

Simulation-Based Prediction and Analysis of Collective Emotional States

Charlotte Gerritsen¹✉ and Ward R.J. van Breda²

¹ Netherlands Institute for the Study of Crime and Law Enforcement, De Boelelaan 1077,
1081 HV Amsterdam, The Netherlands

CGerritsen@nscr.nl

² Sentimentics, NeedForward Research, Rotterdam, The Netherlands
ward@sentimentics.com

Abstract. The emergence of aggressive behaviour in crowds is not uncommon and is typically unpredictable and difficult to control. Since the availability of guardians for the prevention of aggression is typically limited, it would be useful if the emergence of unwanted behaviour could somehow be predicted. This paper presents a novel approach to predict the emergence of aggressive behaviour in crowds, based on a combination of sentiment analysis and agent-based simulation. The main idea is to use information from social media (in particular Twitter), analyse the sentiment of the messages and feed the resulting sentiments into a simulation model of social contagion. Based on these simulations, the system should be able to predict when and where unwanted behaviour will occur.

Keywords: Social media · Sentiment analysis · Crime prevention · Prediction · Agent-based modelling · Emotions · Social contagion

1 Introduction

It is April 30th, 2009. Today it is Queen's Day in the Netherlands, a day of national celebration of the Queen's birthday; a day and night in which Dutch people traditionally have fun and party. At one of these parties, in a bar in Amsterdam, a couple of friends are dancing. Suddenly, another visitor bumps into one of the girls and starts making accusations that the girl was in his way. In an attempt to protect the girl, her boyfriend responds that the man should go away. However, this makes the man even more aggressive, and a small riot emerges. The people surrounding the two men start to participate by pushing, name calling and eventually throwing their drinks at the man and his two friends, who have joined the riot. Because of the heated atmosphere and the limited space in the bar, other people start panicking and attempt to flee from the room. As a result, several people get injured by being pushed to the walls. The situation escalates and security has to enter the bar, stop the fight and escort several of the rioting people out of the bar.

Although this is just one example, it illustrates a problem that is often observed in events involving large crowds: the emergence of aggressive behaviour in crowds is often unpredictable and difficult to control. Whether it concerns sports festivities,

demonstrations, or even dance events, in such contexts large-scale aggressive outbursts are not uncommon and are usually triggered by relatively small stimuli, such as a dispute between two people, or people trying to challenge the police [6, 14]. The amount of people involved in such incidents may grow quickly, and the larger the size of the group, the harder it is for the police to control the outburst.

Although significant progress has been made in understanding the mechanisms behind crowd behaviour [6], the (often) large number of people present makes it difficult for security personnel to oversee everything and to intervene in time to prevent unwanted behaviour.

In the Simulation-based Prediction and Analysis of Collective Emotional States (SPACES) project, that problem will be addressed by developing an intelligent system to help guardians predict and control the emergence of unwanted behaviour in crowds. This system makes use of a combination of two main techniques, namely sentiment analysis and agent-based simulation.

Regarding the former, sentiment analysis is a relatively recent method to extract meaning from text messages, and when applied to social media data this can provide an interesting new angle on monitoring the status of large crowds. Nowadays, many hi-tech solutions are already available to monitor crowds, e.g. using infrared cameras and sound detection. Unfortunately, these are both expensive, and (currently) have to be placed at a fixed location. In the current project, we propose the use of data available via social media, which constitutes a cost-effective and flexible solution. Social media are an interesting source of qualitative information and by analysing this information it should be possible to predict potential behaviour. Information posted on social media is available at low cost and analysing such data is a cost-effective means to get insight in what is exactly going on in a crowd, before, during and after large mass events.

Using this input based on sentiment analysis, SPACES will make a prediction of the development of a scenario in the near future. The main technique for making this prediction is agent-based simulation. Specifically, an agent-based simulation model will be developed that contains knowledge about the intra- and interpersonal dynamics of the (mental) states and actions of individuals in a crowd, such as emotion contagion, and (group) decision making. This model will be built upon an existing agent-based model for social contagion within crowds, called ASCRIBE [2, 5]. Based on this model, the system will be able to predict, for example, the emergence of aggressive outbursts, panicking behaviour, or congested areas in parts of a crowd.

