A Virtual Testbed for Critical Incident Investigation with Autonomous Remote Aerial Vehicle Surveying, Artificial Intelligence, and Decision Support

David L. Smyth, Sai Abinesh, Nazli B. Karimi,
Brett Drury, Ihsan Ullah, Frank G. Glavin, Michael G. Madden

School of Computer Science, National University of Ireland Galway, Ireland
michael.madden@nuigalway.ie

Abstract. Autonomous robotics and artificial intelligence techniques can be used to support human personnel in the event of critical incidents. These incidents can pose great danger to human life. Some examples of such assistance include: multi-robot surveying of the scene; collection of sensor data and scene imagery, real-time risk assessment and analysis; object identification and anomaly detection; and retrieval of relevant supporting documentation such as standard operating procedures (SOPs). These incidents, although often rare, can involve chemical, biological, radiological/nuclear or explosive (CBRNE) substances and can be of high consequence. Real-world training and deployment of these systems can be costly and sometimes not feasible. For this reason, we have developed a realistic 3D model of a CBRNE scenario to act as a testbed for an initial set of assisting AI tools that we have developed.[1]

1 Background and Related Research

We have developed a bespoke virtual environment (VE) model of a critical incident using a state-of-the-art games engine. We use this model to test a range of assisting AI technologies related to information gathering, real-time analytics and decision support.

We developed the VE with the core purpose of using it as a testbed for the development of a range of investigation assisting AI tools. VEs have also been used to train first responder personnel in near photo-realistic yet safe conditions. Chroust and Aumayr [2] note that virtual reality can support training by allowing simulations of potential incidents, as well as the consequences of various courses of action, in a realistic way. There are virtual reality training systems which solely focus on CBRN disaster preparedness. Some of these are outlined by Mossel *et al.* [9]. Other example uses of virtual worlds include *Second Life* and *Open Simulator* [4,3].

CBRNE incident assessment is a critical task which poses significant risks and endangers the lives of human investigators. For this reason, many research projects focus on the use of robots such as Micro Unmanned Aerial Vehicles (MUAV) to carry out remote sensing in such hazardous environments [7,1]. Others can include CBRNE mapping for first responders [6] and multi-robot reconnaissance for detection of threats [12].

---

2   A Virtual Testbed for Critical Incidents

We have developed and implemented a baseline set of decision support systems for investigating critical incidents. In order to test these in an efficient and cost effective manner, we have developed 3D world models of typical CBRNE incidents using a physics-based game engine. These models include virtual representations of Robotic Aerial Vehicles (RAVs).

After identifying the area of interest, multiple RAVs are deployed to survey the scene. The RAVs, which are fitted with sensors and cameras, operate as a multi-agent robot swarm and divide the work up between them. All information is relayed to a central hub in which our Image Analysis module uses a Deep Neural Network (DNN) to detect and identify relevant objects in images taken by RAV cameras. It also uses a DNN to perform pixel-level semantic annotation of the terrain, to support subsequent route-planning for Robotic Ground-based Vehicles (RGVs). Our Probabilistic Reasoning module assesses the likelihood of different threats, as information arrives from the scene commander, survey images and sensor readings. Our Information Retrieval module ranks documentation, using TF-IDF, by relevance to the incident. All interactions are managed by our purpose-built JSON-based communications protocol, which is also supported by real-world RAVs, cameras and sensor systems. This keeps the system loosely coupled, and will support future testing in real-world environments.

This work was undertaken as part of a project called ROCSAFE (Remotely Operated CBRNE Scene Assessment and Forensic Examination) and this demonstration overview is based on Smyth *et al.* [14].

2.1   Modelling a Critical Incident Scenario

To facilitate the development and testing of our AI tools, we have designed, developed and publicly released a VE [15] using the *Unreal Engine* (UE). This is a suite of tools for creating photo-realistic simulations with accurate real-world physics. UE is open source, scalable and supports plugins that allow the integration of RAVs and RGVs into the environment. For this demonstration, we chose an operational scenario to model that consists of a train carrying radioactive material in a rural setting. We used Microsoft's *AirSim* [13] plugin to model the RAVs. AirSim exposes various APIs to allow fine-grain control of RAVs, RGVs and their associated components. We have replicated a number of APIs from real-world RAV and RGV systems to facilitate the application of our AI tools to real-world critical incident use-cases in the future, after firstly testing them in the VE.

