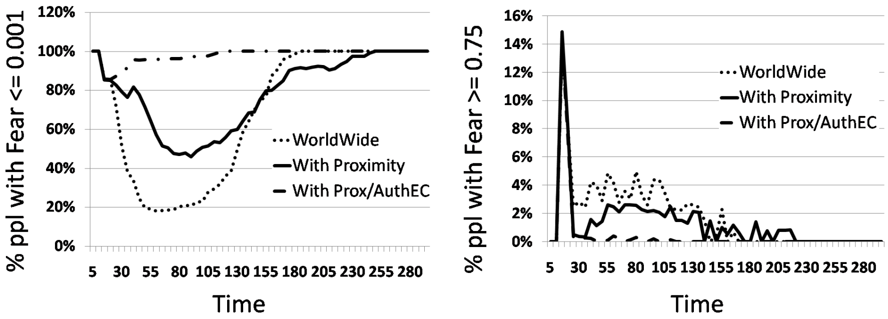

(a) Low fear pedestrians

(b) High fear pedestrians

Fig. 6. Comparison of model with additions

As with the VU model, he effect of proximity on the Durupinar model provides far more realism in the contagion of fear than does the base model. This time we show agents with fear greater than 0.001, but the same dramatic increase in realism remains.

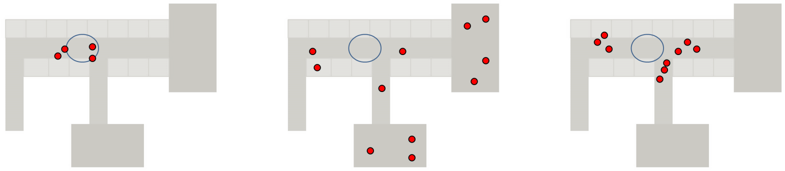

(a) Both models at 20s

(b) Base model at 30s

(e) Proximity at 30s

Fig. 7. Effect of proximity in Durupinar model.

9.3 Safety Analysis

Now we evaluate the impacts of the model augmentations on the actual predictions of evacuations as measured in the ESCAPES system. Again, we show the evacuation rates and average number of collisions of pedestrians in the simulation.

Figure 8a shows the percentage of people that remain unevacuated in the simulation on the y-axis and time step on the x-axis. As can be seen, the evacuation rate remains relatively unchanged. Figure 8b, however, shows very noticeable differences between the models in the number of collisions caused on average. In the 'worldwide' case, as with VU, the number of collisions slopes up substantially faster than the other two models for the same reasons. Next, the model with only proximity follows the same overall trend, but with a lower peak due to the fewer number of fearful people for the majority of the simulation. Finally, the model with both proximity and authority effects shows the lowest peak due to the lower number of infected people in addition to the authority calming effect slowing the pace of pedestrian travel and, therefore, making it easier for agents to avoid collisions with each other. Thus, although the model augmentations do not appear to impact overall evacuation time, the prediction of safety as measured by the number of collisions is again strongly affected.

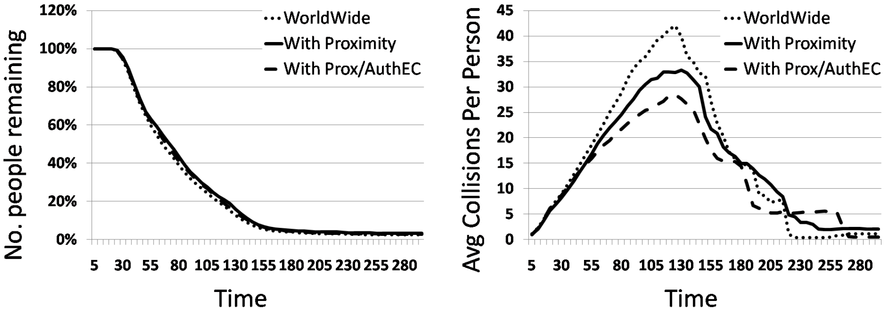

(a) Percent pedestrians unevacuated (b) Avg no. of collisions per person

Fig. 8. Comparison of safety between models

10 Conclusions

Although both Durupinar and VU attempt to model emotional contagion, the underlying mechanisms differ drastically, with the VU model using an independent interaction-based approach and the Durupinar model using a threshold framework inherited from epidemiological studies. In the tests conducted, the VU model seemed to reproduce the contagion phenomenon with higher fidelity.

