An interactive simulation for testing communication strategies in bushfires

Charles Bailly
Grenoble-INP, LIG, Grenoble, France
charles.bailly@grenoble-inp.org

Carole Adam
Univ. Grenoble-Alpes, LIG
F-38000 Grenoble, France
Carole.Adam@imag.fr

ABSTRACT

Australia is frequently hit by bushfires. In 2009, the "Black Saturday" fires killed 173 people and burnt hectares of bush. As a result, a research commission was created to investigate, and concluded that several aspects could be improved, in particular better understanding of the population actual behaviour, and better communication with them. We argue that agent-based modelling and simulation is a great tool to test possible communication strategies, in order to deduce valuable insight for emergency managers before new fires happen. In this paper, we extend an existing agent-based model of the population behaviour in bushfires. Concretely, we added a communication model based in social sciences, and user interactivity with the model. We present the results of first experiments with different communication strategies, providing valuable insight for better communication with the population during such events. This model is still preliminary and will eventually be turned into a serious game. . . .

Keywords

Agent-based modelling and simulation, communication, crisis management, GAMA platform, serious game

INTRODUCTION

Nowadays, the number of crisis events is continuously increasing, be they natural crises (fires, floods, earthquakes, tsunamis...) or man-made crises (industrial accidents, terrorism, refugees flow...). Societies can manage crisis and emergency situations in several ways: adopt urban and territory planning policies to reduce the risks (*e.g.* forbid construction in exposed areas); raise awareness and prepare the population in advance; or create efficient emergency management policies to deal with crises when they happen.

Modelling and simulation offer tools to test the effects and complex interactions of these different strategies without waiting for an actual crisis to happen, without putting human lives at risk, with limited cost, and with a great degree of control on all conditions and the possibility of reproducing exactly the same situation as many times as needed. When modelling human behaviour, mathematical, equation-based models are too limited (**parunak1998**). On the contrary, **agent-based models**, where autonomous entities (agents) interacting with each other represent the humans involved, offer many benefits (**bonabeau2002**). They allow capturing emergent phenomena that characterise such complex systems; they provide an intuitive and realistic description of their behaviour; they are flexible, offering different levels of abstraction by varying the complexity of agents.

In this paper, we are particularly interested in the bushfires that strike the state of Victoria in Australia every summer, burning many hectares of forest, causing many deaths and injuries, and destroying property. (**Dutta2016**) have recently shown a 40% increase in the number of bushfires per week in Australia over 5 years (from 3284 per week in 2007 to 4595 events per week in 2013). The current state policy is "Prepare, stay and defend, or leave early", so the population is given a choice between: evacuating early, before fire reaches their area of residence, because "many people have died trying to leave at the last minute" (**cfa-guide**); or stay and defend their house, only if very well physically and mentally prepared. In both cases, the decision must be made and a plan prepared well in advance. But in the summer 2009, serious bushfires devastated a part of Victoria, culminating on the Black Saturday 7th February when 173 people died despite all efforts at raising awareness. The cost of these bushfires was estimated to 4.4 billion dollars, with 98932 hectares of Victorian parks damaged, an estimation of one million animals dying, cost to agriculture due to loss of cattle and pastures, etc.

Several reports (**royalcomm2009**; **murrindindi2011**) have tried to explain the reasons for this heavy death toll, and have identified different inconsistencies: in **behaviour** (the population does not react as expected by decision-makers), in **information** (received information is not always considered as relevant by the population), and in **communication means** (inefficient, specifically information *broadcast*). As said by (**why2014**): "agencies need to change from an expert authoritative approach to one that seeks to understand community needs and expectations".

Adam *et al.* (**AdGa2016ssc**) have designed a model of the behaviour of the Australian population in bushfires from interviews gathered after the 2009 "Black Saturday" fires by the Victorian Bushfires Research Commission (**statements2009**) to explain the inconsistencies in behaviour in terms of a gap between objective and subjective evaluations of risk and capabilities to deal with it. In this paper, we are interested in the other inconsistencies noted in the report, namely the **communication problems** *i.e.* the fact that the content of information communicated to the population, and the means of communication used, were not optimal.

Our goal is to turn Adam *et al.* 's model into an interactive simulation where the user can **test different communication strategies** (changing the source, media, content or recipient of messages) and obtain indicators of their relative success. For such an interactive simulation to lay valid results, it is important that the underlying human behaviour model be as realistic as possible (**Ruijven2011**). The underlying agent-based model was proven valid (**AdGa2016ssc**) and successfully compared with another more complex model (**AdTaDu2017hicss**). Here we enrich it with a communication model in order to allow the agents to receive messages and to reason on their content and sender. To ensure the validity of these additions, our communication model is grounded in theories from social sciences that have studied communication issues for a long time.

The paper is structured as follows: we first introduce some relevant literature regarding serious games, communication theories, and cognitive biases. We then proceed to presenting our methodology, data, simulation platform, and the model we are extending. The next section then describes our model of communication and its implementation in this existing simulator. The following section exposes our experiments and results with this interactive simulator. Finally we conclude the paper by discussing the limitation of our approach and its future prospects.

STATE OF THE ART

Participatory simulation for raising awareness

Computer simulation is a great tool for crisis management that offers many benefits. Compared to full-scale simulation exercises, it is much less costly, less dangerous, and easier to organise. Yet it still allows to discover knowledge by exploring several "what-if" scenarios before an actual crisis happens, with complete control on all parameters. Participatory simulation is a type of simulation where human users interact with the simulated world by controlling some of the agents in the system. Participatory simulation is therefore a type of serious games, *i.e.* games that are used not for entertainment but for learning, training, or understanding mechanisms (**MIC06**).