Hence, the SPACES project envisions development and testing of a combination of innovative techniques, including dedicated methods to extract relevant information from the dataset in real-time; to match data patterns that are found to psychological states (e.g., fear, stress, anger); and to perform simulation and reasoning about these states to derive predictions about potential future developments. The intelligent system will use this output to provide support for police officers and (formal) guardians in both detecting, decreasing and avoiding (the number of) incidents in large-scale events. The current paper presents an overview of the project and explains our goals and methods to achieve these goals.

The remainder of this paper is structured as follows. Section 2 describes the sentiment analysis method used to extract emotional states from social media feed which provides

the input of the social contagion model. Section 3 provides an overview of the ASCRIBE model for social contagion [2, 5]. Section 4 explains how both models are connected, and illustrates the functioning of the system by means of an example. Section 5 concludes the paper with a discussion.

2 Sentiment Analysis

2.1 Data Processing

As mentioned, the main source of information used within the project consists of messages posted on the Twitter platform [19]. This popular micro-blogging service provides a framework for people to publish and exchange short messages known as *tweets*. The tweets can be sent using a wide array of mobile devices. The tweets have a maximum length of 140 characters, making them fast to publish and to the point, and given the functionalities of the Twitter platform, enable fast distribution of information. Furthermore, it provides services for capturing and analysing tweets in real-time using their API [17]. This makes Twitter an excellent source of input for the project. In a later stage we are interested in using other social media platforms as well (e.g., Facebook [20] or Instagram [18]).

The data processing framework used for the current study filters tweets from the Twitter environment using clusters of keywords related to topic domains that are identified as interesting for analysis. These topic clusters consist of one or more synonyms of the topic under analysis, such as names for an important sporting event, national holiday, or a city name. Also, it is possible to add feature clusters that represent different characteristics of the topic cluster, such as the name of a stadium related to a sporting event, or street names related to a city name.

Given the set multidimensional clusters of topics and related characteristics, the framework stores and derives a number of variables related to the tweets and related to units of time. Some of the important variables that are stored, for each tweet, related to each cluster, are frequency of occurrence, geographical information, sentiment, and keyword related sentiment, and 32 distinct emotions, such as anger, disgust, aggressiveness and disapproval. For more information about the computation and accuracy of the sentiment and the distinct emotions, see [16]. Furthermore, other variables related to the tweets and its senders are stored.

Related to certain units of time, the framework further processes the tweet information. In the current project, the system needs to generate back information about dangerous situations on short notice; the data processing framework uses time units of one minute.

2.2 Data Gathering

Because the data collection process is in real-time, the framework needs to be configured for specific events on forehand. There are two possible approaches for collecting data, namely a top-down and a bottom-up approach.

In the top-down strategy, synonyms for certain high risk events are selected to be contained in topic clusters, such as (national holidays, sport events). Subsequently

feature clusters can be selected containing keywords that represent the physical areas surrounding the events, such as city names, or street names, and related to these clusters, feature clusters representing negative situations, such as violence, intimidation, threat, and injury.

Using the bottom-up strategy, separate topic clusters to contain synonyms for negative situations are configured. Doing this, first negative situations in general are intentionally captured. Subsequently, high activity occurrences are detected, possibly containing identifiers for emergency situations. Although it is unknown if any irregularities will occur for our predefined set of events, using this approach we are still sure to end up with data containing negative situations.

After collecting sufficient amounts of temporal data containing high activity negative situations, we can explore which variables are possible predictors, and whether our agent model can accurately identify emergency situations using them. Here, the effect of the mathematical relations and corresponding weights on the performance of the predictions will be investigated. Also, the objectives of the model will be identified, i.e. the most important variables that it must predict. To further investigate the performance potential of the model, parameter tuning strategies will be used to optimize the agent model.

