

Using Qualitative Reasoning for Social Simulation of Crowds

NATALIE FRIDMAN and GAL A. KAMINKA, The MAVERICK Group and Bar Ilan University

The ability to model and reason about the potential violence level of a demonstration is important to the police decision making process. Unfortunately, existing knowledge regarding demonstrations is composed of partial qualitative descriptions without complete and precise numerical information. In this article we describe a first attempt to use qualitative reasoning techniques to model demonstrations. To our knowledge, such techniques have never been applied to modeling and reasoning regarding crowd behaviors, nor in particular demonstrations. We develop qualitative models consistent with the partial, qualitative social science literature, allowing us to model the interactions between different factors that influence violence in demonstrations. We then utilize qualitative simulation to predict the potential eruption of violence, at various levels, based on a description of the demographics, environmental settings, and police responses. We incrementally present and compare three such qualitative models. The results show that while two of these models fail to predict the outcomes of real-world events reported and analyzed in the literature, one model provides good results. We also examine whether a popular machine learning algorithm (decision tree learning) can be used. While the results show that the decision trees provide improved predictions, we show that the QR models can be more sensitive to changes, and can account for *what-if* scenarios, in contrast to decision trees. Moreover, we introduce a novel analysis algorithm that analyzes the QR simulations, to automatically determine the factors that are most important in influencing the outcome in specific real-world demonstrations. We show that the algorithm identifies factors that correspond to experts' analysis of these events.

Categories and Subject Descriptors: I.2.0 [Artificial Intelligence]: General—*Cognitive simulation*

General Terms: Human Factors

Additional Key Words and Phrases: Demonstrations, social simulation, qualitative reasoning

ACM Reference Format:

Fridman, N. and Kaminka, G. A. 2013. Using qualitative reasoning for social simulation of crowds. *ACM Trans. Intell. Syst. Technol.* 4, 3, Article 54 (June 2013), 21 pages.
DOI: <http://dx.doi.org/10.1145/2483669.2483687>

1. INTRODUCTION

A violent demonstration with many casualties among its participants, police forces, and innocent bystanders is unfortunately not a rare phenomenon. Can we know the potential nature of an event before it ends with tragic outcomes? Accurate predictions regarding the potential violence level of a demonstration can improve the police decision making process, and decrease the number of casualties.

In general, there are several ways we can improve the police decision making process: usage of expert consultants, numeric simulations (e.g., Patrick et al. [1999]), and agent-based simulations [Jager et al. 2001]. Usage of expert psychologists and sociologists

The authors thank IMOD and ISF grant 1357/07 for partial support of this research.

Authors' addresses: N. Fridman (corresponding authors), The MAVERICK Group, Computer Science Department, Bar Ilan University, Israel; email: natalie.fridman@gmail.com; G. A. Kaminka, The MAVERICK Group, Computer Science Department and Gonda Multidisciplinary Brain Research Center, Bar Ilan University, Israel.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2013 ACM 2157-6904/2013/06-ART54 \$15.00

DOI: <http://dx.doi.org/10.1145/2483669.2483687>

for consultations is a very common approach. However, such usage before every demonstration is not a practical and affordable solution. Usage of numerical simulation (e.g., Patrick et al. [1999]) requires full information regarding the domain, which in many cases (such as reasoning regarding the violence level in demonstrations) simply does not exist. Agent-based simulation approaches (e.g., Jager et al. [2001]) require modeling at micro- (individual) level which in many cases may be impractical. Indeed, in general, one of the biggest challenges in the crowd modeling field is the lack of precise data. Such precise data is even difficult to obtain since crowd experiments are complex and costly.

This is not to say that social sciences do not have deep and extensive knowledge of demonstrations. On the contrary, there exists significant literature on the factors that impact violence during demonstrations. However, the literature mostly reports on partial, macrolevel qualitative descriptions of the influencing factors. Integrating these together is a formidable challenge that requires novel forms of modeling.

To enable modeling and reasoning with imprecise and partial information we propose using qualitative reasoning techniques. *Qualitative Reasoning* (QR) [Kuipers 1989; Forbus 1996] is a subarea of artificial intelligence that enables modeling and reasoning with partial or imprecise numeric information. The QR approach shows that while having a qualitative representation of data and order values (such as little/medium/large) is difficult, it is still possible to draw useful conclusions.

In this article we describe a novel application of QR to modeling and reasoning regarding potential violence levels during demonstrations. We develop qualitative models consistent with the imprecise descriptions in social science literature, allowing us to model the interactions between different factors that influence violence in demonstrations. We develop three separate models, incrementally increasing in complexity and in the number of factors they consider. These three models are evaluated on real-world scenarios, using news reports and Wikipedia entries as the source of information as to the values of different quantities.

By using a simple technique that considers the number of paths leading to different violence outcomes, we are able to provide an estimate of the likelihood of different outcomes, for each test case. We contrast the predictions of the different models and thus demonstrate one important benefit of using QR for social simulation modeling, namely, the ability to easily test social science theories on real-world data.

Moreover, we examine whether decision trees, a popular machine learning approach, can be used for qualitative predictions. While the results show that decision trees provide accurate predictions (slightly better than our QR models) they lack the ability to support hypothetical “what-if” reasoning, because they do not have the explanatory power of a social science model. Thus, we claim that using QR is better for reasoning in this task.

Finally, we develop an algorithm that analyzes the qualitative simulation graph of each test case, to determine the factors that are most important in influencing the outcomes of *the specific case* under consideration. The key to this algorithm is to determine simulation graph nodes with high outcome entropy, that is, nodes which lead to different outcomes, at fairly uniform likelihoods. In the states corresponding to such nodes, it is possible to identify actionable factors that can be used to influence the outcomes. We show that for real-world cases, the algorithm results in identifying causes also identified by experts.

2. BACKGROUND AND RELATED WORK

Usage of computer simulation is considered to be a promising approach for modeling and reasoning regarding different social phenomena [Gilbert and Troitzsch 2005]. There are several micro- (individual) and macro- (group) level techniques that enable

such modeling, for example, usage of agent-based simulation, cellular automata, and system dynamics. This article introduces a novel technique for macrolevel modeling.

Agent-based simulation is a microlevel approach where by social behaviors are simulated by simulating each individual, and his/her interactions. By modeling agents' social cognition, we have the ability to model complicated social interactions. Such simulations have been successfully used in modeling some crowd behaviors [Fridman and Kaminka 2007; Jager et al. 2001], economic phenomena [Wander 2000], and more. However, it is a bottom-up approach in the sense that to receive a macrolevel behavior we must model the microlevel interactions which necessitates detailed individual modeling, and when number of agents is scaled up it may provide significant computational barriers. Furthermore, there are domains in which such resolution may be too high and even unnecessary.

System dynamics [Gilbert and Troitzsch 2005] is a macrolevel approach in the sense that it models an entire system. It uses defined stocks, flows, and feedback loops to model system behavior. The models are basically sets of differential equations that describe changes in the system. For crowds and demonstrations, such accurate and complete mathematical formulation is not available.

There are also techniques that do not require a model to be given as an input. Instead, machine learning techniques can facilitate reasoning regarding social phenomena by inducing a model out of examples of the target concepts. For instance, a decision tree learning algorithm [Mitchell 1997] takes as an input a set of examples and the target classes to which they belong (for instance, examples of properties of demonstrations and their associated violence level) and induces (builds) a model, in the form of a decision tree. This decision tree allows classification of the observed data according to the given properties. However, as we show in this article, prediction (classification) accuracy is not the only requirement for policy decision support. In particular, we show that QR models can be much better at analyzing hypothetical settings.