2.2   Communications

A secure purpose-built JSON-format protocol was developed for the communications between subsystems. We used a RESTful API because of the fewer number of messages at pre-defined intervals [11]. The communication protocol not only provides autonomy to several vehicles but it is also flexible enough to integrate with various components using different standards, protocols and data types. In this demonstration, we concentrate on RAVs. Since decision making may happen within each RAV's single-board computer, we have also facilitated direct communication between the RAVs.

2.3 Autonomous Surveying and Image Collection

Our multi-agent system supports the autonomous mapping of the virtual environment. It involves discretizing a rectangular region of interest into a set of grid points. At each point, the RAV records a number of images and metadata. Four bounding GPS coordinates (corner points of a rectangle) can be passed in through a web-based user interface.

Our planning algorithm develops agent routes at a centralized source and distributes the planned routes to each agent in the multi-agent system [16]. A greedy algorithm is used in the current implementation to generate subsequent points in each agent's path by minimizing the distance each agent needs to travel to an unvisited grid point. Current state-of-the-art multi-agent routing algorithms use hyper-heuristics, which out-perform algorithms that use any individual heuristic [19]. We intend to integrate this approach with learning algorithms such as Markov Decision Processes [18] in order to optimize the agent routes in a stochastic environment, for-example where RAVs can fail and battery usage may not be fully known.

2.4 Image Processing and Scene Analysis

Our Central Decision Management (CDM) system uses the object labels predicted by a deep neural network from images taken by the RAV cameras. Specifically, we fine-tuned an object detection model *Mask R-CNN* [5] with our annotated synthetic images that we collected from the virtual scene. Training on a synthetic dataset has been shown to transfer well to real world data in self-driving cars [10] and object detection [17].
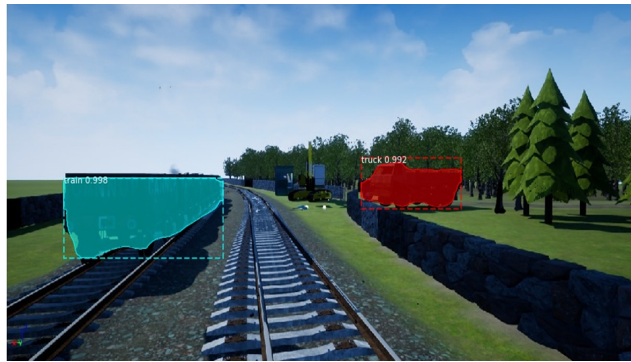

Fig. 1: Object identification from a virtual scene image.

Mask R-CNN is currently a state-of-the-art object detection deep model that detects and localizes objects with bounding boxes and provides overlay instance segmentation masks to show the contours of the objects within the boxes. Figure 1 shows the identification of a train and a truck from our virtual scene. The objective of using this detection algorithm is to highlight objects of interest within the scene to the crime scene investigator's attention. These models can detect objects even if they are overlapping. The predicted labels that are produced are an input for our probabilistic reasoning module. Currently, we are enhancing the performance of this deep learning model by

retraining/fine-tuning the network on other relevant datasets, for example, Object de-Tection in Aerial (DOTA) images [20]. In addition, our plan is to also detect anomalies in the scenes.

2.5 Reasoning and Information Retrieval

We have developed a probabilistic model in the BLOG language [8]. It synthesizes data and reasons about the threats in the scene over time. The objective is to estimate the probabilities of different broad categories of threat (chemical, biological, or radiation/nuclear) and specific threat substances. This information affects the way a scene is assessed. For example, a first responder with a hand-held instrument may initially detect evidence of radiation in some regions of the scene. Subsequent RAV images may then show damaged vegetation in those and other regions, which could be caused by radiation or chemical substances. Another source of information could come from RAVs dispatched with radiation sensors that fly low over those regions. Using keywords that come from sources such as the object detection module, the probabilistic reasoning module, and the crime scene investigators, the CDM retrieves documentation such as standard operating procedures and guidance documents from a knowledge base. This retrieval is done based on rankings (in order of relevance to the current situation). Elastic Search and a previously-defined set of CBRNE synonyms are used for rankings. The documents are re-ranked in real-time as new information becomes available from various sources.

