Intelligently Generating Possible Scenarios for Emergency Management During Mass Gatherings
(PhD Confirmation Report)

Peter Serwylo
PhD Candidate
peter.serwylo@monash.edu

Grace Rumantir
Supervisor
grace.rumantir@monash.edu

Frada Burstein
Supervisor
frada.burstein@monash.edu

13th July 2011

1 Introduction

Synthetic data sets have been used in many different areas. Motivations for synthesising data include stress testing applications with large amounts of data (Gray et al., 1994), anonymising data sets for public release (Reiter, 2002) and training analysts to identify interesting or alarming situations (Whiting et al., 2008). Most of the time the synthetic data is generated because there is not enough real life data to test the system in question (Abtew et al., 1990) or it does not exhibit the required characteristics (Aboulnaga et al., 2001). This project will produce realistic synthetic scenarios able to cater for cases where more data is required, and where specific characteristics are desirable. The resulting synthetic scenarios will contain known ground truths, and will be useful for training, simulating, or decision making.

The domain used to evaluate the research in this project is Mass Gathering Medical Care (MGMC). Many definitions have been provided for the term Mass Gathering. The definition adopted for this research is as follows:

> "A situation (event) during which crowds gather and where there is the potential for a delayed response to emergencies because of limited access to patients or other features of the environment and location." (Arbon, 2007)

Examples of mass gatherings include sporting events, concerts and festivals. At such events, there is generally a higher potential for injuries than would occur in a similar sized population in the general public (Arbon, 2004). To deal with these injuries while not adding undue burden on the local emergency medical services (EMS), event organisers provide first aid staff to treat patients. In the context of mass gatherings, a patient is defined as "an individual presenting for care to on site first aid or medical services at the event" (Serwylo et al., 2011). The volume of patients at an event can be measured by the Patient Presentation Rate (PPR), calculated as the number of patients presenting for care at an event per 1,000 attendees.

This project will result in a new method for generating synthetic patient scenarios, which will be used to enhance planning and preparation for MGMC. The synthetic scenarios will be used as input into computerised simulations of mass gathering events. This simulations will optimise the allocation of resources for events, by identifying when there is not adequate or appropriate resources to deal with particular series of patient scenarios. The more scenarios which are available, the better the simulation will be able to evaluate and optimise the allocation of resources for mass gatherings.

The remainder of this report will be organised as follows: Section 2 will discuss the relevant literature from both the MGMC and synthetic scenario domains. This will be followed by the research goals in Section 3 and a discussion of the methodology to undertake the research in Section 4. A discussion of the contributions to both the MGMC and synthetic scenario domains is presented in Section 5, before presenting the current progress and future plans in Section 6.

2 Background

This section will begin by reviewing the literature in MGMC in Section 2.1, followed by research on generating synthetic scenarios in Section 2.2.

2.1 Mass Gathering Medical Care

During mass gatherings, which typically have more than 1,000 people in attendance, an "over-dependence on the local Emergency Medical Service (EMS) system can result in an inability of the system to respond to routine calls, depriving the local community of the care it needs" (Yazawa et al., 2007). In order to reduce the burden on the local EMS service, medical teams should be present to deal with minor complaints, and transfer more serious complaints to a hospital (Health and Safety Executive, 1999). The issues associated with provisioning this type of medical care is referred to as Mass Gathering Medical Care (MGMC).

Research into MGMC consists of a large number of case studies into different events, describing a range of details including the medical resources provisioned, injuries which were documented, and any other observations which were made. Milsten et al. (2002) analysed over 30 case studies and found that the main factors contributing to injuries at events are the weather, alcohol and drug use, attendance and crowd density, duration, event type and crowd mood.

Building on this literature, researchers have developed models which predict how many patients will present at a given event. Some models are for specific events which occur regularly (Baird et al., 2010, Zeitz et al., 2001), while others are targeted at more general mass gathering events (Arbon, 2002, Kman et al., 2007). Despite research into predicting the PPR for an event, there has been significantly less research towards proper resource provisioning based on the expected number of patients.

> "Resource allocation to these events generally is based on experience and historical knowledge of events retained by individuals. Medical resource allocation traditionally has lacked substantive evidence to support decisions regarding appropriate staffing levels. There is a need to make resource allocation at mass-gathering events more closely mirror actual needs rather than educated guesses." (Zeitz et al., 2005)

The amount of literature regarding medical resource allocation for mass gatherings is quite small. An early case study of several college football games in the United States resulted in suggestions as to the level of resources for future games (Rose et al., 1992), although these suggestions are only applicable to similar events. A more comprehensive guide published by Health and Safety Executive (1999) contains a medical staffing model that considers the risk profile of an event in determining the medical resources required (Smith et al., 2010). However, the model is only for the United Kingdom, and there is no information as to how the risk scoring system used in the model was arrived at. Hartman et al. (2009) provide a similar scoring model, and although they go into some depth discussing how they analysed 55 different mass gatherings, there is also no discussion about how the scores were arrived at. Their model is also quite low resolution, only able to classify events into either minor, intermediate or major, and provisioning resources accordingly. This neglects the fact that mass gatherings exhibit variation in the amount and types of patients presenting, and there are more than three possible outcomes which need to be planned for.

This project will address the resource allocation issue by using computer simulations to find the optimal amount and types of resources for a given event. These simulations will be populated by software agents representing medical resources and patients. When a patient agent requests medical aid, an available medical agent will respond. This request/response paradigm will then identify scenarios where there are not enough medical resources to deal with a request, as well as scenarios where there are idle resources being wasted. By running several such simulations with varying combinations of resources and patient presentations, this will allow the amount and type of medical resources to be optimised.

The simulation process will work in a similar fashion to Airy et al. (2009), which provided an agent based simulation for distributed rescue teams. Agents in this system file requests for certain resources, which are forward to resource provider agents. The resource provider agents provide a cost/benefit analysis of whether it is worthwhile assigning a resource to a particular task or not.

The patient presentation scenarios used for the computer simulation could be used to train event organisers and first responders to deal with multiple different scenarios. This can be accomplished via interactive training simulations. There are three categories of training simulations: live, virtual and constructive simulations (of Defence, 2011). Live simulations involve real people interacting in a real environment (e.g. a real life training exercise), virtual simulations involve real people interacting with simulated systems (e.g. a computer game), and constructive simulations involve the coordination of simulated people in simulated systems (e.g. a game of chess) (of Defence, 2011). Virtual simulations provide a compromise, being more realistic and immersive than a constructive simulation, while cheaper and easier to reproduce than a live simulation (Boldyreff et al., 2011).

