UNIVERSITY OF AMSTERDAM

Machine Learning for Relevance of Information in Crisis Response

C.P.M. Netten

SUMMARY

Efficient communication during crisis response situations is a major challenge for involved emergency responders. Disaster studies show that missing information and information overload are important factors that determine the course of crisis response. It is essential for emergency responders to acquire all information critical to their task execution in time, especially at the beginning of a crisis response situation. Lack of relevant information or too much irrelevant information hampers the emergency responders' decision-making process, workflow and situational awareness.

Crisis response aims to optimise information distribution by engaging in intense training exercises, focussing on familiarity with emergencies, multidisciplinary activities, and resources. In large crisis response operations, an information manager is added to the operational command. This person is responsible for collecting and distributing all information relevant for the common operational picture and for keeping an up-to-date view of the crisis. Despite these efforts to better centralize relevant information during crisis response, a gap still exists between the information supply and information needs of responders.

Our contribution to bridge the information gap is a software system that monitors communication and may send information to emergency responders that were not addressed in the initial communication. The system, Task-Adaptive Information Distribution (TAID) is capable of disseminating information in a timely manner and adapting itself to the fast-changing information needs in a crisis response environment. The TAID system was trained with practical examples for which information relevance is known. Relevance of information was assessed on the basis of workflow information, such as emergency responder task (i.e. task descriptions assigned to roles of emergency responders) and location at a given time. This type of information is provided to TAID by auxiliary information systems (e.g. GPS).

To assess relevance, TAID uses a built relevance model for crisis response using methods from machine learning. In TAID this was set up as a classification task in which input information and knowledge about (ir)relevance of information was used. Machine learning methods are capable of solving these types of complex tasks by coming up with their own program (i.e. model) based on examples provided to them. The alternative to machine learning would be to program (i.e. model and build) such a system directly by hand. Such a hand-crafted approach is rather complex and time-consuming, since it takes much effort to define and maintain all relevance rules for crisis response.

The input for the classification task was communication between emergency responders and the situational information of all responders (i.e. task description and distances to other responders). The output is a decision pertaining to (ir)relevance of the information for each responder, who did not take part in the initial communication. We investigated the feasibility of our machine learning method to assess relevance of information in crisis response. Important generic attributes for assessing the relevance of information were identified and used to build a relevance model, for example the contextual information of responders in relation to what has been communicated.

A vast amount of the information used as input for the learning method consisted of spoken dialogues converted to text. This meant the text messages used to assess as relevant were relatively short. Therefore, we investigated whether message segments could be automatically detected from the flow of dialogue communication, which in turn would yield larger and more comprehensible text messages. For the detection task we again used methods form machine learning to build a model. This model was trained with examples that indicated at which location in the flow of dialogue communication were the transitions from one segment to another. These detected segments improved the classification task for relevance and yielded more comprehensible messages to

forward. All models were built based on data from emergency response practice and were technically evaluated.

We also investigated which effects the relevant information distributions of TAID could have on emergency response. Since evaluation in a real crisis response setting was not possible, we modelled in Brahms a crisis response scenario based on a real exercise, and embedded the TAID system in the simulation model. Brahms proved to be an appropriate environment to model and simulate crisis response work practice. Based on the simulation-experiments we could measure the TAID-effect on crisis response and give an impression of how big their impact is, and a sense of how TAID's effect seems to play out. Our analysis of these effects focus on the time spent by the emergency responders on executing tasks and communicating information. The technical evaluation results of the built relevance models are promising, and the simulation experiments show that the expected effect of the TAID system on crisis response is restricted to specific situations. Only a small number of cases within the used scenario show clear impact of the information distributions of TAID. In spite of this sporadic impact of TAID, it results in a consistent improvement in the situational awareness of emergency responders, who have more access to relevant information than before. In a number of cases the intervening of TAID reduced the duration of the crisis with very low additional costs, which results from time taken due to reading additional messages.