Qualitative Reasoning (QR) [Kuipers 1989; Forbus 1996] is a macrolevel approach to modeling and reasoning with partial and imprecise numeric information. The traditional usage of QR is in modeling common-sense reasoning, for example, in physics [Forbus 1996]. Indeed, for several decades there have been extensive studies of QR techniques in physics to enable reasoning about physical systems. However, while QR is usually associated with physics, other branches of science such as ecology [Salles and Bredeweg 2006], social science [Kamps and Péli 1995], politics [Forbus and Kuehne 2005], etc., are also starting to adopt this approach. Even though social science is much less formal than physics, QR approaches have been found just as powerful.

For example, Kamps and Péli [1995] present an application of QR to social science. They show that the QR approach is successfully applicable for modeling and reasoning of density dependence theory of organizational ecology. They simulate the growth pattern of a population and enable prediction regarding the population size.

Salles and Bredeweg [2006] use QR techniques in ecological modeling. Based on theoretical and common-sense knowledge, they build a qualitative model of population and community dynamics in the Brazilian Cerrado vegetation. They show that QR models capture successfully ecological knowledge and enable valid prediction of different behaviors in ecological systems.

Brajnic and Lines [1998] apply a QR technique in complex, socio-economic allocation problems, in particular the allocation of national income between several important resources such as capital investment, social services, etc. They show that with relatively weak quantitative information about functional relationships, useful prediction regarding the behavior of economic society can be drawn.

Fuzzy Cognitive Maps (FCM) [Kosko 1986; Glykas 2010] is a different macrolevel inference technique that enables causal reasoning using fuzzy directed graphs. Similarly

to QR, FCM enables imprecise and qualitative representation of the model. However, while FCM graph structure allows causal propagation (similar to QR), the two techniques differ in output: A QR simulation graph represents the set of all possible values that the state variables might converge to, while FCMs, when propagated, converge to a single resulting value for each variable.

Our goal is to model and predict the violence level during demonstrations. There are several theories in social science regarding the factors influencing violence level. These factors are described partially and qualitatively, at macrolevel without full and precise information. Moreover, we are not aware of any simple model that attempts to be comprehensive. Thus, we use a qualitative reasoning approach for integrating together these partial models. To the best of our knowledge, QR techniques have never been used for modeling crowd behavior phenomena. We present here a detailed description of such modeling and compare three different models on different real-life scenarios and also for what-if reasoning.

3. QUALITATIVE REASONING AND SIMULATION

A comprehensive discussion of QR is beyond the scope of this article. We provide a brief description here, and refer the reader to Kuipers [1989] and Forbus [1996] for additional details.

Qualitative simulation enables reasoning about possible system behaviors that can emerge from an initial world state. A qualitative model is like any other in the sense that it has variables which hold values, and there are rules or relations between the variables that govern how a given value of a variable influences another variable. Once an initial state (value assignment) is given, the model can be used to determine all possible values that can result from valid inference (i.e., given the known relations between variables).

The key idea in qualitative reasoning and simulation, however, is that the values of variables and the rules that govern their evolution are given in descriptive, qualitative terms, rather than numerical. A qualitative model is made from a set of *quantities* which are variables. Quantities are the lowest-resolution representation for continuous parameters and are each composed of a pair of values: *magnitude* and *derivative*. The magnitude represents the amount of quantity and derivative represents the direction of change. The set of possible values is described by Quantity Space (QS) which is a finite and ordered set of qualitative landmark values. Landmark values represent important real-number values.

For example, suppose we want to represent the amount of water in a pipe. The *magnitude* represents the amount of water and we may define its Quantity Set (QS) with the following landmarks: $QS = \text{zero, plus, max}$, which represent *no water in the pipe*, *some water in the pipe*, and *water in the pipe to full capacity*. The *derivative* represents the direction of change and it has the following quantity set: $QS = \text{min, zero, plus}$ which represents three possible directions accordingly: decreasing, steady, and increasing. If there is a state with the following quantity $\langle \text{plus, plus} \rangle$ it means that in the current state, there is some amount of water in the pipe and it is increasing.

In addition to quantities, a qualitative model is made from a set of causal relationships, that relate quantities to each other, given both their magnitude as well as their derivative. There are two types of casual relationship between quantities: *direct influence* and *indirect influence*. The direct influence, denoted I , signifies that A causes B ; if B happens then A might be the reason for it. The indirect influence, denoted P , signifies that A causes B only if A causes C and C causes B . A quantity that is not influenced by any process is considered to be a constant. Each influence may be positive ($I+$, $P+$) or negative ($I-$, $P-$), meaning the derivative of the target quantity increases or decreases accordingly.

A qualitative simulation takes as input the initial state of the world (i.e., a value assignment for the quantities) and qualitative descriptions of processes that change the state. It produces a state transition graph. Each state is a possible unique behavior that the model develops, containing a unique set of values and inequality statements (quantities) that describe the current behavior of the system.

To construct the state graph, qualitative simulation uses the partial models to drive changes in the quantity magnitudes and derivatives, until they no longer change. Each path of transitions, representing application of the models to the quantities, represents a single possibility of how the state may develop from the initial state until it converges to a final state. The collection of all such paths, leading from the initial state to a terminal state, is the entire set of possible system behaviors.

Thus a state graph captures the set of all possible behaviors that the model may manifest. It consists of a set of states and the transitions between them (state transitions). State transitions transform one state into another, by specifying the changes in values and in inequality statements. Each state may be the origin for multiple transitions which lead to multiple possible developments of the state. A sequence of states connected by state transitions, where each state is the immediate successor of the one before, is called a behavior path.

In each cycle and on each quantity, all influences (direct and indirect) are combined. When positive and negative influences are combined, ambiguities may occur. The qualitative simulation considers all the possible combinations, thus, when qualitative description is incomplete, it provides a nondeterministic prediction.

QR is a well-established technique in artificial intelligence and computer science. There exist free, user-friendly software packages for specifying qualitative models and for reasoning with them. Indeed, in most of the experiments we report on in this article, we utilized an available free software package called GARP3 [Bredeweg et al. 2009], that allows for visual entry and running of models.

4. QUALITATIVE SIMULATION OF DEMONSTRATIONS

QR modeling techniques are a good match for the current state of knowledge in social sciences regarding demonstrations. Existing theories are inaccurate or incomplete. There are many microtheories regarding the influencing factors on the violence level: Each such theory focuses on a small subset of factors. Integrating all of them into a single unified model is a real challenge.

The Israeli police initiated a comprehensive study to address this challenge, resulting in a report [Carmeli and Ravid-Yamin 2006] that provides a collection of factors and their influence on the violence level and also on each other. Their goal was to classify and analyze different kinds of demonstrations in order to propose appropriate methods for the police force in dealing with the mass. They studied 102 crowd events (in particular demonstrations) during the years 2000–2003 and interviews with 87 policemen and police officers. They analyzed a variety of factors that may affect violent behavior, as well as relevant literature. This report is thus a comprehensive, yet informal and imprecise, collection of factors that affect demonstrations.

The first contribution we make is in translating the textual results in the report into an executable model, using QR. Indeed, we developed three separate models, incrementally increasing in complexity and size, of the different components influencing violence in demonstrations. These are described next.