References

1. A Baums. Response to cbrne and human-caused accidents by using land and air robots. *Automatic Control and Computer Sciences*, 51(6):410–416, 2017.
2. Gerhard Chroust and Georg Aumayr. Resilience 2.0: Computer-aided disaster management. *Journal of Systems Science and Systems Engineering*, 26(3):321–335, 2017.
3. Daniel Cohen, Nick Sevdalis, Vishal Patel, Michael Taylor, Henry Lee, Mick Vokes, Mick Heys, David Taylor, Nicola Batrick, and Ara Darzi. Tactical and operational response to major incidents: feasibility and reliability of skills assessment using novel virtual environments. *Resuscitation*, 84(7):992–998, 2013.
4. Daniel Cohen, Nick Sevdalis, David Taylor, Karen Kerr, Mick Heys, Keith Willett, Nicola Batrick, and Ara Darzi. Emergency preparedness in the 21st century: training and preparation modules in virtual environments. *Resuscitation*, 84(1):78–84, 2013.
5. Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *Conference on Computer Vision (ICCV)*, pages 2980–2988, 2017.
6. Piotr Jasiobedzki, Ho-Kong Ng, Michel Bondy, and CH McDiarmid. C2SM: a mobile system for detecting and 3d mapping of chemical, radiological, and nuclear contamination. In *Sensors, and Command, Control, Communications, and Intelligence (C3I)*, 2009.
7. Mario Monteiro Marques, Rodolfo Santos Carapau, Alexandre Valério Rodrigues, V Lobo, Júlio Gouveia-Carvalho, Wilson Antunes, Tiago Gonçalves, Filipe Duarte, and Bernardino Verissimo. Gammaex project: A solution for cbrn remote sensing using unmanned aerial vehicles in maritime environments. In *OCEANS–Anchorage*, pages 1–6. IEEE, 2017.
8. Brian Milch, Bhaskara Marthi, Stuart Russell, David Sontag, Daniel L Ong, and Andrey Kolobov. BLOG: Probabilistic models with unknown objects. *Statistical relational learning*, page 373, 2007.

9. Annette Mossel, Andreas Peer, Johannes Göllner, and Hannes Kaufmann. Requirements analysis on a virtual reality training system for cbrn crisis preparedness. In *59th Annual Meeting of the ISSS*, volume 1, pages 928–947, 2017.

10. Xinlei Pan, Yurong You, Ziyan Wang, and Cewu Lu. Virtual to real reinforcement learning for autonomous driving. *arXiv:1704.03952*, 2017.

11. Leonard Richardson, Mike Amundsen, and Sam Ruby. *RESTful Web APIs*. O'Reilly Media, Inc., 2013.

12. Frank E Schneider, Jochen Welle, Dennis Wildermuth, and Markus Ducke. Unmanned multi-robot cbrne reconnaissance with mobile manipulation system description and technical validation. In *13th International Carpathian Control Conference (ICCC)*, pages 637–642. IEEE, 2012.

13. Shital Shah, Debadeepta Dey, Chris Lovett, and Ashish Kapoor. Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and Service Robotics*, pages 621–635, 2017.

14. David L Smyth, James Fennell, Sai Abinesh, Nazli B Karimi, Frank G Glavin, Ihsan Ullah, Brett Drury, and Michael G Madden. A virtual environment with multi-robot navigation, analytics, and decision support for critical incident investigation. *IJCAI, 27th International Joint Conference on Artificial Intelligence. arXiv:1806.04497*, 2018.

15. David L Smyth, Frank G Glavin, and Michael G Madden. UE4 Virtual Environment: Rural Rail Radiation Scenario. In *https://github.com/ROCSAFE/CBRNeVirtualEnvMultiRobot/releases*, 2018.

16. David L Smyth, Frank G Glavin, and Michael G Madden. Using a game engine to simulate critical incidents and data collection by autonomous drones. In *IEEE Games, Entertainment and Media*, 2018.

17. Yonglin Tian, Xuan Li, Kunfeng Wang, and Fei-Yue Wang. Training and testing object detectors with virtual images. *Journal of Automatica Sinica*, 5(2):539–546, 2018.

18. Marlin W Ulmer, Justin C Goodson, Dirk C Mattfeld, and Barrett W Thomas. Route-based markov decision processes for dynamic vehicle routing problems. Technical report, Braunschweig, 2017.

19. Yue Wang, Min-Xia Zhang, and Yu-Jun Zheng. A hyper-heuristic method for uav search planning. In *Advances in Swarm Intelligence*, pages 454–464, 2017.

20. Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. DOTA: A large-scale dataset for object detection in aerial images. *arXiv:1711.10398*, 2017.