Although not a replacement for the other forms of training, virtual simulations have been successfully applied in a number of different domains (Crichton et al., 2000). They can be implemented via different technologies, ranging from virtual or augmented reality to virtual worlds such as Second Life (Boulos et al., 2007) or game engines such as the Unreal Engine (McGrath and Hill, 2004). The benefits of the latter two options are that they can be implemented with commercial, off the shelf technology, while virtual or augmented reality may require more specialised hardware and software. Unfortunately implementing virtual simulations for MGMC is out of the scope of this project, and will need to be addressed in further research.

One thing which both resource allocation simulations and virtual training simulations for MGMC have in common is that they both benefit when a larger number patient scenarios are available. Simulations with more patient scenarios will allow for more scenarios to be considered when optimising the allocation of resources. Event organisers and first responders using virtual simulations for training may learn to respond to a wider variety of potential scenarios. Due to differences in policy regarding data collection at mass gatherings (Arbon, 2004), there are events for which there is not adequate historical data to base resource allocation decisions on. As such, this project will develop a method for generating synthetic patient scenarios for use in simulations. The following section discusses previous literature in the field of synthetic scenarios.

2.2 Synthetic Scenarios

There is a collection of research on synthesising data, beginning with simple, random data through to complex, realistic scenarios (Figure 1). This review will analyse the various motivations for performing data generation, the scope of each technique, and how they relate to this project. Each method of synthesising data has some restrictions. Although these restrictions may not impact on how the technique is used for its primary purpose, they may prevent it from being applied in other domains. The review will then conclude with ways in which the current project will differ from and improve on the previous literature.

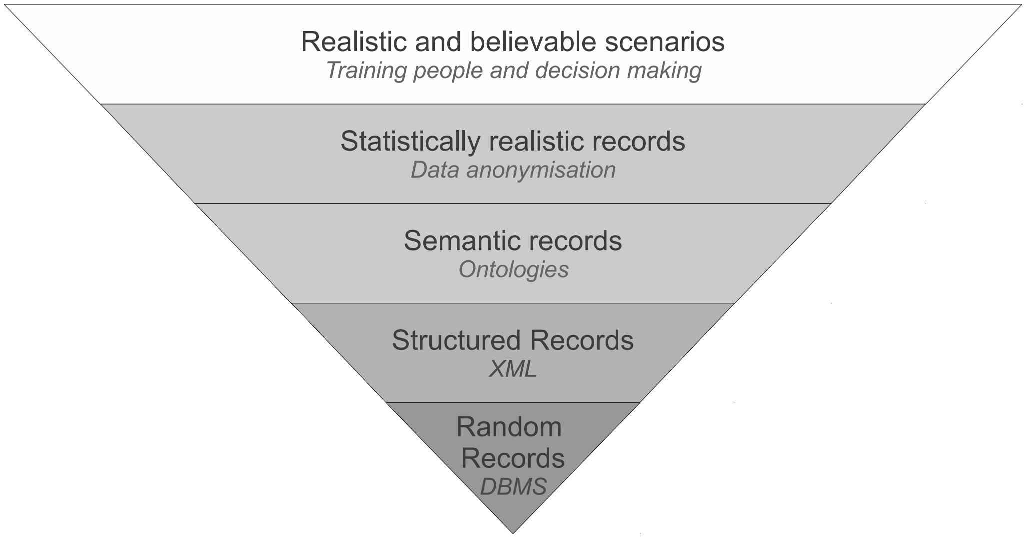

Figure 1: Synthetic data has become more realistic, and been used for more advanced purposes

Testing Software Early research on generating synthetic data focussed primarily on producing large data sets containing mostly random records. This was for the purpose of stress testing and benchmarking software systems. The requirement for synthetic data used in testing DBMS' was mainly that it satisfy certain join constraints (DeWitt, 1991). Other than this, the data could be simply random numbers or text. Similar but more complex relational constraints were placed on synthetic data generated to test XML processing systems (Aboulnaga et al., 2001) and ontological reasoning systems (Wang et al., 2005). Although these systems required the data to exhibit realism in terms of its structure, there was no need to have any degree of realism in the content of the records being generated. Generating synthetic patient scenarios requires a high degree of realism in the content within the scenarios, not just in the relationships between generated scenarios.

Anonymising Publicly Released Data More realistic data has been synthesised for the purpose of anonymising data sets for public release. This process of anonymisation is referred to as statistical disclosure control, and is important to ensure that personally identifiable, or other private information cannot be inferred from the public data set. Sweeney (2002) showed that it is not enough to simply remove identifying attributes such as name and address. They were able to easily obtain the "anonymised" hospital records for Massachusetts, and link it to the publicly available list of voters in the same state, using the unique combination of ZIP code, DOB, and gender (Sweeney, 2002). Adding small amounts of random noise to the attributes is also considered bad practice, because it may remove or falsely add relationships into the data (Sweeney, 2002).

As such, researchers have worked on taking a data set, and then synthesising a new data set which doesn't contain any of the original data in it (Reiter, 2002, Rubin, 1993). Despite this lack of real data, it still exhibits the same statistical properties, and therefore can be used to make inferences (Reiter, 2002, Raghunathan et al., 2003). Unfortunately, what makes this approach inappropriate for other domains is that there needs to be a complete set of data for which to base the synthetic data on. Therefore, it is not useful in situations like MGMC where there may not be enough data available.

Climate Simulations Much of the research in generating realistic data based on incomplete datasets is found in climate science. Often weather records such as rainfall (Abtew et al., 1990), cyclones (James and Mason, 2005) and extreme weather events (Semenov and Barrow, 1997) are not complete. Therefore, the datasets may not be ideal for making predictions about future weather patterns and behaviour. Generating realistic weather information based on these incomplete records ensures there is enough data to perform simulations of future weather events. Many researchers have had success with such synthetic data generation, but there are still limitations which prevent these techniques being applied elsewhere. The data generated by these models, although realistic, are fairly simple. Data such as wave height (James and Mason, 2005) or precipitation levels (Abtew et al., 1990) require methods compared to the complex scenarios that are described by several dependent attributes, such as those found in MGMC.

Identifying Terrorism Threats To find examples of complex, realistic, synthetic scenarios which exhibit similar characteristics to patient presentations (i.e. multiple dependent variables), we need to look at the Threat Stream Generator project (Whiting et al., 2008). This focuses on synthesising data pertaining to terrorism threats, and is able to produce realistic synthetic scenarios which have been described as similar to "mystery novels" (Whiting et al., 2008). The scenarios are constructed with lots of complex, multimedia data such as news articles, photographs, audio and video. Despite being a leap forward in creating realistic, synthetic scenarios, the process is still very manual, even after considering the tools included by the authors to help in generate the scenarios.

