

Title	Monitoring emergency first responders' activities via gradient boosting and inertial sensor data
Authors	Scheurer, Sebastian; Tedesco, Salvatore; Manzano, Óscar; Brown, Kenneth N.; O'Flynn, Brendan
Publication date	2019-01-18
Original Citation	Scheurer S., Tedesco S., Manzano Ò., Brown K.N., O'Flynn B. (2019) Monitoring Emergency First Responders' Activities via Gradient Boosting and Inertial Sensor Data. In: Brefeld U. et al. (eds) Machine Learning and Knowledge Discovery in Databases. ECML PKDD 2018. Lecture Notes in Computer Science, vol 11053. Springer, Cham, pp. 691-694. doi: 10.1007/978-3-030-10997-4_53
Type of publication	Conference item
Link to publisher's version	https://link.springer.com/chapter/10.1007/978-3-030-10997-4_53 - 10.1007/978-3-030-10997-4_53
Rights	© Springer Nature Switzerland AG 2019. This is a post-peer-review, pre-copyedit version of an article published in Lecture Notes in Computer Science. The final authenticated version is available online at: http://dx.doi.org/10.1007/978-3-030-10997-4_53
Download date	2024-06-23 11:45:41
Item downloaded from	https://hdl.handle.net/10468/9649



UCC

University College Cork, Ireland
Coláiste na hOllscoile Corcaigh

Monitoring Emergency First Responders’ Activities via Gradient Boosting and Inertial Sensor Data

Sebastian Scheurer^{1,3}[0000–0002–2163–2168], Salvatore Tedesco², Òscar Manzano², Kenneth N. Brown¹, and Brendan O’Flynn²

¹ Insight Centre for Data Analytics, Dep. of Computer Science,

² Tyndall National Institute,

University College Cork, Cork, Ireland.

³ sebastian.scheurer@insight-centre.org

Abstract. Emergency first response teams during operations expend much time to communicate their current location and status with their leader over noisy radio communication systems. We are developing a modular system to provide as much of that information as possible to team leaders. One component of the system is a human activity recognition (HAR) algorithm, which applies an ensemble of gradient boosted decision trees (GBT) to features extracted from inertial data captured by a wireless-enabled device, to infer what activity a first responder is engaged in. An easy-to-use smartphone application can be used to monitor up to four first responders’ activities, visualise the current activity, and inspect the GBT output in more detail.

Keywords: Human Activity Recognition · Machine Learning · Boosting · Inertial Sensors

1 Introduction

Emergency first responders, such as firefighters, typically enter a building in pairs or small teams, each of which maintains contact to its leader using the portable radio communications device carried by one of its members. As a team makes its way into the building they report their status and progress to, and confirm their next steps with, their leader who tracks their location and status, usually on paper, at every point of the way. Because radio is a notoriously noisy medium, messages have to be repeated frequently to ensure accurate communication, but every minute taken up by these communications adds another minute to the time a person in distress is waiting for help.

The Sensor Technologies for Enhanced Safety and Security of Buildings and its Occupants (SAFESENS) project [8] is developing a novel monitoring system for emergency first responders designed to provide first response team leaders with timely and reliable information about their team’s status during emergency response operations, thereby reducing the amount of time taken up by

radio communications, accelerating response operations, and improving first responders’ safety. The system consists of components for monitoring vital signs, indoor localisation, and human activity recognition (HAR), each with its own set of sensors—vital signs are captured by an instrumented glove, localisation uses ranging data from pre-deployed ultra-wideband (UWB) anchors [1,3], and the HAR component relies on the data captured by the inertial measurement units (IMU) worn by first responders—and related algorithms. A video which demonstrates the execution of the system can be found at bit.ly/iharpkdd.

There is a substantial body of HAR work using wearable IMUs for assisted living, industrial, or sports applications [7], but these do not necessarily translate to first response operations. The few commercial first responder monitoring systems that do exist, such as Medtronic’s Zephyr™ offerings [4], rely on a host of wearable sensors built into vests, boots, or gloves, which—in addition to the fact that first responders are not keen to add to their already cumbersome equipment—tend to be expensive. Our system demonstrates how the state-of-the-art in HAR can be used to monitor emergency first responders’ activities using only one wearable device per first responder. This system can be customised to recognise a different set of activities by re-training the classifier with appropriate sample data.

2 System Architecture and Operation




Fig. 1. System architecture. Note only one first responder plus equipment (WIMU and Smartphone), team leader, and client application are shown.