4.1. The Base Model

The first (*base*) model was developed based on the report's literature review [Carmeli and Ravid-Yamin 2006] (see Figure 2). It was proposed there as a first attempt at building a baseline, purely based on literature review. According to the base model

Participants:
Population

Quantities:
 Hostility for police: QS = {high, low, zero}
 Personal price: QS = {high, low, zero}
 Punishment: QS = {high, low, zero}
 Group cohesiveness: QS = {high, low, zero}
 History of violence: QS = {yes, no}
 Violence: QS = {high, low, zero}

Influences:
 I+(Hostility for police{high}, Violence)
 I-(Hostility for police{low}, Violence)
 I+(Personal price{high}, Violence)
 I-(Personal price{low}, Violence)
 I-(Punishment{high}, Violence)
 I+(Punishment{low}, Violence)
 I+(Group cohesiveness{high}, Violence)
 I-(Group cohesiveness {low}, Violence)
 I+(History of violence{yes}, Violence)

Fig. 1. Base model: description.

the most influential factors on the violence level during a demonstration are: (1) the crowd's a priori hostility towards the police; (2) willingness to pay personal price (such as willingness to be arrested); (3) the chance of punishment for violent actions (e.g., a belief that police will or will not respond strongly); (4) group cohesiveness; (5) previous history of violence. All of these directly increase the level of violence. However, not all have an opposite effect when reversed. For instance, the existence of previous history of violence among the specific group of demonstrators increases the potential violence level, but lack of such history does not decrease the violence level (i.e., has no effect).

Figure 1 presents a model description and Figure 2 presents the model's graphical structure. We define one entity (called population) with six quantities, five of them based on the presented theoretical description and another one the violence level which is the outcome. For each quantity we defined a Quantity Space (QS) and defined direct Influence (I) between them and the violence level. $I+$ represents the derivative of the target quantity (violence level) increases and $I-$ means it decreases. For example, high hostility for police increases the violence level while low hostility for police decreases it. There are also quantities with one directional influence such as previous history where existence of previous history of violence increases the violence level while lack of history does not decrease the violence level and it actually has no effect on it.

4.2. The Police Model

The *police* model, described by Karmeli and Ravid-Yamin [2006] (Figure 4), significantly expanded the base model, based on interviews with police officers and their investigation into 102 demonstrations. In addition to the factors from the base model, the police model adds more variables that can be divided into several groups: *environment variables* (such as time, weather, and space characteristics where the event occurs), *participants variables* (sociological and behavior characteristics of the participants such as

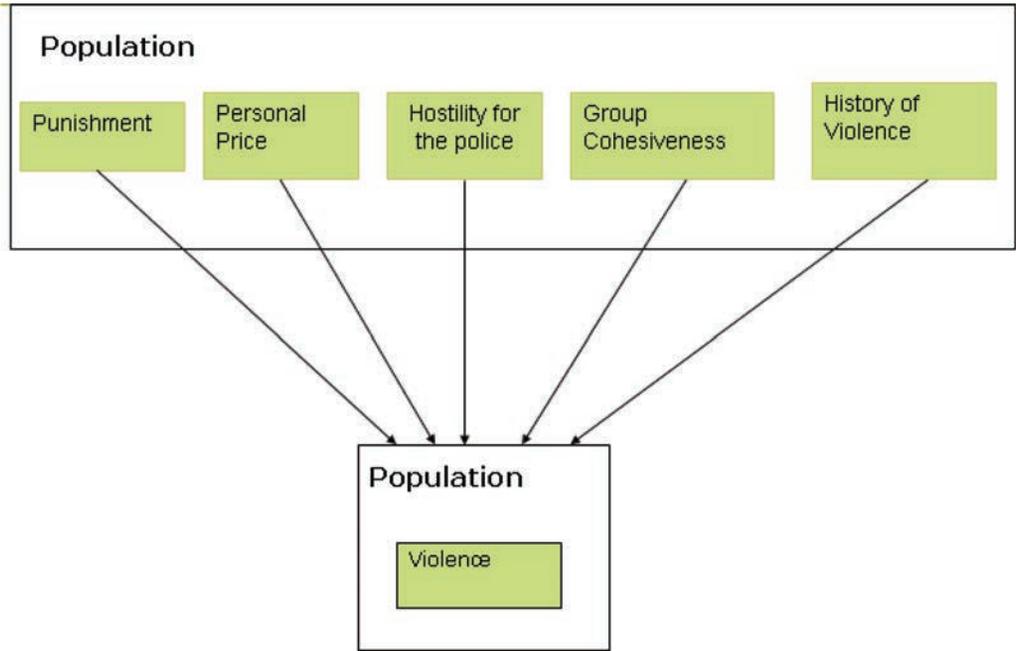


Fig. 2. Base model: structure.

number of participants, existence of speaker or leader, etc.), *police variables* (referring to behavior and organizational characteristics of the police such as intervention time and intervention strength), *procedural variables* (referring to dynamic characteristics of the event), and *outcome variable* (referring to the outcome of the event).

The additional variables to the base model that were also found to be influencing the demonstration outcome are: (1) number of participants, (2) the existence of group representative (such as group leader or speaker), (3) existence of a demonstration license (permit; i.e., whether the demonstration is legal), (4) existence of violent core among the participants, (5) participants' united identity (such as racial minority, social groups, etc.), (6) event purpose (such as emotional event or rational), (7) police time intervention, (8) police intervention strength, (9) weather, (10) time of the day (such as morning, night, etc.), (11) demonstration location sensitivity (a high-sensitivity place, such as a mosque or synagogue, or a low-sensitivity place, such as a city square or piazza), and (12) time of year sensitivity (e.g., Christmas). The research results also showed significant relations between these variables and also their impact on the event outcome (the violence level). For example, political or social demonstrations usually end with low level of violence (usually without casualties). However, demonstrations with nationalistic flavor that intend to express emotions (letting off steam) are characterized by much more violent outcomes. It was also found that the time of the day has impact on the violence level, such as the fact that more violent demonstrations occur at night than during the day.

Figure 3 presents a model description and Figure 4 presents the model's graphical structure. We defined three entities (population, nature, and police) and 18 quantities, where 6 are similar to the base model and 12 are additional ones. Moreover, based on the research conclusions we defined Influence (I) among different variables. As in the previous model, $I+$ represents that the derivative of the target quantity (such as

Participants:
 Population
 Nature
 Police

Quantities:
 Population:
...Similar to Base model...
 License: QS = {have, don't have}
 Group speaker: QS = {exist, not exist}
 Number of participants: QS = {0-1000/1000-10000/10000-100000}<
 United identity: QS = {racial minority/religious/ethnicity/social&political gr.}<
 Violent core: QS = {exist, not exist}<
 Event purpose: QS = {emotional, rational}<

Nature:
 Weather: QS = {extreme, medium}<
 Time: QS = {morning, noon, evening/night}<
 Place sensitivity: QS = {high, low}<
 Time sensitivity: QS = {high, medium, low}<

Police:
 Time intervention: QS = {premature/late, on-time}<
 Intervention strength: QS = {low, medium, high, exact}<

Influences:
...Similar to Base model...
 I+(License{don't have}, Violence)
 I-(License{have}, Violence)
 I-(Group speaker{exist}, Violence)
 I+(Number participants{100-1000/1000-10000/10000-100000}<, Violence)
 I-(Number participants {0-100/10000-100000}<, Violence)
 I+(United Identity {racial minority/religious/ethnicity/social&political gr.}<, Violence)
 I-(United Identity {racial minority/religious/ethnicity/social&political gr.}<, Violence)
 I-(United Identity {social&political gr.}<, Violence)
 I+(United Identity {social&political gr.}<, Violence)
 I+(Violent core{exist}, Violence)
 I-(Violent core{not exist}, Violence)
 I+(Event type{emotional}, Violent core)
 I-(Event type{rational}, Violent core)
 I+(Weather{medium}, Violence)
 I-(Weather{extreme}, Violence)
 I+(Time{evening/night}, Violence)
 I-(Time{morning/noon}, Violence)
 I+(Place sensitivity{high}, Event purpose)
 I-(Place sensitivity{low}, Event purpose)
 I+(Time sensitivity{high/medium}, Violence)
 I-(Time sensitivity{low}, Violence)
 I+(Time intervention{premature/late}, Violence)
 I-(Time intervention {on-time}, Violence)
 I+(Intervention strength {low/high}, Violence)
 I-(Intervention strength {exact}, Violence)

Fig. 3. Israeli-police model: description.