Gaps Identified The type of data that is required for training and planning in MGMC is realistic complex scenarios with many dependant variables. Ideally it should be able to be generated based on as little historical data as possible, so that events without any historical data can make use of it. To improve on the Threat Stream Geneartor work (Whiting et al., 2008), the process should be as automated as possible. In light of these requirements, Section 3 will discuss the specific aims of this research.

3 Research Aims

This section will discuss the specific aims of this research project. It will begin by specifying the goals in Section 3.1, followed by the questions which will be answered by the project in Section 3.2.

3.1 Research Goals

The main goal of this research is to devise a new method of generating synthetic scenarios. By applying this scenario generation to the field of MGMC, we will be able to simulate mass gathering events to help optimise resource planning.

The scenarios that will be generated in this project will be patient presentations. At any given event, it is expected that a certain number of patients will present requiring medical treatment. Each patient has certain characteristics (e.g. age, weight, state of mind, etc.), as does their medical situation (e.g. abrasion, heat stroke, etc.) and the environment which caused the injury (e.g. location, time, weather, etc).

Providers of medical care at mass gatherings in Australia and other countries often collect data relating to patient presentations at events. This data provides a rich description of how events have unfolded at past documented mass gatherings, and has already been used to predict how many patients will present at future events (Section 2.2). One of the goals of this project is to utilise that data to enhance the planning phase in ways other than predicting the number of patients.

The way in which this project aims to improve planning for mass gatherings by utilising patient presentation scenarios is through computer simulations of mass gathering events. These simulations will provide insight into whether the allocated resources are able to deal with the volume of patients. This would be an autonomous simulation where intelligent software agents would pose as patients and others as first aid responders. The first aid responder agents would follow a set of rules in order to respond to patient presentations, and then report back at the end of the simulation as to whether they were able to cope with the demand. This will help identify if there were specific patients scenarios which they were ill-equipped to deal with, and could suggest possible modifications to the allocated staff and resources for subsequent runs of the simulation.

The computer simulation for optimising resource allocation would benefit from synthetic but realistic patient presentation scenarios. For example, using only historical data, resource planners could learn from previous events to see if the current allocation of resources are adequate. Adding new, realistic but synthetic scenarios with known ground truths to the historical data sets would allow planners to consider whether their resources are adequate for as yet undocumented scenarios. This provides an extra level of robustness to the planning stage. Not only will the plans be built on past knowledge and events, but also on potential future events or those which have not been documented yet. Synthetic patient presentation scenarios will also be beneficial for simulating events which don't have any historical data about patient presentations.

3.2 Research Questions

This thesis will attempt to answer three main questions:

1. What method can be formulated to allow generation of realistic synthetic scenarios for use in mass gathering medical care?
2. In what way can the synthetic scenarios be effectively used to improve resource allocation for mass gathering medical care?
3. What characteristics must a domain have to enable generation of realistic synthetic scenarios using the proposed method?

The questions relate to formulating a new method for generation of synthetic scenarios and potentially generalising it across multiple domains. The following paragraphs will address each question in turn, why it is relevant, and how this thesis will answer it.

1. What method can be formulated to allow generation of realistic synthetic scenarios for use in mass gathering medical care? Although Section 2.2 discusses several methods for generating synthetic scenarios, most of the generated scenarios are single numerical values, such as rainfall levels (Abtew et al., 1990) or wave heights caused by cyclones (James and Mason, 2005). As such, the methods used are not applicable to the domain of MGMC, where we are interested in generating synthetic patient injury scenarios, each consisting of several dependent variables (Section 2.2). The few methods which focus on generating complex scenarios require a large amount of manual input in order to produce the synthetic data sets (Whiting et al., 2008).

In order to generate synthetic scenarios for this domain, a new method of generation is required. An ontology will be used because it can provide, at the very least, the set of concepts in a domain, and the structural relationships between them. Given an ontology encodes this knowledge in a machine understandable format, there is less need for an expert to be present for the synthetic scenario generation, allowing for a higher level of automation. Ontologies themselves do not contain enough information to successfully synthesise realistic data about a domain. Although an ontology describes relationships between concepts in a domain, it does not provide a mechanism for specifying dependencies and causality between these concepts. It also fails to define the probability that a particular scenario will arise. In order to incorporate this sort of knowledge into the scenario generation, other sources of information are required. The historical patient data collected from previous mass gatherings and a Bayesian network will be used for this role, ensuring that the chance of a given scenario being present is comparable for both the synthetic and real datasets. While Bayesian Networks are often used for causal reasoning tasks, in this case it will make use of the same causal probabilities in order to ensure the appropriate distribution of particular scenarios. It is anticipated that together, these three sources of information (the ontology, historical data and a Bayesian network) contain enough knowledge to produce realistic synthetic scenarios.

In order to show that the synthetic scenarios are able to be used in place of real historical scenarios for situations where little historical data is available, it is important to evaluate the realism and quality of the scenarios. For this project, this will be done both quantitatively and qualitatively.

The quantitative evaluation will use statistical methods to ensure that the properties and patterns found in the synthetic scenarios are comparable to real scenarios. This will ensure that any synthetic scenarios used in decision making exercises exhibit realistic properties. Previous research in the field of statistical disclosure control (Section 2.2) makes use of statistical techniques to evaluate the properties of a synthetic data set. For example, sampling synthetic scenarios from the generated set and comparing distributions of certain attributes to those of historical data sets (Reiter, 2002) can help ensure that the synthetic scenarios are distributed in a realistic manner. Techniques from the literature may be appropriate for evaluating synthetic patient presentation scenarios with slight modifications.

For qualitative evaluation, we must to involve domain experts in the field of MGMC to make judgements about how accurate and reasonable the synthetic scenarios are. A previous research project achieved this by holding competitions at conferences, where participants were asked to analyse the synthetic data in the same way that real data would be analysed (Whiting et al., 2008). After conducting their analysis, they were surveyed about the quality and realism of the synthetic scenarios. Another possible way for domain experts to quantitatively evaluate the synthetic scenarios is via a blind test, where they are asked to choose whether a particular scenario is synthetic or real. This could then be followed by a questionnaire asking about the specific characteristics of the data which led them to that conclusion. The responses from this blind test could then be used to modify and enhance the scenario generation method, ensuring more realistic and believable synthetic scenarios.