Serious games have several benefits over more classical approaches to teaching or raising awareness. They follow a constructivist logic in which the players build their own knowledge by confronting a problem in a simulated world. A meta-analysis gathering 193 articles about serious games (**SAU07**) has shown many benefits such as: favouring the development of social and human relationships and communication skills; increasing learning motivation, self-esteem and self-confidence, engagement and persistence; developing problem-solving skills; helping learners to structure, build and represent knowledge; and helping learners to integrate information by developing the capability to build links and transfer knowledge from other contexts.

Simulation-based serious games are particularly interesting for raising awareness of various types of risks (**CRO16**). By being placed in a risky situation and allowed to try several ways of managing it, the players can better comprehend the risks and their possibility of occurrence, but also the consequences of their actions of these risks. For major risks such as bushfires, exploring different strategies and their impact in a serious game provides players with some experience, simulated but close to the real world mechanics. Such experience would be hard to acquire from real crisis in such a short time, due to the long duration between events (several months to years), and the stakes involved that prevent from trying blindly. An important aspect of serious games and participatory simulations is to rely on a pedagogical scenario integrated in the game design to answer a specific pedagogical objective (**CHA15**), and on well-specified rules (objectives, conditions of victory, possible interactions...) to guide the player.

Communication theories for crisis management

In his well-known Communication Theory, Shannon (**shannon1948**) devised a model of communication with the following components: information source, transmitter (encoder), transmission channel, receiver (decoder), destination, and message. Among these components of a message, the **channel** has already been largely studied in crisis management. For instance some works show a significant "channel effect" of social media (**SUG2011**; **USG2013**), where the same message has a different effect depending on which channel it is delivered on (social media vs more traditional channels). Trust in the **source** of the message is also important and depends on their trustworthiness, expertise and attractiveness (**PeCa1986elm**).

Regarding the **message** itself, Speech Acts Theory (**searle1969**; **vanderveken1990**) exhaustively lists the 5 types of messages that can be communicated: assertive (state a fact, provide an information), promissive (commit to perform an action), directive (ask or order the hearer to perform an action), expressive (express an emotion), and declarative (formal institutional action such as declaring someone married, or guilty). Most relevant in crisis management are assertive and directive speech acts, in order to provide the population with information and recommendations.

Regarding the **recipient**, (**devito2000**) has divided the listening process into 5 sequential phases that occur after actually hearing the message: receiving (or attending, *i.e.* actually focusing on the message), understanding (getting the meaning of the message), remembering, evaluating (forming an opinion about the validity of the message), and responding (*i.e.* provide feedback regarding acceptance of the message). This final stage can be in the form of direct feedback, or just by changing behaviour as a result of the message.

Communication can fail at any stage of this process. The message might not be heard if the recipient is not monitoring the channel (TV or radio is off). Even if the message is heard, it might not be attended to or remembered if the hearer is overwhelmed by receiving too much information at the same time (information overload); the hearer might miss relevant data that is drowned in too many irrelevant messages, which might lead them to stop listening to a given emitter because they cannot deal (**APF1999**). If attended to, the message might be evaluated as irrelevant or inaccurate and discarded. Finally, even if the message is considered accurate, it might not lead to the expected behaviour change.

Behaviour change

Messages sent by emergency managers during disaster often aim at changing the population's behaviour towards what they consider to be the best response (evacuate to a safe area after an earthquake, stay confined inside during a chemical incident, etc). Behaviour change has been extensively studied (see (**prager2012**) for a review), but a lot of works focus on medium or long-term changes (*e.g.* non-healthy habits like smoking) while bushfires are short-term emergency situations.

The Elaboration Likelihood Model (**PeCa1986elm**) is a theory describing how attitudes can be changed by persuasive stimuli. It distinguishes two types of processing: central (cognitive, high effort) and peripheral (heuristic, low effort).

- The **central route of processing**, used when the individual is motivated and able to process the message carefully, is a more cognitive, high-effort elaboration of the message received, based on its actual logical value. As a result, it leads to more resistant attitude change, and is more predictive of behaviour change.
- The **peripheral route of processing**, used to reduce mental efforts when the individual is not motivated or unable to process the message, is a less thorough elaboration of the message that relies more on cues and heuristics. For instance it focuses on the credibility, attractiveness or familiarity of the source, at the expense of the actual logical value of the content. It is therefore more influenced by mood, or by emotions towards the emitter.

Individuals try to reduce their mental efforts and will thus tend to use the peripheral route, unless they are sufficiently motivated and able to elaborate on the message.

- **Motivation** is affected by the relevance, interest and consistency of the message with the recipient's current beliefs (contradictory messages are more easily rejected); it is also affected by the recipient's personality (do they like thinking, whatever the subject).
- **Ability** is affected by the recipient's knowledge (is it sufficient to critically evaluate the content of the message) and familiarity with the subject, and the availability of their cognitive resources (how busy or distracted they are, time pressure).

As we can see, peripheral processing is more likely to happen during disasters due to the stress and time pressure. However, it leads to less lasting attitude changes and is less likely to trigger behaviour change. Crisis communication should therefore be adapted so as to favour central processing, by sending only clear, understandable and relevant information.