The output of the sentiment analysis model will be provided as input to the ASCRIBE model [2, 5]. This model will be explained in the next section.

3 The ASCRIBE Model for Social Contagion

The proposed model to predict the development of the crowd's mental states over time is an extension of the ASCRIBE model [2, 5]. The ASCRIBE model (which stands for Agent-based Social Contagion Regarding Intentions, Beliefs and Emotions) describes how mental states can spread over large crowds based on social contagion processes [2, 5]. The model is rooted in theories from the literature in neuroscience, including Damasio's Somatic Marker hypothesis [3] and the concept of mirror neurons [10].

The main ingredients of ASCRIBE are three types of mental states, namely *beliefs*, *intentions* and *emotions*. Beliefs refer to certain information about the world (e.g., an agent may have the belief that a fight is going on). Intentions may refer to planned actions (e.g., an agent may have the intention to join the fight or to escape). Emotions may refer to several types of emotional states (e.g., panic or aggression). All of these mental states are represented in terms of real-numbered variables in the domain $[0,1]$: for instance, an agent with an 'aggression state' of 0.1 is rather calm, whereas an agent with an 'aggression state' of 0.9 is highly aggressive. Next, all mental states of agents can be influenced by intra-agent processes (e.g., an agent's level of panic influences its beliefs) as well as inter-agent processes (e.g., agent A's panic level influences agent B's panic level). These influences are represented in terms of differential equations. Details are given in [2, 5], but the general idea is that the extent to which a mental state influences another mental state (of the same or another agent) depends on the current values of the mental states, as well as several additional parameters, such as the distance between the agent that transmits the mental state and the agent that receives it.

In [2], the ASCRIBE model was extended with mechanisms to determine individual agent movements, hence resulting in a crowd simulation model that goes beyond traditional approaches, in the sense that it contains sophisticated mechanisms to simulate the mental states of crowd members. Additionally, in [2] the model was applied to simulate a real world incident that took place in Amsterdam in 2010, and it was shown to outperform the influential model by Helbing and colleagues [4] due to its ability to address scenarios in which the mental processes of individual agents are a key determinant of the crowd's behaviour (for instance, a scenario where some people stop panicking as soon as they understand that the situation is not threatening anymore). Also in another study by different authors [11], ASCRIBE has been shown to perform equal and mostly superior to other models in the context of a protest scenario in Greece. Finally, ASCRIBE has been applied successfully to simulate various other cases in the emergency domain, including the London Bombings in 2005 [2].

The ASCRIBE model will form the basis of our simulation model. The input for the model will be provided by messages posted on Twitter [19], as mentioned in Sect. 2. In the next section the connection between both models will be explained.

4 Connection Between Models

In this section, the intended functionality of the combined model will be illustrated by means of a motivational example. Imagine a crowded pop concert, which takes place in a squared area¹ with a stage, a bar and some toilets (see Fig. 1a). For the purpose of the analysis, a grid structure is imposed onto the area, dividing it into a number of cells of equal size (e.g., 5×5 m); see Fig. 1b.

Once the pop concert starts, the sentiment analysis module described in Sect. 2 becomes active. By extracting sentiment out of the runtime Twitter data, and connecting this to the geographical information extracted from the same data, the module is able to create a two-dimensional map of particular states of the crowd. As an example, suppose we are interested in the distribution of the state of 'aggression' over the crowd during the pop concert. Then, the sentiment analysis module will focus on keywords related to aggressive behaviour (e.g., 'fight', 'beat', 'kick', etc.). Based on such information, the module assigns a 'level of aggressiveness' to each cell in the area. For instance, assume at a certain point in time a small riot emerges near the restrooms, and as a result bystanders start sending tweets about this, like "things are getting out of hand at the #toilets people getting mad! #pinkpop". By processing such tweets, the sentiment analysis module concludes that the cells in the vicinity of the restrooms have a relatively high level of aggressiveness (e.g., 0.8 on a scale from 0 to 1), whereas other cells relate to parts of the crowd with lower levels of aggressiveness.² As a result of this analysis, each cell can be assigned a

¹ For the sake of the example, it is assumed that the area consists of a single room with a squared shape. However, in principle the approach can be applied to any type of area (e.g., sports stadiums, streets, or even entire cities).