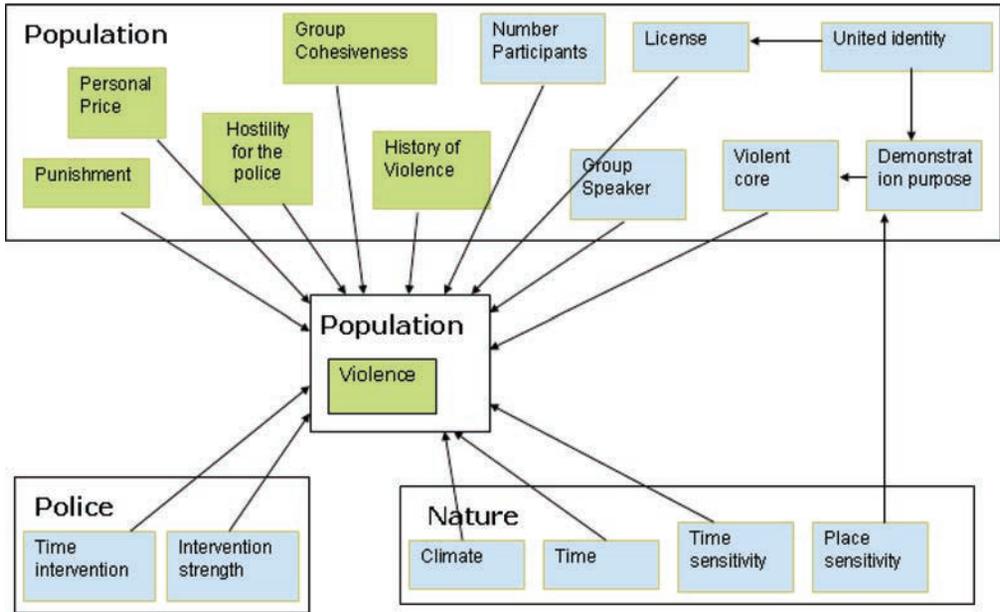


Fig. 4. Police model: structure.

violence level) increases and $I-$ means it decreases. For example, an emotional event increases the existence of a violent core among the participants.

4.3. The Bar-Ilan University Model

The third model, *BIU* (Bar Ilan University), is shown in Figure 6. It is our own novel extension of the police model. Based on interviews with social and cognitive scientists, as well as additional literature surveys [Schweingruber and McPhail 1999; McPhail 1991], we added four additional variables, and updated 19 influences (relations) among the variables. The added factors are: (1) event order (indicates amount of preparation that was made before the event, such as delineation, disposition of the police forces, etc.); (2) participants' anonymity (indicates whether the participants believe that they can be recognized and identified); (3) participants' visual cohesiveness (such as similar outfit as among football fans); and (4) light.

Figure 5 presents a model description and Figure 6 presents the model's graphical structure. We used the same entities (population, nature, and police) as in the Israeli-police model and added four additional quantities to the existing ones. In addition we updated the quantity space of the police intervention strength. We also updated the Influences (I) among the variables as fully described in the model description in Figure 5. As in previous models $I+$ represents that the derivative of the target quantity increases and $I-$ means it decreases.

We provide here several examples for updated influences. First, we updated the influence of police's intervention strength, thus instead of direct impact on violence level as in the police model, it impacts the participants' belief that they may be punished, and their hostility for the police. In the BIU model, high intervention strength increases participants' hostility towards the police and increases the participants' chance for punishment. However, low intervention strength just decreases the participants' chance for punishment without a change in hostility for the police factor. Another example is that existence of a group speaker and the request (and acceptance) of a demonstration

Participants:
 Population
 Nature
 Police

Quantities:
 Population:
 ...*Similar to Israeli-Police model*..
 Intervention strength: QS = {low
 Event order: QS = {low high}
 Anonymity: QS = {low, high}
 Visual cohesiveness: QS = {low,
 Nature:
 ...*Similar to Israeli-Police model*..
 Light: QS = {low, high}
 Police:
 ...*Similar to Israeli-Police model*..
Influences:
The influences in following quantities did.
 [Hostility to police, Personal price, History
 Event purpose, Violent core, United identit
 Time sensitivity, Place sensitivity, Weather

I+(Group speaker{exist}, Order)
 I+(License{ have},Order)
 I-(License{don't have},Order)
 I+(Order{ low},Violence)
 I-(Order{high},Violence)
 I+(Visual cohesiveness{high}, Group cohesi
 I+(Group cohesiveness{high}, Anonymi
 I-(Group cohesiveness {low}, Anonymi
 I+(Anonymity {low}, Punishment)
 I-(Anonymity {high}, Punishment)
 I+(Time {morning/noon}, Light)
 I-(Time {evening/night}, Light)
 I+(Light{low}, Anonymity)
 I-(Light{high}, Anonymity)
 I+(#participants{100-1000/1000-10000/
 I-(#participants{0-100}, Anonymity)
 I+(Intervention strength {medium/high}
 I-(Intervention strength {low},Punishme
 I+(Intervention strength {high},Hostility

Fig. 5. BIU model: description.

license increase the maintenance of order, which decreases the violence level. In contrast, in the police model, license and group speaker variables had a direct influence on the violence level. Moreover, for the variable *number of participants*, we no longer allow direct influence on the violence level as in the police model, but instead have it influence the participants' anonymity level ("the more participants around me, the less recognizable I am"). Another example of addition to the BIU model is: participants'

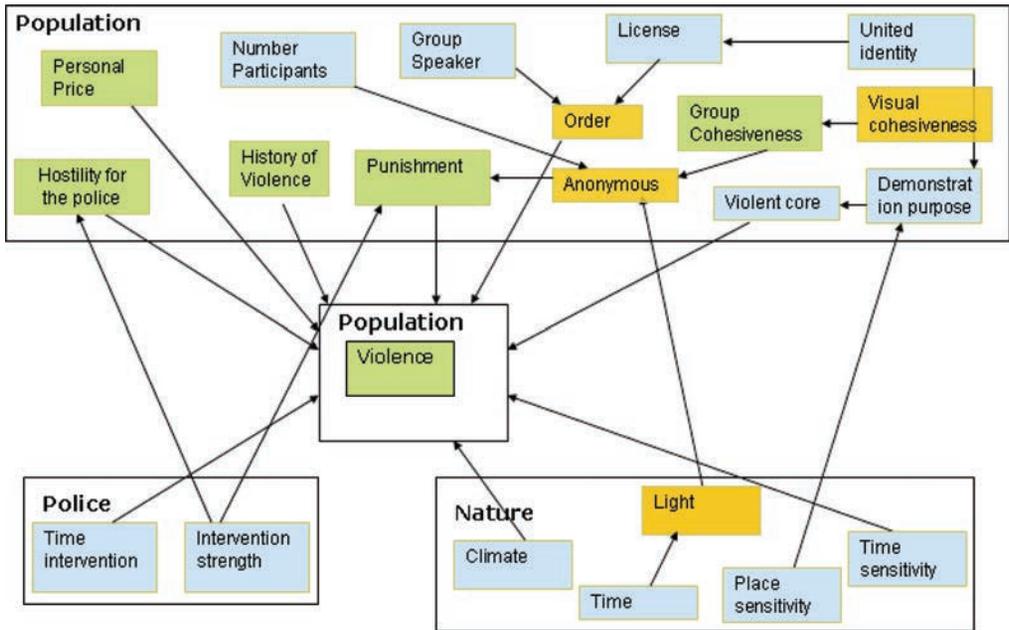


Fig. 6. BIU model: structure.

visual cohesiveness has an impact on group cohesiveness, it actually increases the sense of belonging to the same group.