Cognitive biases

Message acceptance and behaviour change is also impacted by various cognitive biases, phenomena described in psychology and social sciences that twist reasoning towards ”irrational” shortcuts in order to make faster decisions. They are particularly relevant during crises (**yudkowsky2008**; **KKRP2014**) when decisions are made under high stress and time pressure.

Previous studies (**AdGa2016ssc**) have found occurrences of these cognitive biases in the population interviews performed after the bushfires (**statements2009**). For instance :

- the *confirmation bias* is a tendency to give more credit to information confirming existing beliefs and to discard inconsistent information (*e.g.* interpret the presence of firemen as a cue that everything is safe in order to confirm motivation to stay);
- over-estimation of danger to others and under-estimation of danger to self (residents often report they knew there were going to be fires but felt they were not going to be impacted; however many residents worried for their friends, neighbours and relatives);
- the *anchoring effect* is an excessive focus on the first information received that prevents from changing one's initial decision even when receiving further (possibly contradictory) information (the authorities report an over-commitment of residents to their defense plan even when the fires were known to be much too strong to be fought);
- the *bandwagon effect*: doing and believing the same as others around (*e.g.* residents who believe it is safe because their neighbours stay);
- the *sunk cost fallacy* consists in refusing to abandon a goal even when new information would require it, because of having already invested in it. It is therefore self-reinforcing as more actions are performed to reach the goal (*e.g.* residents who have invested a lot in building and preparing their house to resist the fires are less likely to abandon it, even when informed they should evacuate).

Behaviour profiles in bushfires

Alan Rhodes (**why2014**) has extracted 6 behaviour profiles of the population in bushfires.

- **Can do defenders:** skilled and experienced defenders, relying nearly only on their own abilities to manage the situation;
- **Considered defenders:** good defenders with several possible plans to deal with fires; they are more likely to listen to warnings from the authorities than Can do defenders;
- **Livelihood defenders:** will protect their property whatever the danger, because it is their source of income;
- **Threat monitors:** focused on defense but will immediately escape if they feel a real danger;
- **Threat avoiders:** focused on escape, will decide to run away as soon as they are aware of fires;
- **Unaware:** do not feel concerned by fire risk, and do not known how to react in case of a fire.

However, he provides no information about the distribution of these profiles in the population, and states that they are not linked with demographic features.

OUR MODEL

Based on these elements, we designed an interactive simulation of the population behaviour in response to various communication strategies during bushfires. Our simulation is implemented in GAMA, and based on an existing simulator, extended here with a communication model grounded in the theories presented above.

GAMA platform

GAMA (**GAMA2013**; **Drogoul2013prima**; **Drogoul2013paams**) is an open-source platform for agent-based modelling and simulation, offering an integrated programming language and development framework to develop elaborated models with up to several millions of agents. The GAma Modelling Language (GAML) is a high level agent-based language based on Java, specifically designed to be easy to use even for non-computer scientists, allowing domain experts to create and maintain their own models. GAMA also provides native management of GIS (Geographical Information Systems) data allowing to integrate geographical data files into simulations. Finally, GAMA offers interactive functions (user commands) enabling the use of the participatory dynamics required in our interactive simulation.

Existing simulator

Adam and Gaudou (**AdGa2016ssc**) have implemented in GAMA an agent-based simulator of the behaviour of the Australian population in bushfires, based on the population interviews gathered after the 2009 fires (**statements2009**). In their simulator, the world is a grid of 50*50 cells inhabited by four species of agents: fires, houses, shelters and residents.

Fires are reactive agents; they are initially placed randomly on a free cell and then grow up randomly at each cycle, increasing their intensity (which directly increases the damage they deal) and size. People cannot go through fires when escaping but can cross their area of effect (representing the smoke and heat zones around the fire). When all fires are extinguished, the simulation stops automatically.

Houses are inhabited by exactly one resident (no families). They offer some amount of protection but residents can still be hurt if fires are close enough. Houses can be reinforced by their owner up to a given point while fires are still far away. When fires are close enough, they deal damage to the house until possibly destroying it.

Shelters are safe areas where people cannot be harmed by fires. Residents know the location of some shelters, and when choosing to escape they aim at the closest one they know (which might not be the absolute closest one if they ignore its location).

Finally **residents** are the most complex agents in this simulator. They have various attributes to represent their health, their motivations (to defend or to escape), their risk perception (awareness of fire, assessment of danger), and their abilities. Their possible actions are to prepare their house and themselves, to escape towards a shelter, to defend against the fire, or to take cover in their house. The choice of action is determined by a finite state machine architecture (see Figure 1): in each state, the corresponding action is performed, and the transitions to another state are constrained by the values of attributes and the position of fires.