2. In what way can the synthetic scenarios be effectively used to improve resource allocation for mass gathering medical care? At most mass gatherings, event organisers provide medical staff, such as physicians, paramedics and first responders. These staff are responsible for treating patients which present for medical treatment. During treatment, they may make use of equipment such as bandages or medication. Deciding on the correct level of resources (medical staff and equipment) has historically relied on the knowledge and past experience of the event organisers (Zeitz et al., 2005). This project aims to address this issue through the use of computerised simulations.

When planning for an event, the organiser will be able to run several simulations to optimise the allocation of resources. Each simulation will represent a potential set of patient scenarios which may unfold at an event. Each simulation will also be provided with a number of resources to treat the patient scenarios. At any point in the simulation, if there are not enough resources to treat a patient, then that is an indication that not enough resources have been allocated.

There are certain things which need to be considered before optimising the number of resources based on a single simulation of some patient scenarios. For example, how common are the specific scenarios which were presented to the simulation? If the scenarios are quite common and often occur in real events, then an under-allocation of resources becomes a problem. If, however, the scenarios are very uncommon, then it only becomes an issue if the patient scenarios have severe repercussions such as serious injuries. This information may be able to be identified from querying the ontology to identify the severity of a given scenario. Another consideration is to balance the importance of being able to deal with patient presentations with the cost of over-staffing an event. This will need to be decided in consultation with domain experts, to ensure that the balance is correct.

3. What characteristics must a domain have to enable generation of realistic synthetic scenarios using the proposed method? After a method has been formulated, evaluated and validated, it is worthwhile investigating if that method can be applied to other domains in addition to MGMC. The input which is required to generate synthetic scenarios will be historical data, a domain ontology, and a Bayesian network. Given that there are other domains which have this type of information available, they could potentially benefit from the proposed method of generating realistic synthetic scenarios.

Although other domains may have ontologies and Bayesian networks available, this does not automatically qualify them to be able to make use of the proposed method for data generation. This part of the project aims to identify the specific characteristics that must be present within a domain or a specific problem in order for the proposed method of synthetic scenario generation to be generalisable.

4 Methodology

This research will be framed by the *design science* methodology, in comparison to research in fields such as social science and cognitive psychology which are traditional framed around *behavioural science* methods. Whereas the behavioural science paradigm endeavours do find out "what is true", design science focuses instead on "what is effective" (Hevner et al., 2004). The two main activities undertaken during design science research are building artefacts and subsequently evaluating them. This research will focus on building and evaluating a method for generating realistic synthetic scenarios. It will then evaluate the proposed method both quantitatively and qualitatively. March and Smith (1995) also include the additional research activities of *theorizing about* and *justifying* the artefacts in their design science framework. Theorizing has been described as "the construction of theories that explain how or why something happens" (March and Smith, 1995), and justification refers to "theory proving, and requires the gathering of scientific evidence that supports or refutes the theory" (March and Smith, 1995). This project will theorize about the method for generating synthetic data, by investigating what is required to generalise the technique for domains other than MGMC. And although out of the scope of this research, justifying this theory would require applying it to other domains to see if in fact the method can be generalised, given a certain set of characteristics present in a domain.

		Research Activities			
		Build	Evaluate	Theorize	Justify
Research Outputs	Constructs				
	Model	✔			
	Method	✔	✔	✔	
	Instantiation	✔	✔		

Table 1: The main artefacts being designed during this research (table adapted from March and Smith (1995)).

Table 1 shows a matrix which defines the research outputs and activities which can be undertaken during research. The ticked cells represent the artefacts which will be produced during the course of this project. The following paragraphs will discuss the artefacts specified in Table 1. They will begin with a definition of the artefact in the design science context, and then a description of what this research expects to produce for that artefact.

Building* the *Constructs In the design science paradigm, *constructs* provide the language required to define and communicate problems and their solutions (Schön, 1983). When discussing methods of synthetic data generation, projects are invariably multidisciplinary. As such, the constructs used to frame previous research in synthetic data tend not to revolve around the process of generating the data, but rather the domain for which the data is being generated. This will be the same approach taken for this research. The constructs used to define the problem will be framed in the MGMC domain, rather than the synthetic data domain (for which there are very few constructs used in the literature).

In order to ensure conformity with previous research, this project will make use of an ontology which unambiguously defines the terms and features of the MGMC domain. Delir Haghighi et al. (2010) developed the Medical Emergency Management in Mass Gatherings (MEMMG) ontology (Figure 2), an extension of the Domain Ontology for Emergency Management (DOEM) (Sujanto et al., 2008). The main concepts identified in the ontology were collected from domain experts in MGMC, in addition to other sources of information such as public reports, journals, conference papers and the Emergency Management Australia manual series for mass gatherings.

Building* the *Model A *model* is a "set of propositions or statements expressing relationships among constructs" (March and Smith, 1995). In design science

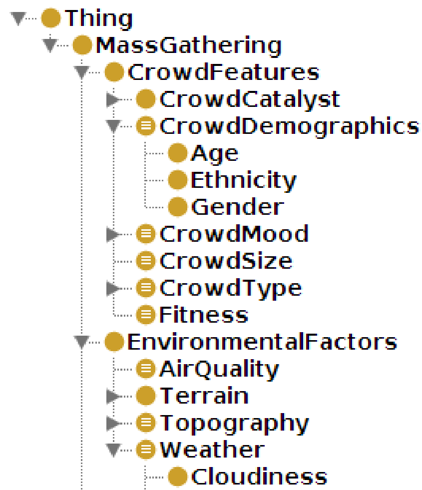

Figure 2: Part of the MGMC ontology from Delir Haghighi et al. (2010)

research, models describe a design problem and its solution space (Hevner et al., 2004).

> "Focus of attention is the key to success - focussing on the particular features of the situation that are relevant to the problem, then building a problem space containing these features but omitting the irrelevant ones". (Simon, 1996)

So once the appropriate constructs are available for critically evaluating literature on data generation, then *models* can be built which describe the relationships between the constructs for different domains.

This project will first define models of previous domains that have utilised synthetic data, e.g. climate modelling (Semenov and Barrow, 1997) and data anonymisation (Raghunathan et al., 2003). It will then follow by describing the problem of generating synthetic scenarios for MGMC. This will include both the details of generating the scenarios, and also using them for training first responders and using simulations to specify resource requirements. This model will help with the following stage, where a specific method for generating data will be detailed.