The system architecture (fig. 1) shows the different components and how they can be used during a first response operation. Each first responder wears a wireless-enabled IMU (WIMU) that captures inertial data and transmits them via low energy bluetooth (BLE) to the smartphone carried by each first responder. The phone posts the data in batches of configurable duration (default: 10 s) via HTTP to an API exported by the SAFESSENS server, which is running on a PC in the Command & Control Centre (CCC). The WIMU also connects to any reachable UWB anchors, which compute ranging data for each connected WIMU, and posts them to the API. The API receives and stores both (inertial and ranging) types of data in a relational database, where they are available to client applications as well as the localisation and HAR algorithm.

3 Recognising and Monitoring First Responder Activities




Fig. 2. Screenshots of the team leaders' application showing a) the dashboard, and b) the screen with details of the current HAR estimate.

In this demo we shall limit ourselves to the HAR component, for which we have developed an Android application whose initial screen is shown in panel a) of fig. 2. This screen illustrates the status of up to four first responders, showing, for each of them, the most recent activity, according to the output of the HAR algorithm, and highlighting it if and when appropriate—e.g., to issue a “firefighter down” alert if a firefighter is thought to have ceased moving and be lying on the ground. Team leaders can tap on a first responder’s icon to access a second screen (panel b) that provides more detail about the HAR estimate, and illustrates the most likely activity by means of a 3D model.

The HAR component, once triggered, operates as follows. First, the most recent batch, by default covering 10 s of inertial data, is loaded from the database. Then, the signals are resampled to their mean sampling frequency, using linear interpolation to fill any gaps that might have been introduced by the resampling, before a moving median filter is applied over a window of 3 samples to smooth the signals. Next, the smoothed acceleration signals are separated by means of a low pass filter [2] into their respective gravity and body components, which thenceforth replace the original accelerometer signals. After this, the two components and the smoothed gyroscope signals are segmented into 3 s sliding

windows with 1 s overlap, and a set of time- and frequency-domain features is extracted from each window. Finally, the extracted features are passed to the HAR inference algorithm, and the resulting probability estimates for each of the target activities averaged, producing a final estimate for the batch which is returned as response to the client’s API request.

The HAR inference algorithm is a gradient boosted ensemble of decision trees (GBT) which has been trained, for the sake of this demo, to recognise six activities, namely standing/sitting, crawling on hands & knees or stomach, walking, falling, and lying down. The GBT hyper-parameters, such as the number of iterations (750), learning rate (0.02), or maximum depth (3) of the trees, have been tuned via leave-one-subject-out cross-validation to minimise the average mean absolute error (MAE) across the target classes. More details on the training data, pre-processing, feature extraction, as well as tuning, training, and evaluation of the GBT can be found in [6] and [5], where this approach achieved a MAE of < 4% and Accuracy of > 90% when evaluated on data from an unseen individual.

Acknowledgements

This publication emanated from research supported by research grants from Science Foundation Ireland (SFI) and the European Development Fund under grant numbers SFI/12/RC/2289 and 13/RC/2077-CONNECT, and the European funded project SAFESSENS under the ENIAC program in association with Enterprise Ireland (IR20140024).

References

1. Brahmi, I.H., Abruzzo, G., Walsh, M., Sedjelmaci, H., O’Flynn, B.: A fuzzy logic approach for improving the tracking accuracy in indoor localisation applications. In: *Wireless Days Conf. IEEE* (4 2018)
2. Karantonis, D.M., Narayanan, M.R., Mathie, M., Lovell, N.H., Celler, B.G.: Implementation of a real-time human movement classifier using a triaxial accelerometer for ambulatory monitoring. *Trans. Inf. Technol. Biomed.* **10**(1) (1 2006)
3. Khodjaev, J., Tedesco, S., O’Flynn, B.: Improved NLOS error mitigation based on LTS algorithm. *Progress In Electromagnetics Research Letters* **58** (2016)
4. Medtronic: Zephyr™ performance systems for first responders and industrial safety. online (2017), <https://www.zephyranywhere.com>
5. Scheurer, S., Tedesco, S., Brown, K.N., O’Flynn, B.: Human activity recognition for emergency first responders via body-worn inertial sensors. In: *Int. Conf. Wearable & Implant. Body Sensor Netw. IEEE* (5 2017)
6. Scheurer, S., Tedesco, S., Brown, K.N., O’Flynn, B.: Sensor and feature selection for an emergency first responders activity recognition system. In: *Sensors. IEEE* (10 2017)
7. Sreenivasan, R.R., Nirmalya, R.: Recent trends in machine learning for human activity recognition—a survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **8**(4) (3 2018)
8. Tedesco, S., Khodjaev, J., O’Flynn, B.: A novel first responders location tracking system: Architecture and functional requirements. In: *Mediterranean Microw. Symp. IEEE* (11 2015)