² In case no information about a particular cell is available, by default the average of the surrounding cells can be used.

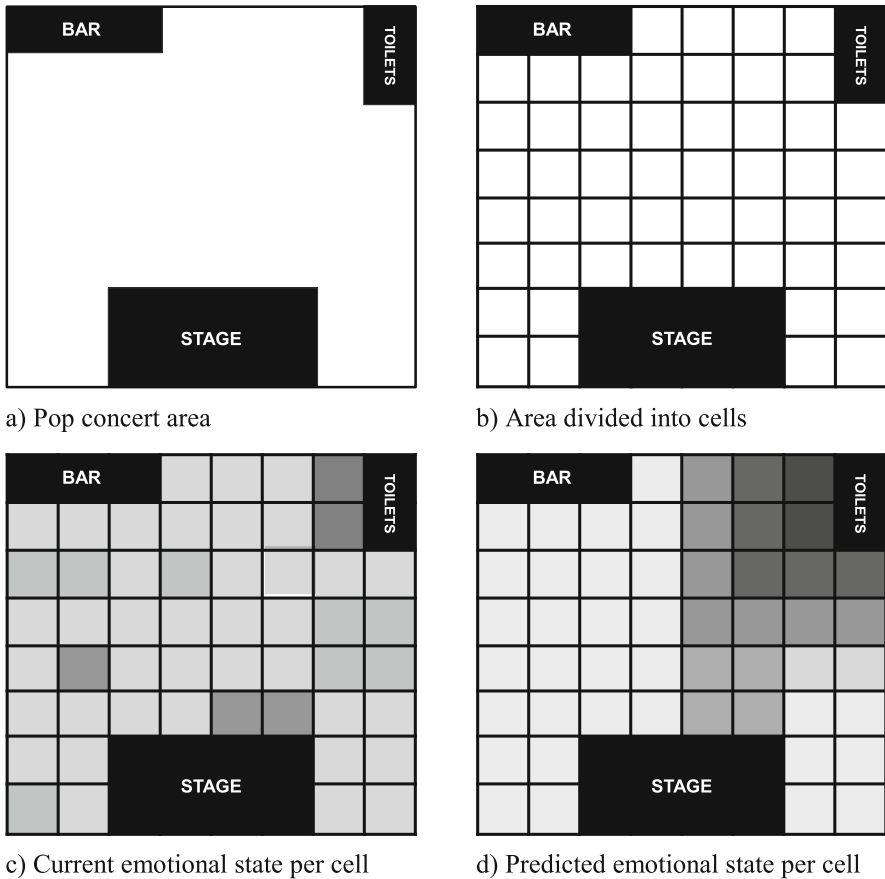


Fig. 1. Illustration of the dynamics of the combined model for an area at a pop concert

numerical value that estimates its level of aggression (or any mental state, as will be explained below). This ‘landscape of aggression’ can be visualized as shown in Fig. 1c (where the darkness of the cells is proportional to the level of aggressiveness).

Next, the output of the sentiment analysis module is passed to the ASCRIBE model for social contagion, which makes a prediction of the dynamics of the mental state of the crowd in the near future. For instance, based on the situation depicted in Fig. 1c, ASCRIBE may predict that the initial riot that was detected next to the restrooms is likely to expand, thereby also increasing the state of aggressiveness of the adjacent cells (e.g., because the people in these cell are susceptible for aggression). On the other hand, the state of aggressiveness in other cells may decrease (e.g., because people are moving away from the riot). Hence, the social contagion model is able to calculate the dynamics of the mental states of the crowd over time, and as a result a similar picture can be created of the predicted situation in the near future (say, after 5 min), as shown in Fig. 1d. In this example, the riot has clearly expanded to a larger area in the neighbourhood of the restrooms and the stage.