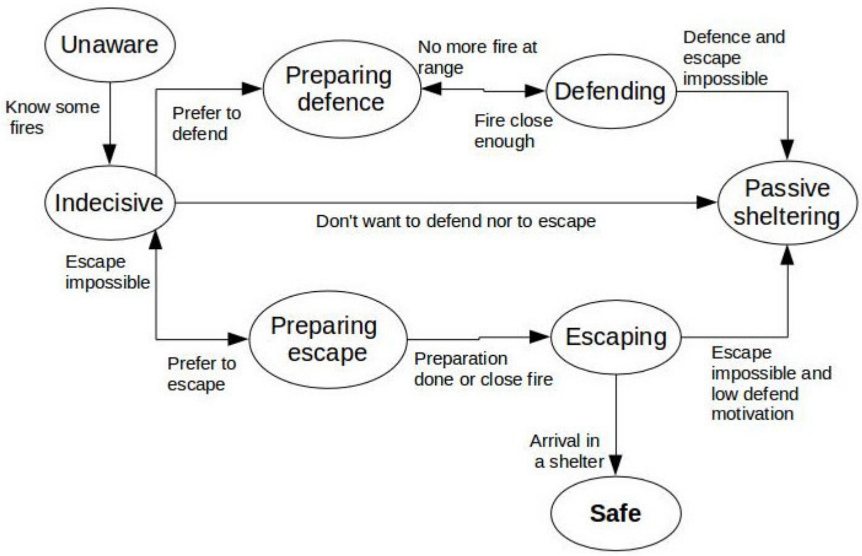

Figure 1. Finite state machine architecture of residents agents

The following paragraphs describe the extensions we performed on this model.

Communication strategies

Based on the communication theories presented above, we modelled several types of communication strategies, that concern the different components of a message, namely the choice of its source (authorities, fire soldiers, general media...), its content (information, recommendations), and its target (all the population, a precise geographic area, or a specific category of residents).

- *Content-based* strategies are focused on what is concretely told to residents. We modeled two of the 5 types of speech acts: information about fires (*e.g.* position) and recommendations (advice about the appropriate behaviour, *e.g.* evacuate). Indeed, these are the ones most frequently related by the residents in the interviews, either because they received them, or because they wish they had received them.
- *Target-based* strategies concern the accuracy of messages. We modeled three possibilities: global broadcast (target all residents), geographical-based (target people in a specific area), and plan-based (target residents based on their declared fire plan: defend or escape).
- *Source-based* strategies concern the emitter of the message. We have implemented different possible sources (firemen, authorities...) as well as indirect communication strategies where global authorities send messages to local managers (*e.g.* mayors) who filter them and spread relevant data to their neighbourhood. *Composed* strategies consist in sending a sequence of messages in a precise order. This type of strategies aims at determining if it is efficient to send a combination of several messages, and which order of messages is most efficient (for instance inform about fires before or after giving recommendations).
- Finally we also implemented various *shelter-based* strategies in order to compare the efficiency of building many shelters vs communicating more about the existing ones. These are therefore not strictly communicative strategies as the player can also choose to create new shelters. Shelters are designated safe areas (cricket oval, community house, etc) where residents are invited to gather in case of a fire.

We implemented these communication strategies in the existing GAMA simulator in the form of user actions: the user can right click in the simulation window at any time during the simulation; they are then invited to select a communication strategy and specify its features (source, channel, etc) before executing it. Each communication action has a different cost (for instance it is more costly to accurately target communication than to broadcast, and more costly to build shelters than to advertise them). In the next paragraph, we describe how we also updated the residents model to allow these agents to receive messages, interpret them, reason on them and make relevant decisions so as to change (or not) their subsequent behaviour.

Psychological aspects of decision making

In order to obtain a realistic model of residents and of their handling of messages, we implemented several psychological phenomena based on the theories exposed above.

Profiles of behaviour

We implemented the different profiles of behaviour in terms of ranges of values of their attributes, and modified transitions in the finite-state machine. For instance livelihood defenders cannot escape but are more likely to shelter if needed; they also have a high ability to fight fire and a high trust in local sources. The type of population can be selected as a simulation parameter.

Trust in message source

Trust is one of the most intuitive aspects involved in communication. If the source of the message is not trusted, the receiver is not likely to take its content into account. In our model, residents have a "trust probability" attribute: it is a table matching each source with a probability to trust the messages it sends. The values in the trust table depend on the resident's profile. For example, *Can do defenders* have a higher trust in local sources of information than *Threat monitors*, but a lower trust in firemen.

Message acceptance

To represent the information overload phenomenon, we added an acceptance probability for each source in the pedestrian attributes. The initial values are based on the resident's trust in each source. When a resident receives a message, they might accept or reject it based on the resident's acceptance probability for its source. The resident then updates their acceptance according to the message accuracy and its perceived relevance. Thus, if the message is perceived as inaccurate (too global or imprecise) or useless, the resident's acceptance probability for messages from this source will decrease; on the contrary relevant messages will increase this acceptance probability.

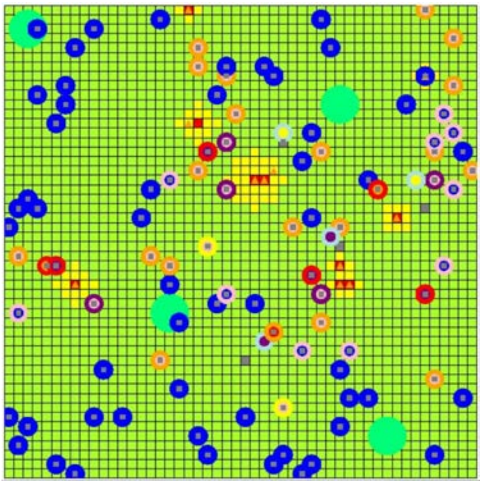

Figure 2. Screenshot of the graphical experiment in GAMA

Cognitive biases

A number of biases are taken into account in our model. For instance the confirmation bias and anchoring effect are implemented by the fact that the resident's motivation influences their risk assessment and vice versa: a resident motivated to defend will under-estimate risk (and discard cues of a higher risk), and low estimation of risk will increase their defense motivation (defense becomes an anchor). The sunk-cost fallacy is implemented thanks to the feedback from actions: the defense motivation gradually increases with successful defense actions, making it harder to give up defense.