5. PREDICTION AND ANALYSIS

For different demonstration cases, one can set the initial state quantities to their qualitative values, based on the demographics and environment values as known at the time. Then qualitative simulation is used to expand all possible outcomes possible based on the initial values. The resulting violence outcomes are used as the basis for prediction. Then, the simulation graph itself is used to point out specific settings in which intervening is particularly important.

5.1. Estimating the Likelihoods of Different Outcomes

A qualitative simulator takes as input an initial setting of the world state (partial state information is acceptable) and produces a simulation state transition graph. Each sequence of states, following transitions from the initial state and ending with a different outcome state, is a possible system trajectory, that is, a possible sequence of qualitative state changes that may occur given the initial state and the qualitative dynamics specified. The end state of each such path is where the system dynamics allow no further evolution (i.e., the system is stable). Taking the value of the outcome variables (in our case, violence level) in these final states allows categorical predictions.

The violence-level variable can take three categorical values: zero, low, and high. The zero value represents demonstrations that ended without any casualties and also without property damage. The low value represents demonstrations that ended with property damage but without any casualties, and the high value represents demonstrations that ended with casualties.

However, it is not enough to know whether a demonstration might be violent; in a sufficiently complex model, all three possible outcomes will have at least one stable state in which they appear. Instead, our goal is also estimate the likelihood of different

outcomes. To do this, we use the received state graph as an input. Based on this developed graph we calculate the likelihoods of different outcomes by counting the number of behavior paths that lead to a specific violence level, and divide it by the total number of paths. The result is a distribution over possible violence outcomes.

For instance, suppose that there are 345 total paths leading from the initial simulation state to stable states (leaves in the simulation graph). Suppose further that 123 of these paths end up in leaves with violence level *high*, 121 of the paths end up in leaves with violence level *low*, and the remaining 101 paths end up with violence level *zero*. Then the distribution of predicted violence levels is $\langle high, low, zero \rangle = \langle 123/345, 121/345, 101/345 \rangle = \langle 0.36, 0.35, 0.29 \rangle$.

5.2. Determining Important Factors in Specific Settings

The a priori predictions of the model, given initial values, do not provide decision makers with information about factors that, in the particular case, influence the level of violence. Thus we do not know, out of the many different factors that may increase the level of violence, which are important in the specific case being simulated.

For instance, a perception of anonymity by the demonstrators may reduce their fear of being punished for breaking the law, and this in turn can increase the chances of violence erupting during a demonstration. Perception of anonymity can be addressed by the police through various means: segregating the demonstrators into smaller disconnected groups, shining bright lights (if the demonstration is held when dark), etc. However, a priori, there are few indicators of the potential anonymity perceived by the crowd. Moreover, we do not know whether tackling such perception can be effective. It could be that there are so many factors increasing the violence that anonymity (being an indirect influencing factor) is just not worth treating. Or likewise, it could be that violence is highly unlikely, and thus bringing in bright lights is just an overkill that may incite the crowd. Thus anonymity should be addressed only in specific settings where it becomes a determining factor in promoting violence.

To aid in this decision making, we describe an algorithm for determining the k most important factors in influencing the outcomes of the simulation, and also determining the conditions under which they should be addressed. This is carried out as follows.

First, we traverse the simulation graph bottom up (from leaves to the root, which is the initial state). In each node, we count the number of paths resulting from it which end up in high violence, low violence, or zero violence. This process is in fact a generalization of the prediction process described before for making predictions. The number of paths of each type that is associated with the initial state is exactly the outcome distribution which we described earlier. Here, we are simply generating the same count for all nodes in the graph.

Now, we identify the k nodes with the highest level of *outcome entropy*, that have more than a single child¹. The outcome entropy measures the uniformity of the distribution of different potential violence outcomes. A perfectly uniform distribution $\langle 0.33, 0.33, 0.33 \rangle$ will have maximal entropy; a perfectly nonuniform distribution where all paths lead to the same outcome will have minimal entropy (0).

The reason for seeking simulation nodes with high entropy is that these are nodes where there is a difference to be made, that is, they are potentially *actionable*. Nodes with low entropy are those in which outcome is essentially decided already. Changing their outcome will necessarily involve making multiple changes to the state, that is, they involve more complex intervention. In contrast, nodes with high entropy are nodes

¹A parent with a single child will have the same count of paths going through it as its child, and thus the same entropy. But we seek the state where the divergence into multiple outcomes occurs, hence we prefer the child.

whose outcome is far from decided, and thus offer a good opportunity for relatively simple intervention.

Given the k highest-entropy nodes, we can now identify the factors that influence the outcomes. We do this by examining the simulation information saved at the node and contrasting it with that of its children. We thus determine which qualitative relations are at work at the node and how they interact to lead towards the different outcome. This significantly narrows the list of factors that are relevant to the different outcomes and also unravels the conditions in which these factors are important.

For instance, we may see that a node splits into different children because of the interaction between two opposing forces: low anonymity which decreases violence (it increases the chances of punishment, as perceived by the crowd) and the lack of police responses to events (i.e., the police are responding too late, or too weakly) which increases violence. Both these factors interact to cause multiple possible outcomes. Acting (e.g., by increasing police force) can countermand the interaction, and cause the outcomes leading from this node to converge towards low or zero violence. Moreover, the state represented by the simulation node tells us the conditions under which increasing the police force will be effective (as this is not always the correct response to violence!).

6. EVALUATION

To evaluate the approach described earlier we implemented the three models in GARP3 [Bredeweg et al. 2009], a QR engine including a user-friendly visual interface which enables building and simulating qualitative models and was successfully used in many domains [Salles and Bredeweg 2006; Bredeweg and Salles 2009]. We also developed 24 test cases, real-life demonstrations reported on by a variety of sources. 22 of these taken from Hebrew Wikipedia under category demonstrations [Wikipedia: The Free Encyclopaedia 2010]. The cases were taken both from the main category and from the subcategories “demonstrations in Israel” and “massacres in demonstrators”. We disqualified general descriptions which did not describe a specific event (e.g., descriptions of recurring demonstrations) and also omitted two cases due to lack of information (for a total of 20 cases). An additional three cases are well-known events which were extensively analyzed and described [Carmeli and Ravid-Yamin 2006; Lewis 1989; Useem 1997; Stott and Drury 2000] by experts. The last event was a peaceful demonstration that we videotaped.

To initialize the test cases, we utilized the information appearing in their descriptions in the literature and in Wikipedia. We initialized only the quantities for which we had explicit information. Quantities for which we had no information or ambiguous information were removed from the initial set. Qualitative simulation can work with such partial information.

6.1. Prediction Accuracy

Each model was examined on the 24 test cases described before. We use the simulation state graph for our calculation of the numeric probability as presented earlier. Figure 7 presents the example of transitions state graph built by GARP of one of the experiments. Figure 7(a) presents the base model built state graph, Figure 7(b) presents the police model state graph in the same experiment, and Figure 7(c) presents the BIU model state graph in the same experiment. The circles represent states and the arrows represent state transitions. The end path circles are the final states with one of the possible outcomes: zero violence, low violence, or high violence. The figure shows that the BIU model results in a much more complex and richer state graph than the base and police models; this demonstrates the complexity of reasoning afforded by a rich qualitative model.

