

Towards Distributed Robotic Casualty Assessment using Multimodal, Non-Contact Perception and Probabilistic Inference

Zachary Bortoff¹, Srijal Shekhar Poojari², Kleio Baxevani¹, Joshua Gaus³,
Christopher Titus³, Ahmed Ashry¹, and Derek A. Paley⁴

Abstract—Mass-casualty incidents demand rapid and accurate triage, but the scale and acuity of injuries often overwhelm available medical personnel. To address this, we present a system that enables ground and aerial robots to localize and assess casualties using non-contact sensors, including color and thermal cameras, millimeter wave radar, and microphones. Injury and vital sign measurements from modality-specific classifiers are fused using a probabilistic model that captures correlations between injury states and supports distributed, asynchronous evidence accumulation. We validate the system through a series of timed mass-casualty field experiments using custom-built drones and Boston Dynamics Spot ground robots customized for robotic medical triage, demonstrating reliable estimation of casualty states and robustness to noisy conditions and sensor drop out.

I. INTRODUCTION

Prioritizing casualties based on injury severity is essential to saving lives during mass casualty incidents (MCIs), where the number and acuteness of casualties leave medical personnel overwhelmed [1]. The potential to scale triage operations by providing medics with a casualty’s location and injury severity is made possible due to advances in non-contact vitals sensors, image and audio classifiers, and ground and aerial robots.

However, deploying robotic platforms in real-world triage settings introduces challenges. Injuries may manifest across different sensing modalities, necessitating the integration and fusion of multiple sensor streams. Environmental noise such as smoke obscuring thermal signatures, poor lighting degrading color camera imagery, and loud ambient sounds masking casualty vocalizations significantly reduces sensor data quality. At the same time, failing to estimate injuries accurately risks not prioritizing life-saving care appropriately.

Addressing these challenges requires principled models that integrate heterogeneous data sources and provide calibrated estimates of injury likelihood and triage urgency. The

*This work was supported by Defense Advanced Research Projects Agency (DARPA) Grant No. HR00112420304. Approved for public release; distribution is unlimited.

¹Zachary Bortoff, Kleio Baxevani, and Ahmed Ashry are with the Department of Aerospace Engineering, University of Maryland, College Park, MD, USA. {zbortoff, kleio, aashry}@umd.edu

²Srijal Poojari is with the Department of Electrical and Computer Engineering, University of Maryland, College Park, MD, USA. srijal@umd.edu

³Joshua Gaus and Christopher Titus are with University of Maryland UAS Research and Operations Center, California, MD, USA. {jgaus, ctitus}@umd.edu

⁴Derek A. Paley is with the Department of Aerospace Engineering and the Institute for Systems Research, University of Maryland, College Park, MD, USA. dpaley@umd.edu



Fig. 1: (Top) Three customized Boston Dynamics Spots and one custom drone performing medical triage of casualties in simulated mass-casualty incident. (Bottom) Spot and drone injury assessments of same casualty.

System Competition Track of the DARPA Triage Challenge (DTC) aims to promote the development of robotic platforms that can assist emergency medical services in quickly and accurately performing triage of casualties during MCIs [2]. This paper presents a system for assessing casualties developed by our team, RoboScout DTC, for participation in the second year of the three-year Challenge.

Recent work has demonstrated the feasibility of deploying deep learning models on aerial and ground robots for contact-free injury classification and vital sign estimation using color and infrared cameras, microphones, and millimeter-wave (mmWave) radar. Carrion et al. proposes an object detection model for detecting wounds in color imagery for use in medical treatment facility (MTF) settings [3]. West et al. shows that simulated blood is visible from a thermal image taken by a drone, but do not propose a method of automatically detecting blood [4]. Saeed et al. compares two machine-learning models for human scream detection deployed on a mobile ground robot for use in fire emergencies [5]. Zhao et al. and Chen et al. each propose algorithms for using a mmWave radar sensor deployed on a robot to estimate heart rate [6], [7]. While these approaches enable medical assessments by robotic platforms, the challenges