EXPERIMENTS AND RESULTS

This section presents our first experiments with this model. Below we define the scenarios that we want to test (*e.g.* compare broadcast with personalised communication), as well as the indicators used to measure success of the compared strategies, before describing our results.

GAMA experiments

The GAMA platform allows two types of experiments:

- Graphical experiments: we run one simulation and observe the behaviour of the agents "live" to have a quick overview of their reaction to various strategies. (See a screenshot on Figure 2.)
- Batch experiments: automatically run many iterations of each simulation (with the same parameters). In that case there is no graphical display but GAMA outputs graphs of the average values (over all iterations) of selected indicators, which allows to smooth out the randomness of the simulation.

Scenarios.

In this paper we describe the results of batch experiments for 4 scenarios that compare communication strategies with each other, depending on the target population. Here we focused on 2 of the 6 profiles of behaviour, namely the can-do defenders (planning to defend their property) and the threat avoiders (planning to escape fires).

- Scenario 0 is our baseline: we compared the values of all indicators on the two populations without any communication.
- Scenario 1: compare possible **contents** of strategies (give information about fires vs recommendations) on these 2 populations
- Scenario 2: compare the **accuracy** of communication campaigns (broadcast vs geographically-targeted vs plan-targeted)
- Scenario 3: compare **shelter-based** strategies on the 2 populations

Indicators. To be able to compare the different strategies, we defined and implemented the following indicators measuring their relative success or failure: number of deaths; number of injuries; total damage to houses; total cost of communication actions.

Settings. For each scenario (comparing several simulations), we ran 60 iterations of 200 cycles of each simulation (which represents over 3 hours of simulated time with one minute long cycles). Between the compared simulations, we only varied the communication strategy tested, with the other parameters being exactly the same (number and strength of fires, population, availability of communication channels...).

Output. The output of our batch experiments consists in graphs showing the comparative average values of these indicators for different strategies. The graphs obtained for the scenarios defined above are discussed below.

Scenario 0: baseline, no communication

We first compared the impact of fires on the two populations when there is no communication actions performed. With the threat avoiders population focused only on escape, and since we did not implement firemen in our simulation, we expected the fires to grow out of control and lead to a high number of victims and great amount of damage. There was indeed a huge gap (in terms of damages, injuries and deaths) compared to the can-do defenders population. The next scenarios compare the impact of various strategies on these two different populations.

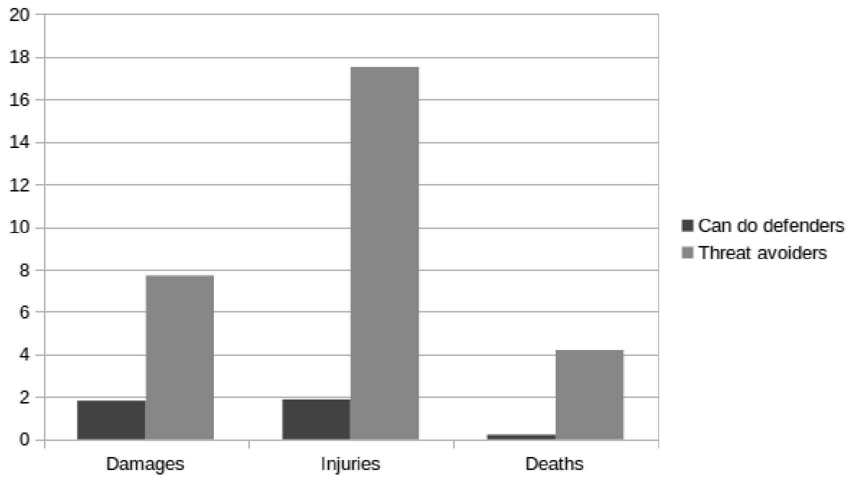

Figure 3. Comparing Can-do defender and Threat avoider populations - no communication

Scenario 1: comparing information vs recommendation messages

Can-do defenders. As expected with such skilled defenders, values of building damage, injuries and deaths are quite low. As shown on Figure 4, there is no big differences between broadcasting fire information and recommendations in terms of damage or cost. There are slightly more injured and dead with recommendations only. Our data allows to explain this observation: in the absence of information messages, many residents remain unaware of the fires (too far to be perceived directly) so they do not feel concerned and ignore recommendations. As a result they do not prepare and end up being more vulnerable when the fire arrives. Since less people defend, the fire grows faster and later blocks or injure escapers. On the contrary with information only, residents are aware of all fires even far away, which increases their subjective risk perception and in turn influences their behaviour: more people escape early and stay safe. The difference between the two strategies remains small because can-do defenders rely mainly on themselves rather than the authorities, so they are less likely to accept the messages anyway.