To conclude, note that, although the current example focuses on the mental state of aggressiveness only, in reality a similar analysis can be made for a variety of mental states. In particular, ASCRIBE distinguishes between three categories of mental states, namely beliefs, intentions and emotions. Aggression can be considered one type of emotion, but when it comes to riots, other emotion types may play a role as well, in particular the emotion of fear (or panic). Additionally, several beliefs may influence the behaviour of the crowd as well. For instance, the belief ‘that a riot is taking place near the restrooms’ may trigger certain people to avoid that area, whereas it triggers others to approach it. Similarly, it may be valuable to have information about people’s intentions (e.g., to party, to fight, to flee, etc.). Hence, with the help of ASCRIBE, ultimately a number of ‘mental state maps’ as shown in Fig. 1 can be produced. One step further, ASCRIBE not only predicts the dynamics of these states in isolation, but also takes the interplay between these states into account. For instance, the belief that a riot is taking place may trigger someone’s emotion of fear, which in turn may lead to an increased intention to escape. Eventually, ASCRIBE is also able to predict actual actions of the people in question, which can be used to predict the dynamics of the density of the crowd. For example, if many people intend to move away from the riot, certain areas may become congested (see [2]) for a more elaborated illustration).

5 Discussion

In this paper the main ideas were put forward of the recently started project called SPACES. The main goal of our work is to make events involving large crowds safer, by proposing a method to decrease the number of riots while using a limited amount of (formal) guardians. To achieve this, we will integrate the existing ASCRIBE model for social contagion [2] with a sentiment analysis module. Based on the outcomes of the sentiment analysis we are able to perform agent based simulations regarding the spread of mental states, like beliefs, intentions and emotions. This gives us predictions of the development of the mental state of a crowd over time and space, and serves as an input to guide the behaviour of guardians. For instance, guardians can use this information to determine where they should go to de-escalate situations, or to decide which areas are in need for other types of support (e.g., providing directions on where to go in case of congestion [1] or alterations in light or music to calm people down [15]).

This is not the first project regarding this societal problem [8, 9, 12]. Li et al. [7] monitored emotions in a crowded area by using an app. The visitors had to rate their emotion at fixed time periods by using the EmoApp [7]. Their findings revealed that both movement and emotions were consistent with the activities at the festival indicating that the users answered truthfully. The main difference between this work and our project is the way of collecting information about the mental state of the crowd. While Li et al. [7] asked specifically for the level of emotion we are able to unobtrusively extract emotions (as well as other mental states) from social media. In Ahuja et al. [1] the authors present a different way of monitoring the crowd. They propose the use of drones when certain areas are overcrowded or congested. These drones are able to give suggestions to the visitors, to escape the crowd. The main goal of Ahuja et al. [1] is somewhat

different from the goal in the SPACES project but it is an interesting approach of controlling the crowd. Both Li et al. [7] and Ahuja et al. [1] visualize the data at run time but they do not combine the data on emotions with a predictive model as we do with the ASCRIBE model [2]. The work by Van Dyke Parunak et al. [13] comes closest to the SPACES project. They make predictions of crowd behaviour using Twitter data as well, and also combine this information with a predictive model. The main difference is that we provide a more detailed analysis of the different mental states of the crowd, and their interactions.

The proposed project is still in an early stage. Although the conceptual framework has been developed, much of the actual implementation work remains to be done. As a first step we will refine the sentiment analysis and the ASCRIBE model, by extending them and tuning them to the needs of our project. Once that has been established the connection between both models will be implemented and the system will be tested using simple case studies. The last phase of the project consists of an evaluation of the system based on a real life event, and the development of support methods.

Acknowledgments. The authors are grateful to Tibor Bosse for fruitful discussions about the project.