Building* the *Method A *method* is a "set of steps or an algorithm used to perform a task" (March and Smith, 1995). The method produced by this research will generate realistic synthetic patient scenarios. The method will make use of a Bayesian Network to infer the likelihood of certain scenarios occurring. Research Question 1 (Section 3.2) discussed the use of historical data about patient presentations and a Bayesian network, used to ensure the probability of a given scenario arising in the generated data set is accurate. The historical data will be obtained from collaborators from Flinders University and

St Johns Ambulance Australia. They have previously documented data about patient injuries, and used it in published research regarding PPRs at events (Zeitz et al., 2001, Arbon, 2002, Serwylo et al., 2011). Although some data sets are up to 10 years old, there has been discussion about using recent data sets.

The method being proposed for synthetic data generation will depend heavily on one main source of information, the Bayesian Network describing the domain. Currently we are not aware of any published Bayesian Networks for the domain of MGMC, and as such one will be developed for this project (Figure 3). In order to minimize the level of input required from domain experts, the structure of the Bayesian Network will be partially learned from the ontology (Helsper and Van der Gaag, 2002), and partially from the historical data regarding patient presentations (Korb and Nicholson, 2004). The prior and conditional probabilities required for the Bayesian Network will be derived primarily from the historical patient data. It is expected that there will be some nodes in the network which were constructed from the ontology, which do not have a corresponding entry in the historical data, and these nodes will require domain expertise to calculate the required probabilities to complete the network. The step labelled “Build Bayesian Network” in Figure 3 will be the main focus of this stage of the research.

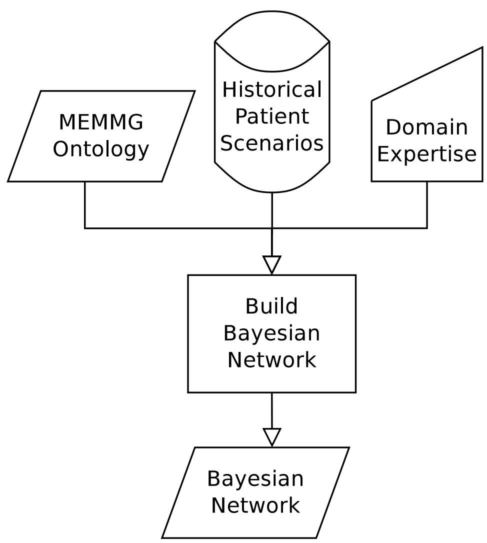

Figure 3: Process for creating the Bayesian Network

Instantiating* the *Method Hevner et al. (2004) stress the importance of being able to prove that an artefact (in this project, the method for generating synthetic scenarios) is able to be implemented in a working system, is feasible, and suitable for its intended purpose. This project will therefore implement a system for generating synthetic scenarios, based on the method described above, allowing generation of possible patient scenarios at mass gatherings. Although the method used is general, this instantiation of the method will be generating scenarios for MGMC. Specifically, it will generate patient presentation scenarios for use in simulations to optimise allocation of resources.

Once the MGMC Bayesian Network has been designed, different sampling methods will be investigated to identify the best way to generate realistic scenarios. Some existing Bayesian Network software such as Netica (Norsys Software Corp) provide implementations of stochastic simulation. Stochastic simulation uses the Bayesian network to generate scenarios based on the network distribution (Korb and Nicholson, 2004). This entails traversing the network and at each node, assigning a random value to the variable, weighted by the conditional probability of the node given any parent node states. There are different techniques for this type of sampling, and this project will assess the strengths and weaknesses of each, for generating realistic scenarios.

In addition to implementing the method for generating synthetic data, the computer simulation for optimising allocation of resources at mass gatherings will also be implemented. The project will make use of existing frameworks for developing agent based software, such as the Java Agent DEvelopment Framework (JADE) (Bellifemine et al., 2005). The main types of agents in the simulation will be patient agents and medical resource agents. Each event which is simulated will be assigned a specific set of medical resource agents, such as physicians, paramedics, ambulances or first responders. During the course of an event, generated patient scenarios will manifest themselves as patient agents in the simulation, and thus will need to be treated by an available and appropriate medical resource agent. The metrics used to optimize the allocation of resources are how many times a patient is unable to be treated due to lack of resources, and how long excess medical resources are left idle. The goal of the optimisation problem is to minimize the number of patients who are not treated, and minimize wasted resources. The resource allocation simulation will also form one component of *evaluation* for the generated scenarios, as discussed in the next section.

Evaluating* the *Method and Instantiation Evaluation of the method and instantiation for generating synthetic scenarios will be completed both analytically and experimentally. The analytic evaluation will involve quantitatively and qualitatively analysing the generated scenarios to ensure that they exhibit the desired characteristics. Research Question 1 (Section 3.2) discusses ways in which quantitative and qualitative evaluation will be undertaken using statistical methods and domain experts. The experimental evaluation will consist of inputting the generated scenarios into a simulation to optimise resource allocations. This experimental evaluation is not directed at validating the correctness of the synthetic data, but rather the utility of it: is it useful for its desired purpose?

Theorise* about the *Method The final stage of the research will be to theorise about the generalisability of the method for generating synthetic scenarios. Can it be applied in domains other than MGMC? This is discussed in Research Question 3 (Section 3.2), and may be addressed by analysing the types of dependencies and causal relationships represented by the Bayesian Network. This could be done by modifying the network in some way, generating a new set of data, and then evaluating it using the same techniques mentioned in the previous section to see if it is still valid. Any modifications which result in invalid, incorrect or unrealistic data would imply that the particular relationships that existed prior to the modification is required for the data generation to be successful. Other ways in which the generalisability of the method can quantified will be explored in the latter part of this project.

5 Contribution

This project is inherently interdisciplinary, and as such, there will be contributions to the field of MGMC and synthetic scenario generation. This section discusses the contributions to both fields.

Synthetic Scenario Generation The part of the project relating to enhanced methods of realistic scenario generation will build on literature in data generation (Section 2.2). The main contribution here will be the introduction of a new method for generating synthetic scenarios, catering for more complex types of scenarios than previously allowed. In order for this to eventuate, a new method will have to be developed which uses ontologies, Bayesian networks and historical data. Also, in the process of developing the scenario generation method, it is anticipated that contributions to validating ontologies, deriving Bayesian networks from ontologies, and validating synthetic data sets will also be made.

The MEMMG ontology being used (Delir Haghighi et al., 2010) is relatively young and is still undergoing improvements. In order to construct the Bayesian network, the ontology and a historical data set will be combined. In the process of analysing the data set, new relationships may be discovered which are not yet in the ontology, or relationships specified in the ontology may not be supported by the data set. These observations can be passed on to domain experts familiar with the ontology in order to see if it is worthwhile adjusting the ontology to reflect the newly discovered details.