Threat avoiders. The difference between information and recommendations is much more visible on threat avoiders since they are more likely to accept messages (Figure 5). With the information messages, there are more escapers and less defenders because of danger perception distortion (knowing more fires leads to over-estimating danger), and therefore more damage to buildings. Furthermore, threat avoiders have a lower capability than

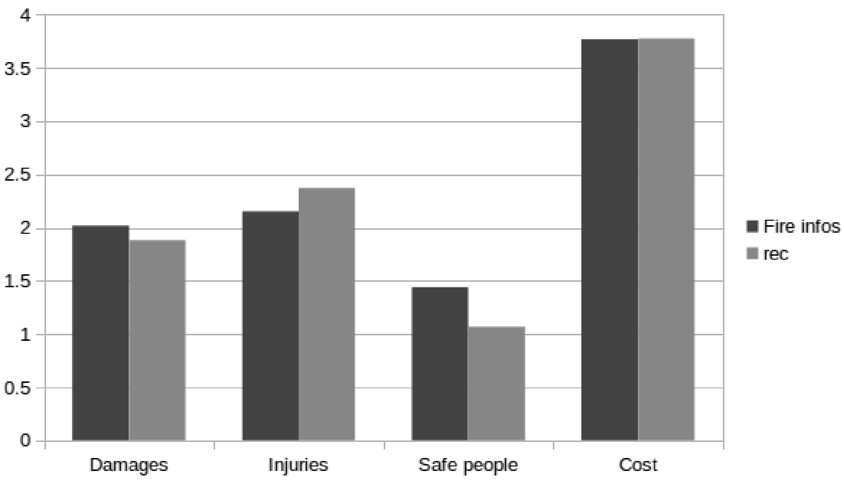

Figure 4. Comparing impact of content-based strategies (information vs recommendations) on Can-do defenders

can-do defenders, their escape will be less efficient so they are more likely to get hurt, leading to more injuries. Recommendations are therefore more efficient than generic information.

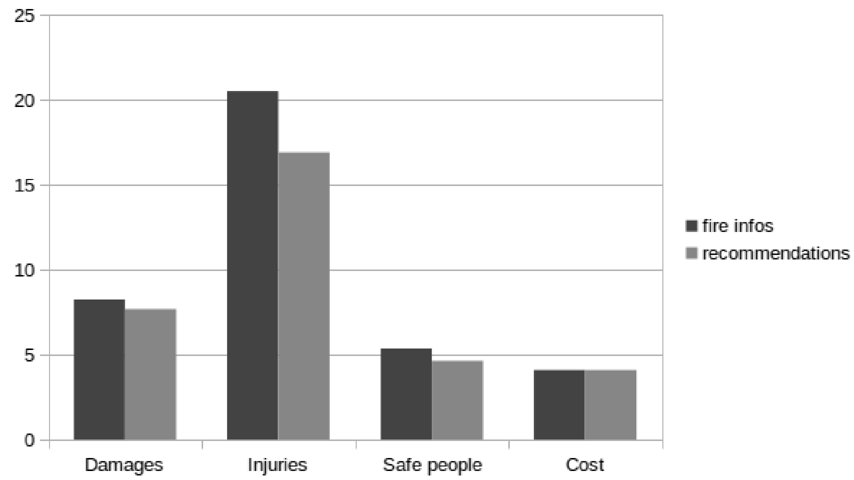

Figure 5. Comparing impact of content-based strategies (information vs recommendations) on Threat avoiders

Scenario 2: impact of accuracy of messages

Can-do defenders. The accuracy of messages has a more significant impact than their content on can-do defenders. Broadcasting is the worst strategy, geographical targeting the best but most costly, and plan-based targeting is a compromise (Figure 6). Contrary to what one might think, message acceptance is not the major factor here (can-do defenders are skilled so they can make good survival decisions even if rejecting most messages). What matters most is the number of escapers: again, broadcasting all fires to all residents leads to over-estimation of danger, encouraging more residents to escape. Can-do defenders are somewhat protected against over-estimation of danger due to their lower trust in messages and higher ability to observe the fires, but we still observed about half of the population escaping and half defending. On the contrary, with geographically targeted communication people are informed only about fires close to them and therefore have a more accurate perception of immediate danger. As a result less residents decide to escape without being in real danger, so they are more likely to protect their home and avoid damage and injuries.

Threat avoiders. Not surprisingly, the most efficient strategies are the same as for Can-do defenders (targeted communication is better than broadcast), even though the values of the indicators (damage, injuries) are much higher here (see Figure 7). This is because threat avoiders are more likely to accept and react to the messages even

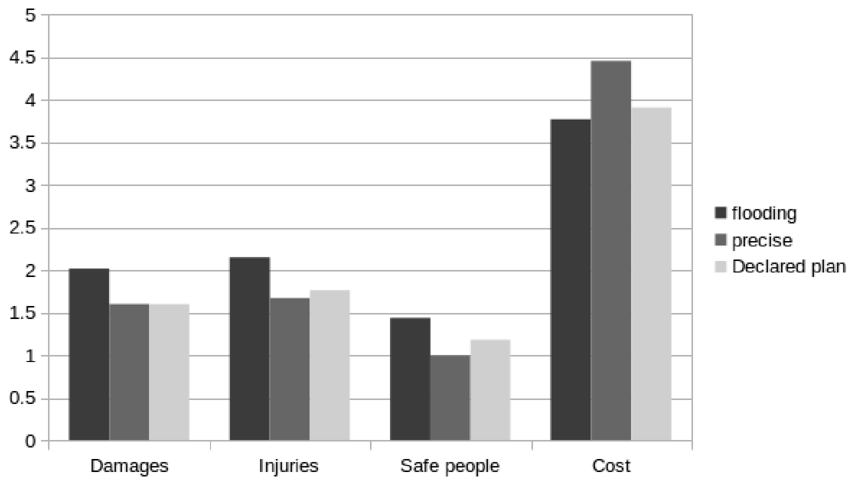

Figure 6. Comparing impact of target-based strategies on Can-do defenders

when not directly concerned (fires too far), and more likely to over-estimate danger due to their lower ability to judge by themselves. As a result we observed much more escapers and less defenders in this population. Therefore less people fight the fires, resulting in more damages to buildings, and the fires grow bigger, injuring more escapers.