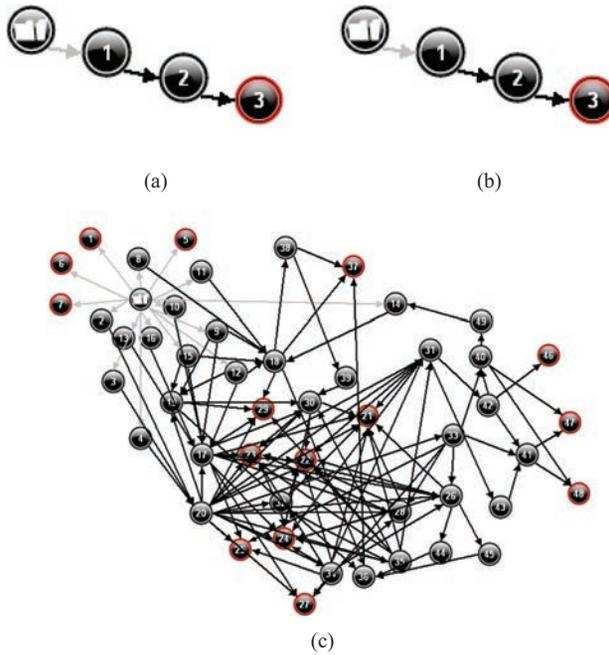


Fig. 7. Transitions state graph.

In evaluating the predictions of the different models, we look at the maximum likelihood prediction of each model, for the 24 different cases. If the maximum likelihood prediction corresponds to the outcome of the event in the real world, we count this as an error of 0. Otherwise, we examine how far off the prediction was from the actual outcome, and we define the following types of errors:

- yes, denoting success;
- one-level error, corresponding to one ordinal-level mistake such as instead of classifying to high violence the model classified to low;
- two-level error, corresponding to two ordinal-levels mistake such as instead of classifying to high violence the model classified to zero.

Figure 8 summarizes the experiment results across the 24 cases. The horizontal axis separates the different models. The vertical axis measures the number of cases. The results of the three models are presented as stacked bars. Their total height is always equal (24 cases), but they are internally divided into 0-level errors, 1-level errors, and 2-level errors. All three models have 19 successes out of 24. While the base and police models have five 2-level errors, the *BIU* model replaces four of these 2-level errors (predicting high levels of violence where there was none) with 1-level errors (predicting low levels of violence where there was none). The Bar Ilan model predictions are thus noticeably closer to the actual outcomes.

6.2. Comparison to a Machine Learning Technique

The field of artificial intelligence has other techniques which can be used to make qualitative (ordinal, in this case) predictions. We decided to compare the QR predictions against those of a popular machine learning algorithm, decision tree learning [Mitchell 1997].

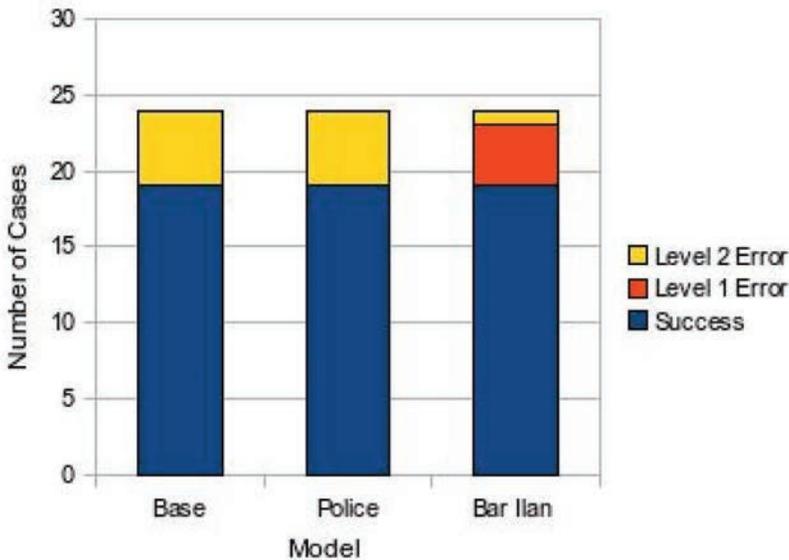


Fig. 8. Model prediction results.

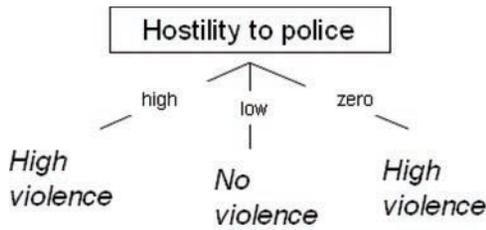
To do this, we used WEKA [Bouckaert et al. 2010], an open-source, user-friendly software that contains a collection of machine learning algorithms and used the J48 decision tree algorithm, which is considered state-of-the-art. We used the algorithm with the default parameters (confidence threshold of 0.25, and a minimum of two instances per leaf). Testing was carried out via ten-fold cross-validation, which is a standard in the field. We built three files that were used as an input to WEKA. Each file contains a collection of attributes with their values and was built based on the quantities initialization set of each QR model (base model, police model, and BIU model). The target class value of attribute violence in each file was set based on the real-life event outcome. The output of the J48 algorithm is comprised of the learned decision tree and classification statistics.

Figure 9 presents the decision trees that were learned based on the each QR model initial quantity set. Figure 9(a) presents the decision tree that was learned based on the quantity set of the base model (base decision tree), Figure 9(b) presents the tree that was learned based on the quantity set of the police model (police decision tree), and the same tree was learned based on the quantity set of the BIU model (BIU tree).

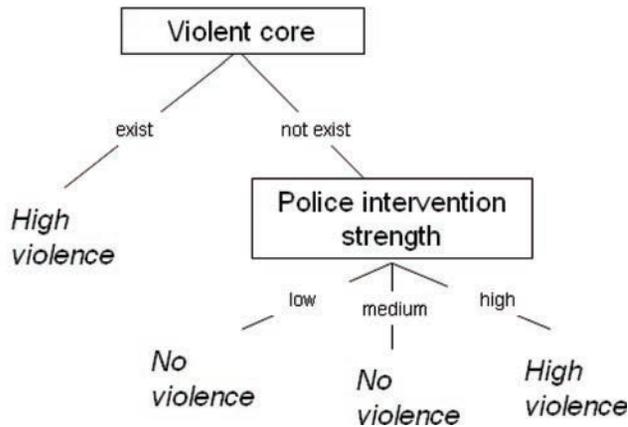
The results show that police tree and BIU tree have 100% of success in classification, while base tree has 70.8% success. While the decision tree technique provides an accurate prediction, a slightly better prediction than the BIU model with QR approach, we will claim, in the next section, that the QR approach is much more sensitive to changes and can account for what-if scenarios. Thus, using a QR approach is better for reasoning regarding the potential violence level to improve the police decision making process.

6.3. Sensitivity Analysis

In the following experiments we demonstrate the use of QR and decision trees for a variety of hypothetical changes. According to experts [Lewis 1989; Useem 1997; Stott and Drury 2000] in several of the events we modeled (Exp. 15–17), the police intervention strength was found to be one of the important factors for violence eruption.



(a) base decision tree



(b) police and BIU decision tree

Fig. 9. Decision trees.

Thus, in this section, we want to examine the presented QR model's prediction and the decision tree technique in what-if scenarios.

First we want to examine whether the presented models with QR approach and decision tree technique are sensitive enough in terms of whether they can account for what-if scenarios. Moreover, we want to examine what influence the police intervention strength has on the event outcome: Could it be the main factor that can prevent the violence, or is the event in essence a violent one and police intervention strength has little to do with it? Then we want to examine the hypothetical situation of changing the chance for violence in several test case scenarios by changing different controlled factors and not just the police intervention strength.

6.3.1. Sensitivity Analysis: Experiment 1. In this experiment we want to examine whether the presented models built with the QR approach and the machine learning technique may account for changes in the police intervention strength. We used the same 24 test cases as described in Section 6 and examined the police intervention strength attribute with all possible values. As in Section 6 we estimated the likelihood of different event outcomes. The model will be considered sensitive to the changes if for different values in the examined attribute, it will provide a different outcome. The change can be one of the following: different distribution but no change in classification, or a different distribution, as well as a change in classification.