There is a small amount of research on using the structual information found in an ontology to help derive the graphical structure of a Bayesian network (e.g. Helsper and Van der Gaag (2002)). This project will build on that research, to incorporate using the available historical data to identify the prior and conditional probabilities for nodes in the network. The process of ascertaining the probabilities for nodes in the Bayesian network will in turn build on previous techniques, such as those mentioned in Korb and Nicholson (2004).

Another anticipated contribution stemming from the new method of generating synthetic scenarios is in the evaluation of the resulting synthetic data sets. Earlier research into synthetic data (Section 2.2) reveals several different techniques for measuring the validity of generated data. These evaluation techniques will need to be modified to cater for the interrelationships and dependencies between variables for patient presentation scenarios.

Finally, once the evaluation is complete, the feasibility of using the scenario generation method in domains other than MGMC will be investigated. This will allow other domains to evaluate whether they can make use of the method. The specifics of this contribution are discussed in Research Question 3 (Section 3.2). Allowing multiple domains to make use of the one method for generating scenarios will make this project stand out from previous research, which typically focuses on one particular domain (Section 2.2).

Mass Gathering Medical Care In the field of MGMC, this project will contribute to the planning stage. This will be done by providing new ways for organisers to decide the correct allocation of resources, based on the results of computer simulations. Where previously organisers had to rely on their own past experience, this thesis will contribute a new tool to compliment the knowledge and experience of organisers.

6 Progress and Future Plans

This section briefly documents the progress of the candidature, along with the plans for the rest of the project. Section 6.1 introduces the literature which was reviewed in order to better understand the domain of MGMC. This section also discusses the resulting publication from that literature review. Section 6.2 discuss the future plans for the thesis.

6.1 Reviewed Mass Gathering Medical Care Literature

The beginning of the candidature was characterised by reviewing literature regarding MGMC. This was required in order to get a better understanding of the processes involved in provisioning medical care at such events, and to get a better insight into the problems which are faced. An invaluable resource for this was Prehospital and Disaster Medicine journal, which contains several studies about MGMC. Reading articles on this topic, and collaborating with Paul Arbon from Flinders University and Grace Rumantir who is a supervisor for this project led to the following publication:

> Peter Serwylo, Paul Arbon, and Grace Rumantir. Predicting patient presentation rates at mass gatherings using machine learning. In Maria A. Santos, Luísa Sousa, and Eliane Portela, editors, *ISCRAM 2011: Proceedings of the 8th International Conference on Information Systems for Crisis Response and Management*. ISCRAM, 2011.

This paper analysed data collected by the St Johns Ambulance in order to identify the factors which have the most impact on the PPR at mass gatherings. The data was previously analysed by Arbon (2002), by using linear regression techniques to try an use variables such as event type, attendance and temperature in order to predict the PPR. The current publication used machine learning techniques to achieve a similar goal, with the best model being able to predict the PPR at several events with about 80% accuracy (Serwylo et al., 2011).

The conference was held in Portugal during May, and the paper was well received. The day preceeding the main conference consisted of a PhD Colloquium in which the current thesis topic was presented to other PhD students, in addition to some of the people on the board of the ISCRAM organisation.

6.2 Timeline

Figure 4 shows the timeline for the expected completion of this thesis. It is broken into two main areas, generating the synthetic scenarios and then using those scenarios in a simulation for resource allocation. The timeline also includes a publication plan for the articles which will be written as part of this project. There may be more articles if the opportunity arises throughout the project.

During Semester 2, 2011, two enabling units have also been enrolled in: FIT4007 - Advanced topics in information systems and FIT4009 - Advanced topics in intelligent systems. FIT4007 relates to Information Systems research methods while FIT4009 has a strong focus on Bayesian Networks.

Glossary

Mass Gathering A situation (event) during which crowds gather and where there is the potential for a delayed response to emergencies because of limited access to patients or other features of the environment and location. This potential delay requires planning and preparation to limit (or mitigate) the hazards inherent in a mass gathering and ensure timely access to appropriate health care is available.

MGMC Mass Gathering Medical Care. In order to reduce the burden on the local EMS service, medical teams should be present to deal with minor complaints, and transfer more serious complaints to hospital.

Patient An individual presenting for care to on site first aid or medical services at the event.

PPR Patient Presentation Rate. The number of patients presenting for care at an event per 1,000 attendees.