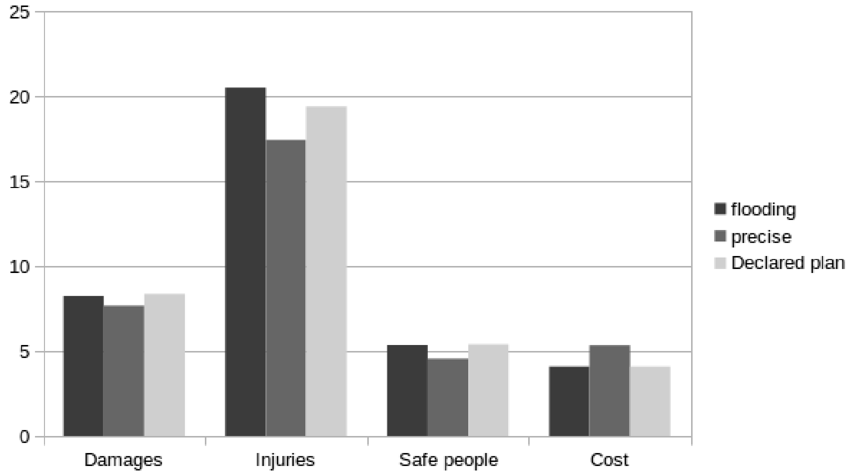

Figure 7. Comparing impact of target-based strategies on Threat avoiders

Scenario 3: impact of shelter communication

The goal here was to compare the efficiency of two opposite strategies: building more shelters but without informing people about them, or relying only a few shelters but advertising a lot about them. Both strategies aim at making sure that all residents know where to escape if needed.

Can-do defenders. We observed no real difference between these 2 opposite strategies on can-do defenders (see Figure 8). Communication is only slightly better but much less costly than building new shelters. This is not really surprising. Can-do defenders are skilled and experienced and therefore know the position of the shelters already (even if there are only few of them) and can reach them easily. Moreover, most of them decide to stay and fight against fires anyway, making shelter-based strategies irrelevant to them. Finally with more defenders, there is less congestion on the roads to the shelters and the fire does not propagate so fast, so the escapers are more likely to reach the shelters uninjured.

Threat avoiders. Shelter-based strategies have a much more differentiated impact on threat avoiders (see Figure 9). The first thing we can notice is the difference in terms of safe people. Threat avoiders have lower abilities to react to

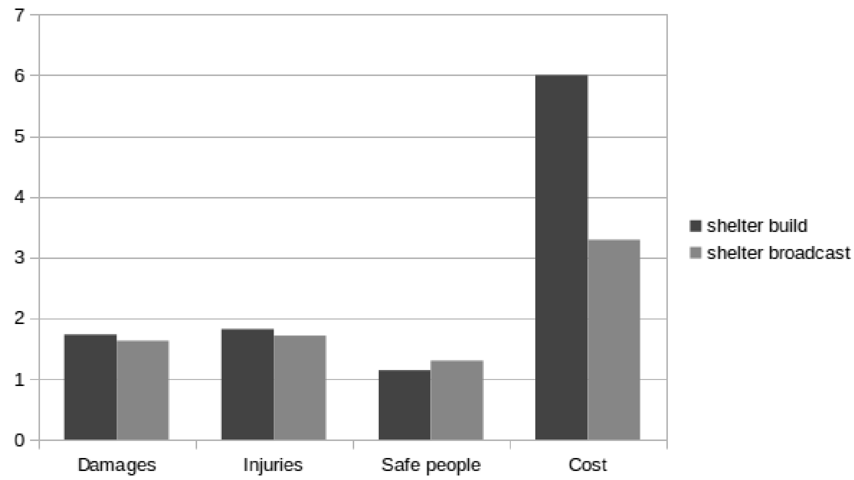

Figure 8. Comparing impact of shelter-based strategies on Can-do defenders

fires, and may not know the position of shelters or not be able to quickly reach them. When building more shelters rather than advertising them, some residents do not know any safe area and just run randomly to avoid fires, leading to more injuries and less safely sheltered people. The longer the simulation time, the bigger the fire grows, and the more randomly running people get trapped. However, the damage is slightly lower in that case as some residents, not knowing where to go, end up sheltering in their own house and defending it. Advertising the existing shelters reduces the time spent on roads looking for one, and therefore the number of injuries incurred by escapers.

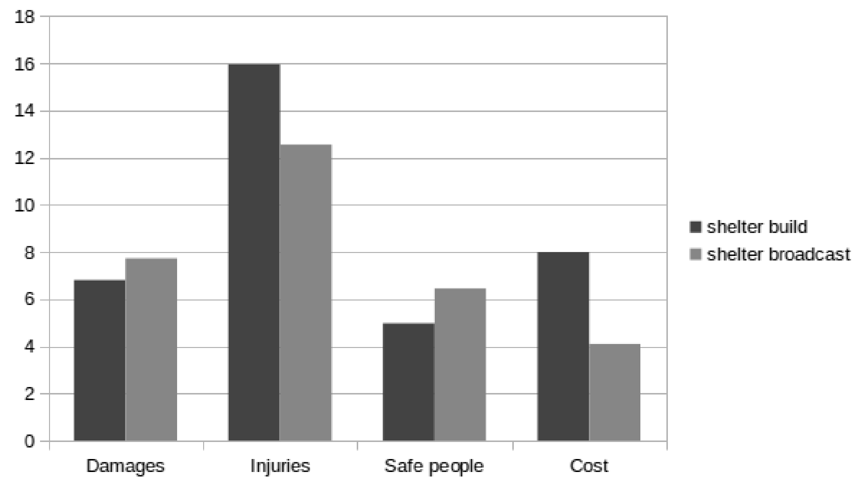

Figure 9. Comparing impact of shelter-based strategies on threat avoiders

Of course these are the two extreme strategies (building only vs communication only) and we expect the best strategy to lay somewhere in the middle, with a balanced compromise of having enough safe areas and making sure everybody is aware of them.