The Durupinar model possesses a number of inherent flaws due to its origins in epidemiological modeling. Its lack of a representation of 'strength' of the emotion means that agents with more fear have the same impact as agents with only slight fear. Qualitatively, this is inconsistent with observations in social psychology [7]. Furthermore, the Durupinar model possesses no mechanism for ‘reverse’ contagion where a fearful agent might be impacted by the *lack* of fear of other agents. This means that a handful of fearful agents entering a room with 100 0-fear agents will not lose their own fear any faster than if they were alone. In fact, if their fear decays slowly enough and the infection sampling is done quickly enough, they will inevitably infect the entire crowd with their fear. While this *may* occur, the Durupinar model unrealistically implies that duration of exposure even to an extreme minority will *inevitably* lead to escalation.

The VU model, however, is not without its shortcomings. In particular, the lack of a decay function for the emotions means agents will *never* lose fear unless they encounter lower fear agents. The base VU model implemented here also never exhibits escalation because it enforces convergence to the weighted average. However, follow-up work has attempted to address this in [2]. Finally, the proximity and authority figure implementations used here, although an improvement over the base model, are but one of the possible ways that they can be done and further exploration is necessary to determine the most accurate, theoretically-based methods.

References

1. S. G. Barsade and D. E. Gibson. Group Emotion: A View from Top and Bottom. In D. Gruenfeld, E. Mannix, , and M. Neale, editors, *Research on Managing on Groups and Teams*, pages 81–102. JAI Press, 1998.
2. T. Bosse, R. Duell, Z. A. Memon, J. Treur, and C. N. V. D. Wal. A Multi-Agent Model for Emotion Contagion Spirals Integrated within a Supporting Ambient Agent Model. In *PRIMA-09*.
3. T. Bosse, R. Duell, Z. A. Memon, J. Treur, and C. N. V. D. Wal. A Multi-Agent Model for Mutual Absorption of Emotions. In *ECMS-09*, 2009.
4. J. Diamond, M. McVay, and M. W. Zavala. Quick, Safe, Secure: Addressing Human Behavior During Evacuations at LAX. Master’s thesis, UCLA Department of Public Policy, June 2010.
5. W. Doherty. The Emotional Contagion Scale: A Measure of Individual Differences. *Journal of Nonverbal Behavior*, 21(2), 1997.
6. F. Durupinar. *From Audiences to Mobs: Crowd Simulation with Psychological Factors*. PhD dissertation, Bilkent University, Dept. Comp. Eng, 2010.
7. E. Hatfield, J. T. Cacioppo, and R. L. Rapson. *Emotional Contagion*. Cambridge University Press, 1994.
8. M. Hoogendoorn, J. Treur, C. v. d. Wal, and A. v. Wissen. An Agent-Based Model for the Interplay of Information and Emotion in Social Diffusion. In *IAT-10*.
9. L.-O. Lundqvist. Factor Structure of the Greek Version of the Emotional Contagion Scale and its Measurement Invariance Across Gender and Cultural Groups. *Journal of Individual Differences*, 29(3):121–129, 2008.
10. J. Tsai, G. Kaminka, S. Epstein, A. Zilka, I. Rika, X. Wang, A. Ogden, M. Brown, N. Fridman, M. Taylor, E. Bowring, S. Marsella, M. Tambe, and A. Sheel. ESCAPES: Evacuation Simulation with Children, Authorities, Parents, Emotions, and Social Comparison. In *AAMAS-11*.