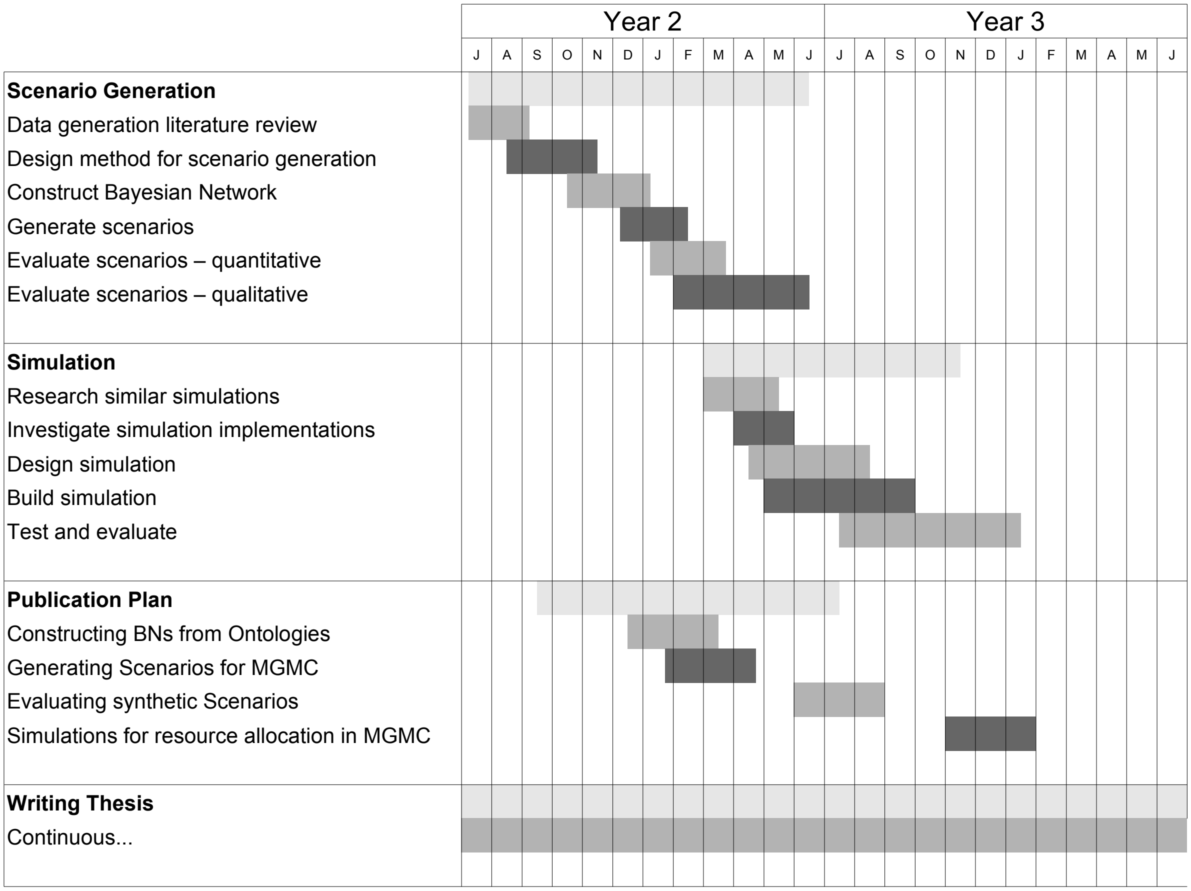

Figure 4: Timeline for thesis. Year 2 begins in July 2011, while the candidature finishes in July 2013.

References

Ashraf Aboulnaga, Jeffrey F. Naughton, and Chun Zhang. Generating synthetic complex-structured xml data. In *WebDB*, pages 79–84, 2001.

Wossenu Abtew, Rafael G. Moras, and K.L. Campbell. Synthetic precipitation data generation. *Computers & Industrial Engineering*, 19(1-4):582 – 586, 1990. ISSN 0360-8352. doi: DOI:10.1016/0360-8352(90)90185-O. URL `http://www.sciencedirect.com/science/article/pii/036083529090185O`.

G. Airy, T. Mullen, and J. Yen. Market based Adaptive Resource Allocation for Distributed Rescue Teams. In *Proceedings of the 6th International Information Systems for Crisis Response and Management (ISCRAM) Conference, Gothenburg, Sweden*, 2009.

P. Arbon. The development of a web-based algorithm for the prediction of patient presentation rates at mass gatherings. *Australian Journal of Emergency Management, The*, 17(1):60, 2002.

P. Arbon. The development of conceptual models for mass-gathering health. *Prehospital and Disaster Medicine*, 19:208–212, 2004.

P. Arbon. Mass-gathering medicine: A review of the evidence and future directions for research. *Prehospital and Disaster Medicine*, 22(2):131, 2007.

Matthew B. Baird, Robert E. O'Connor, Allen L. Williamson, Benjamin Sojka, Kostas Alibertis, and William J. Brady. The impact of warm weather on mass event medical need: a review of the literature. *The American Journal of Emergency Medicine*, 28(2):224 – 229, 2010. ISSN 0735-6757. doi: DOI:10.1016/j.ajem.2008.10.034. URL `http://www.sciencedirect.com/science/article/B6W9K-4YCGW64-S/2/e26b17d432554bebbfa7f3dfc0eefbd7`.

Fabio Bellifemine, Federico Bergenti, Giovanni Caire, and Agostino Poggi. Jade - a java agent development framework. In Gerhard Weiss, Rafael Bordini, Mehdi Dastani, JÃŒrgen Dix, and Amal Fallah Seghrouchni, editors, *Multi-Agent Programming*, volume 15 of *Multiagent Systems, Artificial Societies, and Simulated Organizations*, pages 125–147. Springer US, 2005. ISBN 978-0-387-26350-2. URL `http://dx.doi.org/10.1007/0-387-26350-0_5`.

C. Boldyreff, M. Dastbaz, H. Liu, and Y. Arafa. Engineering advanced training environment for crisis management: The pandora project. In *AC&T 2011: Proceedings of the 6th annual conference on Advances in Computing and Technology*, 2011.

Maged N. Kamel Boulos, Lee Hetherington, and Steve Wheeler. Second life: an overview of the potential of 3-d virtual worlds in medical and health education. *Health Information & Libraries Journal*, 24(4):233–245, 2007. ISSN 1471-1842. doi: 10.1111/j.1471-1842.2007.00733.x. URL `http://dx.doi.org/10.1111/j.1471-1842.2007.00733.x`.

Margaret T. Crichton, Rhona Flin, and William A. R. Rattray. Training decision makers - tactical decision games. *Journal of Contingencies and Crisis Management*, 8(4):208–217, 2000. ISSN 1468-5973. doi: 10.1111/1468-5973.00141. URL `http://dx.doi.org/10.1111/1468-5973.00141`.

Pari Delir Haghighi, Frada Burstein, Arkady Zaslavsky, Paul Arbon, and Shonali Krishnaswami. The role of domain ontology for medical emergency management in mass gatherings. In *DSS 2010 - the 15th IFIP WG8.3 International Conference on Decision Support Systems*, 2010.

D.J. DeWitt. The Wisconsin benchmark: Past, present, and future. *The Benchmark Handbook for Database and Transaction Processing Systems*, 1, 1991.

Jim Gray, Prakash Sundaresan, Susanne Englert, Ken Baclawski, and Peter J. Weinberger. Quickly generating billion-record synthetic databases. *SIGMOD Rec.*, 23:243–252, May 1994. ISSN 0163-5808. doi: http://doi.acm.org/10.1145/191843.191886. URL `http://doi.acm.org/10.1145/191843.191886`.

Nicholas Hartman, Allen Williamson, Benjamin Sojka, Kostas Alibertis, Marjorie Sidebottom, Thomas Berry, Jay Hamm, Robert E. O'Connor, and William J. Brady. Predicting resource use at mass gatherings using a simplified stratification scoring model. *The American Journal of Emergency Medicine*, 27(3):337 – 343, 2009. ISSN 0735-6757. doi: DOI:10.1016/j.ajem.2008.03.042. URL `http://www.sciencedirect.com/science/article/B6W9K-4VXMFRS-N/2/2b978fdc6ca38589345b974fec5b3578`.

Health and Great Britain Safety Executive. *The Event Safety Guide: A Guide to Health, Safety and Welfare at Music and Similar Events*. Guidance Booklet Series. HSE Books, 2 edition, 1999. ISBN 9780717624539. URL `http://books.google.com/books?id=DR6lAAAACAAJ`.