Discussion

Personalisation to the profiles

As expected, communication should be personalised to the different profiles found in the population (or at least to broader categories of defenders vs escapers), in order to provide each resident with information that is relevant and helpful to them without drowning it in a flow of irrelevant messages.

Our results show that the best communication strategy towards can-do defenders is to inform them about fires in their geographical area. Its downside is its high cost (the smaller each target area, the more different messages need to be elaborated and sent), so targeting residents based on their declared plans may also be a good and less costly alternative. Messages about safe areas are secondary for intended defenders and should be used parsimoniously.

Accurate geographical targeting is also better for avoiders (even though more costly), but contrary to defenders recommendations are more useful for them than bare information, as they are not skilled enough to interpret the latter. Moreover, early information about shelters is useful to raise awareness of escape possibilities and trigger earlier evacuation, thus reducing injuries incurred while escaping.

Plan-based targeting

For the plan-based broadcast strategy, we assumed that each resident had previously declared their fire plan (intention to defend or to leave) to the fire authorities. This is not the case in reality (in these scarcely populated areas, the fire brigades might have some information but not necessarily about everybody); however experimenting with this strategy allows us to show what could be done if the population was asked for their fire plan in advance and therefore still provides valuable insight.

Besides, it also shows the interest of indirect communication: global authorities might not know each resident's fire plan when broadcasting messages, but local managers might do. They could act as a filter between the large scale broadcasting and their local residents, receiving all messages and only forwarding the most adapted ones for each resident. Of course this is also a costly strategy, but automation could be investigated.

More experiments needed

A big limitation of our experiments is that they are all performed on a homogeneously profiled population in order to draw relevant results for that population. However in reality, the population is heterogeneous with residents of all 6 profiles listed above, and a continuous range of motivations and abilities. Strategies are yet to be tested on a population with a realistic distribution of these profiles, but we first need to obtain data about what this distribution is in the actual population modelled.

We also tested all strategies independently, while in reality emergency managers may use a combination or sequence of several strategies, for instance broadcasting information and recommendations at the same time. More experiments are yet to be conducted with more realistic combined strategies, possibly in cooperation with emergency managers.

Finally our experiments were quite short: 200 cycles of simulation means about 3 hours of simulated time. In reality these fires can last for days. Longer-term communication strategies thus need to be tested as well, to ensure residents stay out of affected areas, respect possible road blocks, etc.

Simplification

A model always ought to be a simplified version of reality. There are however a number of improvements that could be made in future work to improve the realism of our simulation. Of course the random fire model is the first thing that comes to mind; it should be replaced with a realistic model of fire taking physical and meteorological parameters into account (wind strength and direction, rain, temperature, etc) when computing propagation and growth. Such models already exist (**MHSP2015spark**).

Of more interest to us is the agents involved in the simulation. So far we only modelled residents as autonomous agents, and emergency managers are ”played” by the user but can only send messages. It would be very interesting now to also model firemen and their different actions on the field, from fighting the fire to communicating with the population and helping them. The only presence of firemen also has a great psychological impact on residents (**KKRP2014**) which should be modelled.

CONCLUSION

In this paper we described our extension of an existing simulation of the Australian population in bushfires. The underlying simulation was validated against real data obtained mainly from residents' interviews after the 2009 bushfires (**AdGa2017jasss**). Concretely, we enriched the residents model to allow them to handle messages, and modelled several psychological processes that influence this handling. We also added interactive functionality for the user to test various communication strategies on this simulated population. Finally we ran batch experiments to highlight the pros and cons of different possible strategies on different profiles of residents, and deduced some useful insight for emergency managers. However these first experiments were mainly intended as a proof of concept of our simulation, and much remains to be done to improve our simulator and to test more realistic strategies.

In particular, we intend to dedicate some future work to the following aspects. First, we will model a more realistic population, with a heterogeneous distribution of the different profiles, and also with social relationships and attachment between the agents. Second, we will go deeper in the formalisation of cognitive biases (**ArAdDu2017**); here we focused only on a few of them that affect message handling, but many others play a role in disaster reactions; other interesting psychological factors at play in such situations include emotions. Finally the communication model could also be enriched; in particular we want to add more possible message contents (promissives and expressives might be relevant, for instance expressing fear about the situation, or promising that the fires are under control); also we want to more deeply study trust (in the source and channel), its impact on behaviour change, and its dynamic.

It is important to notice that such work is generic and can be applied to different types of disasters. For instance, we have already developed a similar serious game for raising awareness about coastal floods risk (**AdTa2016sprite**). We believe that agent-based modelling and simulation is a great tool to raise awareness and prepare crisis management plans for any types of crisis, be they natural disasters or man-made events.

ACKNOWLEDGEMENTS

This work was partially supported by project AGIR SWIFT.