We compared the BIU and police models built with QR techniques to the decision tree that was built with BIU initialization set. The base model built with QR techniques

is irrelevant for this experiment since the base model does not account for the factor of police intervention strength and therefore there are no changes in the model's predictions.

The results show that the police model changes its distribution in five test cases (from 24) and in two of them it also changes its classification. The BIU model changes its distribution in all of the examined test cases and in seven of them it also changes its classification. The decision tree cannot provide a distribution for all possible outcomes, it can only provide a final classification, thus unless there was a change in classification the prediction remains the same². Out of 24 examined test cases, the decision tree changed its classification in six.

These results demonstrate that a QR model (the BIU model) is much more sensitive to changes in the input conditions than a competing QR model (the police model) as well as the decision tree. But of course, these are all hypothetical changes; we do not know whether the modified predictions agree with the hypothetical results. In other words, did the predictions change correctly?

6.3.2. Sensitivity Analysis: Experiment 2. To address this question, we used the three test cases that were explored independently by experts. For these, the experts have already discussed the hypothetical results of changes to the police force used, and their conclusions can be used to validate the predictions of the models.

- The first case, Exp. 15, is the 1985 Heysel Stadium Disaster, during the European Cup final. According to Lewis [1989] who analyzed this event, one of the reasons for this violent outcome was the police's lack of intervention in preventing the developing violence.
- The second case, Exp. 16, is the Los Angeles Riots which occurred in 1991. Useem [1997], who analyzed this event, argued that the police were not properly organized and did not react in time with appropriate force to prevent the eruption. This allowed a violent core to grow.
- The third case, Exp. 17, is the London Riot Disaster which occurred in 1990. As opposed to the previous two events, here the police used enormous force against the protests without distinguishing between anarchists and peaceful marchers [Stott and Drury 2000]. The marchers turned to fight back, and this changed the initially peaceful protest into a very violent event with many casualties.

Table I presents the experiment results. The first column corresponds to the examined test case. The second column corresponds to recommended police intervention strength. Then we present the models' predictions for each possible outcome: no violence, low violence, and high violence. In the table, below each prediction, we present the change, if any, in the recommended prediction. *Dist. Change* denotes a change in the distribution, but not in overall prediction; *Classif. Change* signifies a change in the classification.

For example, if in the Heysel Stadium Disaster (Exp. 5), the police would have increased the intervention strength, the results show that there is change in police and BIU distributions. Now the BIU model assigns 83% to high violence instead of 96% (before the recommended change). While there is a change, the max posterior prediction is still high violence. The reason for this is there are other factors in this test case that lead to high violence. Another example is the London Disaster. If we decrease the police intervention strength, now the BIU model assigns only 19% to high violence and 45% to low violence instead of 57% to high violence. In this test case the BIU model changes its classification. Instead of predicting high violence it now predicts

²Note that there exist other machine learning algorithms that result in a complete distribution, and may thus be useful in this task. Exploring their use remains high on our list of future work.

Table I. Experiments Results: Changed Police Intervention Strength

Exp. #	Recommended Change	Model Outcome	Police Model	BIU Model	Decision tree
15	Increase strength [Lewis 1989]	High violence Low violence No violence	66% 0 34%	83% 6% 10%	
	Prediction Change		High Dist. Change	High Dist. Change	High No Change
16	Increase strength [Useem 1997]	High violence Low violence No violence	66% 0 34%	87% 3% 10%	
	Prediction Change		High No Change	High Dist. Change	High No Change
17	Decrease strength [Stott and Drury 2000]	High violence Low violence No violence	80% 0 20%	19% 45% 36%	
	Prediction Change		High No Change	Low Classif. Change	High No Change

low violence. This shows us that in this test case, the police force has a big influence on the event outcome which matches the experts' expectations.

The results demonstrate that the decision tree technique is not sensitive to the examined changes that were claimed by the experts. The police model performed slightly better than the decision tree (changed the distribution in Exp. 15 but failed in two other test cases). However, the BIU model provided good results which match the experts' expectations and it shows that it can account for what-if scenarios. We remind the reader that these results are limited to the use of the decision tree. Other machine learning algorithms could potentially fare better; we leave such evaluation to future work.

6.3.3. Sensitivity Analysis: Experiment 3. In this experiment we want to examine whether we can even further decrease the violence level in test cases 15 (Heysel Stadium Disaster) and 16 (LA Riots). We used the same initializations with several updates as explained shortly. Some factors such as weather or history of violence cannot be changed, while others can be controlled. For example, police's intervention strength, anonymity, and order are examples for features that can be manipulated in the sense that there are actions that can be done to change their values. Police may increase the intervention strength by using more man power or by using different kinds of weapons. The existence of projectors and cameras in the gathering zone decrease the perception of anonymity of the participants.

Table II presents the experiment results. In this experiment we examined the BIU model and the decision tree. The first column corresponds to the examined test case and the second column corresponds to the changed initial values of the quantities. Then we present the models, predictions before the change and after.

Here again the results demonstrate that the decision tree technique is not sensitive to changes. This is not surprising, since the only components of the learned tree which can change its classification are the existence of a violence core and the police intervention strength. However, the BIU model is found sensitive to the changes.

We emphasize that these results do not lead to the conclusion that machine learning algorithms in general cannot be used in a way that allows for some hypothetical reasoning and sensitivity analysis. For example, a novel form of visualization [Mozina et al. 2004] has recently been used to allow some sensitivity analysis with another machine learning method (*naive Bayes*).

Table II. Experiment Results: Hypothetical Manipulations

Exp. #	Changed initializations	Violence level	BIU Model before change	BIU Model after change	Dec. Tree before change	Dec. Tree after change
Exp15	Police strength: medium, Punishment: high, Anonymity: low	High Low No	96% 3% 1%	80% 6% 14%		
	Prediction Change			High Dist. Change	High	High No Change
Exp16	Police strength: medium, Punishment: high, Order: high	High Low No	99% 1% 0%	80% 6% 14%		
	Prediction No Change		High	High Dist. Change	High	High No Change

Table III. Important Factors in Test Cases Exp 15–17

Exp. #	Factors (Algorithm)	Factors (Experts)
15	Police Strength Too Low	Police Strength Too Low [Lewis 1989]
16	Police Strength Too Low, Number of participants > 100	Police Strength Too Low, Too Late [Useem 1997]
17	Hard Core of Demonstrators, Low Perceived Chance of Punishment	Police Strength Too Much [Stott and Drury 2000]

6.4. Determining Influential Factors

The previous section demonstrates that while QR models are sensitive to hypothetical changes to the simulated quantities, not all changes can cause a qualitative change in the predictions of the system. In other words, not all factors have equal weight in affecting the outcome of a particular case. Trying all factors in the hope of identifying those that are important is a long computer-intensive process that does not scale.

We now turn to evaluating the use of the algorithm described in Section 5 for determining important factors influencing the outcome of the demonstrations. We ran the algorithm on the resulting simulation graphs for the three cases (Exp. 15–17) for which we have expert analysis in addition to the predictions of the different models. We requested the 5 highest-entropy nodes. The algorithm analyzed the information associated with them to determine which factors were interacting to cause the different outcomes to form (or more accurately, to create children leading to the different outcomes).

In Table III we report on the top factors increasing violence in cases Exp. 15–17. The second and third columns in the table show the factors determined by the algorithm and the factors determined by experts in the field, who have analyzed these cases.

In case Exp. 15 we see complete agreement between the algorithm and the expert. In case Exp. 16 we see partial agreement: Both the algorithm and the expert agree that the police responded with too little strength, but the expert also points out that the intervention occurred too late. The algorithm, in contrast, reports the number of participants as a significant factor in the violence. This is a factor that cannot typically be changed dynamically, but of course the algorithm cannot differentiate static from dynamically controllable factors (we leave such extension to future work). Note, however, that the algorithm does recognize that the eruption of violence occurs when, in addition to responding too weakly, the police are too late. However, it does not report on it as a key factor in the eruption of violence. In this, it differs from expert opinion.