E.M. Helsper and L.C. Van der Gaag. Building Bayesian networks through ontologies. In *ECAI*, pages 680–684, 2002.

A.R. Hevner, S.T. March, J. Park, and S. Ram. Design science in information systems research. *Mis Quarterly*, 28(1):75–105, 2004. ISSN 0276-7783.

M. K. James and L. B. Mason. Synthetic tropical cyclone database. *Journal of Waterway, Port, Coastal, and Ocean Engineering*, 131(4):181–192, 2005. doi: 10.1061/(ASCE)0733-950X(2005)131:4(181). URL `http://link.aip.org/link/?QWW/131/181/1`.

N.E. Kman, G.B. Russell, W.P. Bozeman, K. Ehrman, and J. Winslow. Derivation of a Formula to Predict Patient Volume Based on Temperature at College Football Games. *Prehospital Emergency Care*, 11(4):453–457, 2007. ISSN 1090-3127.

K.B. Korb and A.E. Nicholson. *Bayesian artificial intelligence*, volume 1. cRc Press, 2004.

Salvatore T. March and Gerald F. Smith. Design and natural science research on information technology. *Decision Support Systems*, 15(4):251 – 266, 1995. ISSN 0167-9236. doi: DOI:10.1016/0167-9236(94)00041-2. URL `http://www.sciencedirect.com/science/article/B6V8S-3Y5FNDY-1/2/2ea948195f581aa5f63b92459958d987`.

D. McGrath and D. Hill. Unrealtriage: A game-based simulation for emergency response. In *The Huntsville Simulation Conference, October*, 2004.

A.M. Milsten, B.J. Maguire, RA Bissell, and KG Seaman. Mass-gathering medical care: a review of the literature. *Prehospital and Disaster Medicine*, 17 (3):151–162, 2002.

Department of Defence. Department of Defence Modeling and Simulation Body of Knowledge, 2011. URL `http://www.msco.mil/`.

T.E. Raghunathan, J.P. Reiter, and D.B. Rubin. Multiple imputation for statistical disclosure limitation. *Journal of Official Statistics - Stockholm*, 19(1): 1–16, 2003. ISSN 0282-423X.

J.P. Reiter. Satisfying disclosure restrictions with synthetic data sets. *JOURNAL OF OFFICIAL STATISTICS-STOCKHOLM-*, 18(4):531–544, 2002. ISSN 0282-423X.

WD Rose, SC Laird, JE Prescott, et al. Emergency medical services for collegiate football games: A six and one half year review. *Prehospital and Disaster Medicine*, 7:157–9, 1992.

D.B. Rubin. Statistical disclosure limitation. *Journal of official Statistics*, 9(2): 461–468, 1993.

D.A. Schön. *The reflective practitioner: how professionals think in action*. Basic Books, 1983. ISBN 9780465068784. URL `http://books.google.com/books?id=ceJIWay4-jgC`.

Mikhail A. Semenov and Elaine M. Barrow. Use of a stochastic weather generator in the development of climate change scenarios. *Climatic Change*, 35:397–414, 1997. ISSN 0165-0009. URL `http://dx.doi.org/10.1023/A:1005342632279`. 10.1023/A:1005342632279.

Peter Serwylo, Paul Arbon, and Grace Rumantir. Predicting patient presentation rates at mass gatherings using machine learning. In Maria A. Santos, Luísa Sousa, and Eliane Portela, editors, *ISCRAM 2011: Proceedings of the 8th International Conference on Information Systems for Crisis Response and Management*. ISCRAM, 2011.

H.A. Simon. *The sciences of the artificial*. The MIT Press, 1996. URL `http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=5579`.

W.P. Smith, V. Wessels, D. Naicker, E. Leuenberger, P. Fuhri, L.A. Wallis, and F.S.A. FCEM. Development of a Mass-Gathering Medical Resource Matrix for a Developing World Scenario. *Prehospital and disaster medicine: the official journal of the National Association of EMS Physicians and the World Association for Emergency and Disaster Medicine in association with the Acute Care Foundation*, 25(6):547, 2010. ISSN 1049-023X.

St Johns Ambulance. URL `http://www.stjohn.org.au/`.

Fonny Sujanto, Andrzej Ceglowski, Frada Burstein, and Leonid Churilov. An integrated framework for comprehensive collaborative emergency management. In *Proceeding of the 2008 conference on Collaborative Decision Making: Perspectives and Challenges*, pages 127–138, Amsterdam, The Netherlands, The Netherlands, 2008. IOS Press. ISBN 978-1-58603-881-6. URL `http://portal.acm.org/citation.cfm?id=1563893.1563909`.

Latanya Sweeney. k-anonymity: a model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 10:557–570, October 2002. ISSN 0218-4885. doi: 10.1142/S0218488502001648. URL http://portal.acm.org/citation.cfm?id=774544.774552.

Sui-Yu Wang, Yuanbo Guo, Abir Qasem, and Jeff Heflin. Rapid benchmarking for semantic web knowledge base systems. In Yolanda Gil, Enrico Motta, V. Benjamins, and Mark Musen, editors, *The Semantic Web - ISWC 2005*, volume 3729 of *Lecture Notes in Computer Science*, pages 758–772. Springer Berlin / Heidelberg, 2005. doi: 10.1007/11574620_54. URL http://dx.doi.org/10.1007/11574620_54.

Mark A. Whiting, Jereme Haack, and Carrie Varley. Creating realistic, scenario-based synthetic data for test and evaluation of information analytics software. In *Proceedings of the 2008 conference on BEyond time and errors: novel evaLuation methods for Information Visualization*, BELIV '08, pages 8:1–8:9, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-016-6. doi: http://doi.acm.org/10.1145/1377966.1377977. URL http://doi.acm.org/10.1145/1377966.1377977.

K. Yazawa, Y. Kamijo, R. Sakai, M. Ohashi, M. Owa, L.T.C.C.R.A. DeLorenzo, C.J. Bonnett, T.R. Schock, K.E. McVaney, C.B. Colwell, et al. Medical Care for a Mass Gathering: The Suwa Onbashira Festival. *Prehospital Disaster Medicines*, 22(5):431, 2007.

K.M. Zeitz, D. Schneider, and BN Dannielle Jarrett. Mass gathering events: Retrospective analysis of patient presentations over seven years at an agricultural and horticultural show. *Prehospital and Disaster Medicine*, 585:1–14, 2001.

K.M. Zeitz, C.J. Zeitz, and P. Arbon. Forecasting medical work at mass-gathering events: predictive model versus retrospective review. *Prehospital and Disaster Medicine*, 20:164–168, 2005.