References

1. Ahuja, G., Karlapalem, K.: Managing multi robotic agents to avoid congestion and stampedes. In: Elkind, B., Yolum, W., (eds.), Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), ACM Press (2015)
2. Bosse, T., Hoogendoorn, M., Klein, M.C.A., Treur, J., van der Wal, C.N., van Wissen, A.: Modelling collective decision making in groups and crowds: integrating social contagion and interacting emotions, beliefs and intentions. *Auton. Agent. Multi-Agent Syst. J.* **27**, 52–84 (2012)
3. Damasio, A.R., Everitt, B.J., Bishop, D.: The somatic marker hypothesis and the possible functions of the prefrontal cortex. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* **351**(1346), 1413–1420 (1996)
4. Helbing, D., Farkas, I., Vicsek, T.: Simulating dynamical features of escape panic. *Nature* **407**(6803), 487–490 (2000)
5. Hoogendoorn, M., Treur, J., van der Wal, C.N., van Wissen, A.: Modelling the interplay of emotions, beliefs and intentions within collective decision making based on insights from social neuroscience. In: Wong, K.K.W., Mendis, B.S.U., Bouzerdoun, A. (eds.) *ICONIP 2010, Part I. LNCS*, vol. 6443, pp. 196–206. Springer, Heidelberg (2010)
6. Knutsson, J., Madensen, T.D.: Preventing Crowd Violence. Lynne Rienner Pub., Boulder (2011)
7. Li, J., Erkin, Z., de.Ridder, H., Vermeeren, A.: A field study on real-time self-reported emotions in crowds. In: proceedings of ICT OPEN, eindhoven, The Netherlands, pp. 80–84 (2013)
8. Mitleton-Kelly, E., Deschenaux, I., Maag, C., Fullerton, M., Celikkaya, N.: Enhancing crowd evacuation and traffic management through ami technologies: a review of the literature. In: Mitleton-Kelly, E. (ed.) *Co-evolution of Intelligent Socio-Technical Systems*, pp. 19–41. Springer, Heidelberg (2013)

9. Nagy, A. Stamberger, J.: Crowd sentiment detection during disasters and crises. In: Rothkrantz, L., Ristvej, J., Franco, Z., (eds.), Proceedings of the 9th International ISCRAM Conference – Vancouver, Canada, April 2012
10. Rizzolatti, G., Sinigaglia, C., Anderson, F.T.: *Mirrors in the Brain How our Minds Share Actions and Emotions*. Oxford University Press, Oxford (2008)
11. Tsai, J., Bowring, E., Marsella, S., Tambe, M.: Empirical evaluation of computational emotional contagion models. In: Vilhjálmsson, H.H., Kopp, S., Marsella, S., Thórisson, K.R. (eds.) IVA 2011. LNCS, vol. 6895, pp. 384–397. Springer, Heidelberg (2011)
12. Tsai, J., Fridman, N., Bowring, E., Brown, M., Epstein, S., Kaminka, G., Marsella, S., Ogden, A., Rika, I., Sheel, A., Taylor, M.E., Wang, X., Zilka, A., Tambe, M.: ESCAPES—evacuation simulation with children, authorities, parents, emotions, and social comparison. In: Tumer, K., Yolum, P., Sonenberg, L., Stone, P., (eds.), Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011). Valencia: Innovative Applications Track. (2011)
13. Parunak, H.V.D., Brooks, S.H., Brueckner, S., Gupta, R.: Dynamically tracking the real world in an agent-based model. In: Alam, S.J., Parunak, H.V.D. (eds.) MABS 2013. LNAI, vol. 8235, pp. 3–16. Springer Verlag, Heidelberg (2014)
14. http://nl.wikipedia.org/wiki/Strandrellen_in_Hoek_van_Holland
15. <http://www.de-escalate.nl>
16. <http://www.sentimentics.com>
17. <https://dev.twitter.com>
18. <https://instagram.com>
19. <https://twitter.com>
20. <https://www.facebook.com>