Finally, in case Exp. 17 there is an apparent disagreement between the expert and the algorithm: The expert believes that the main factor accounting for the violence is that the police acted too harshly, while the algorithm points out the existence of a

core of demonstrators and a low perceived chance of punishment as being key factors. Note, of course, that the algorithm and experts do not provide contradictory results. It could be that both are correct: our algorithm's goal is to discover opportunities for intervention, and it could be that the expert's analysis accounts for a large portion of the state space in which no intervention is possible (since there, the police acted too harshly, but this cannot be taken back).

Indeed, almost as a side-effect of this analysis, we not only discover which factors are important, but also under which circumstances to act upon them. These circumstances are easily determined by examining the state of the qualitative behavior, as denoted by the node in question. For instance, for Exp. 15, the highest-entropy state (where the algorithm recommends increasing the police response) has the following attributes: moderate weather, high cohesion of the demonstrators, emotional event, a hard core of demonstrators present, between 100 and 1000 participants, weak police strength applied too quickly, lack of a spokesman or representative for the demonstrators, evening hours and dark, property damage already caused by the demonstrators (i.e., low violence already erupted). Under these specific settings, the corrective action to take would be to immediately increase the strength of the police response in hopes of preventing the violence from escalating.

7. SUMMARY AND FUTURE WORK

In this article we described a method for modeling and reasoning about social behavior of large groups and applied it to the problem of predicting potential violence during demonstrations. We used Qualitative Reasoning (QR) techniques which to our knowledge have never been applied for modeling crowd behaviors nor in particular demonstrations. Based on social science research, we incrementally presented and compared three QR models for predicting the level of violence in demonstrations: a base model, police model, and BIU model. We evaluated these models on 24 real-life test case scenarios. The results show that the BIU model makes better predictions than the compared models and it also was found sensitive to changes. We also compared our performances to the machine learning method, a decision tree. While the machine learning method made accurate predictions, it failed in the sensitivity analysis. Thus, the BIU model built with the QR approach can account for what-if scenarios as opposed to the decision tree and is more preferable for reasoning regarding the potential violence level to improve the police decision making process.

In our future work we plan to expand our model to account for bidirectional influences (feedback loops). For example, in the BIU model the hostility for the police quantity increases the violence level. However, increasing the violence level has no impact on hostility. We believe that such expansion is necessary to provide a more accurate prediction. We also plan to tackle the next logical step in the use of QR for social simulation, which is to move beyond determining the important factors to determining plans of action that utilize them.

Of course, our overall goal in this article has been to introduce the use of qualitative reasoning—both previous algorithms as well as our novel extensions—in social sciences. The use for modeling demonstrations is but one possible domain of application. The nature of qualitative modeling seems to fit well with the nature of social science knowledge. The technique can therefore be used to validate and test theories against qualitative data.

ACKNOWLEDGMENTS

Thanks to Tomer Zelberstein, Avishay Zilka. As always, thanks to K. Ushi.

REFERENCES

- BOUCKAERT, R. R., FRANK, E., HALL, M. A., HOLMES, G., PFAHRINGER, B., REUTEMANN, P., AND WITTEN, I. H. 2010. WEKA—Experiences with a java open-source project. *J. Mach. Learn. Res.* 11, 2533–2541.
- BRAJNIK, G. AND LINES, M. 1998. Qualitative modeling and simulation of socio-economic phenomena. *J. Artif. Soc. Social Simul.* 1, 1.
- BREDEWEG, B., LINNEBANK, F., BOUWER, A., AND LIEM, J. 2009. Garp3 - Workbench for qualitative modelling and simulation. *Ecol. Inf.* 4, 5–6, 263–281. <http://hcs.science.uva.nl/QRM/software/>.
- BREDEWEG, B. AND SALLES, P. 2009. Mediating conceptual knowledge using qualitative reasoning. In *Handbook of Ecological Modelling and Informatics*, S. Jrgensen, T.-S. Chon, and F. E. Recknagel, Eds., Wit Press, Southampton, UK, 351–398.
- CARMELE, A. AND RAVID -YAMIN, I. 2006. Research report on the subject of crowd events and public order. Tech. rep., Ministry of Public Security, Bureau of the Chief Scientist.
- FORBUS, K. D. 1996. Qualitative reasoning. In *CRC Handbook of Computer Science and Engineering*. CRC Press, Boca Raton, FL.
- FORBUS, K. D. AND KUEHNE, S. E. 2005. Towards a qualitative model of everyday political reasoning. In *Proceedings of the 19th International Qualitative Reasoning Workshop*.
- FRIDMAN, N. AND KAMINKA, G. A. 2007. Towards a cognitive model of crowd behavior based on social comparison theory. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI'07)*.
- GILBERT, N. AND TROITZSCH, K. G. 2005. *Simulation for the Social Scientist*. Open University Press.
- GLYKAS, M., ED. 2010. Fuzzy cognitive maps: Advances in theory, methodologies, tools and applications. In *Studies in Fuzziness and Soft Computing*, vol. 247, Springer.
- JAGER, W., POPPING, R., AND VAN DE SANDE, H. 2001. Clustering and fighting in two-party crowds: Simulating the approach-avoidance conflict. *J. Artif. Soc. Social Simul.* 4, 3.
- KAMPS, J. AND PELI, G. 1995. Qualitative reasoning beyond the physics domain: The density dependence theory of organizational ecology. In *Proceedings of the Qualitative Reasoning Workshop*.
- KOSKO, B. 1986. Fuzzy cognitive maps. *Int. J. Man-Mach. Stud.* 24, 65–75.
- KUIPERS, B. 1989. Qualitative reasoning: modeling and simulation with incomplete knowledge. *Automatica* 25, 4.
- LEWIS, J. M. 1989. A value-added analysis of the heysel stadium soccer riot. *Current Psychol.* 8, 1, 15–29.
- MCPHAIL, C. 1991. *The Myth of the Madding Crowd*. Aldine de Gruyter, New York.
- MITCHELL, T. M. 1997. *Machine Learning*. McGraw-Hill.
- MOZINA, M., DEMSAR, J., KATTAN, M. W., AND ZUPAN, B. 2004. Nomograms for visualization of naive bayesian classifier. In *Proceedings of the 8th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD'04)*. 337–348.
- PATRICK, S., DORMAN, P. M., AND MARSH, R. L. 1999. Simulating correctional disturbances: The application of organization control theory to correctional organizations via computer simulation. *J. Artif. Soc. Social Simul.* 2, 1.
- SALLES, P. AND BREDEWEG, B. 2006. Modelling population and community dynamics with qualitative reasoning. *Ecol. Model.* 195, 114–128.
- SCHWEINGRUBER, D. AND MCPHAIL, C. 1999. A method for systematically observing and recording collective action. *Sociol. Methods Res.* 27, 4, 451.
- STOTT, C. AND DRURY, J. 2000. Crowds, context and identity: Dynamic categorization process in the poll tax riot. *Hum. Relat.* 53, 247.
- USEEM, B. 1997. The state and collective disorders: The los angeles riot/protest of april 1992. *Social Forces* 76, 2, 357–377.
- WANDER, J. 2000. Modelling consumer behavior. Ph.D. thesis, University of Groningen, The Netherlands.
- WIKIPEDIA: THE FREE ENCYCLOPEIDA. 2010. Category: Demonstrations (hebrew). <http://he.wikipedia.org/w/index.php?title=%D7%A7%D7%98%D7%92%D7%95%D7%A8%D7%99%D7%94:%D7%94%D7%A4%D7%92%D7%A0%D7%95%D7%AA&oldid=8167278>.

Received June 2011; revised October 2011; accepted January 2012