Developing an automatic threat detection system for aviation security

On November 10, 1972, three hijackers threatened to fly Southern Airways Flight 49 into a nuclear reactor at Oak Ridge National Laboratory. As a direct response to this incident, the Federal Aviation Administration in the United States required all airlines to begin screening passengers and their carry-on baggage by January 5, 1973. A year later, the Air Transportation Security Act sanctioned the FAA’s universal screening rule, forcing U.S. airports to adopt metal-detection screening portals for passengers and X-ray inspection systems for carry-on bags.

Since then, using X-ray machines for inspecting passengers' luggage is a widely used method, as it can be both fast and highly effective for detecting threats. The common way of deploying this technology is by employing and training people to inspect the produced two-dimensional X-ray images. Although this has been a standard method for detecting prohibited items for many years, it also meets several limitations. Since it relies on individuals' abilities for recognising potentially prohibited items, different operators can achieve varying performances in identifying such items based on their visual knowledge acquired through experience and training. Additionally, a two-dimensional projection of an object captured from the X-ray scanner can differ significantly and be unfamiliar compared to its shape as stored in the visual memory of the screener. Other external factors, such as emotional exhaustion and lack of job satisfaction, can also negatively impact decision time and detection performance.

As the global aviation traffic grows and the threat landscape evolves, the above issues necessitate the development of automatic threat detection (ATD) systems, that will be capable to increase the detection performance and response time, enhancing passengers' transit experience, while increasing airports' security. Recently, machine learning approaches using deep convolutional neural networks (CNNs) have been successfully employed to 2D X-ray image data, for assisting screeners with real-time threat detection. However, the nascent introduction of CT (computerized tomography) scanners for carry-on baggage, which are able to capture 3D representations of the inspected bags, shifts the requirement for ATD capabilities from working with 2D image data to 3D volumetric datasets. This transition involves developing new algorithms that can cope with large amounts of data in short time frames. Moreover, employing 3D deep learning approaches in this field can be particularly challenging, considering the limited availability of annotated data. Commonly, the top-performing deep learning algorithms for recognising images, are trained on datasets of millions of labelled images, however, since data from CT scanners can be expensive and time-consuming to be acquired and labelled, most datasets are private and limited to a few thousand samples.

In order for an automatic threat detection algorithm to be applicable for the airport’s needs, it needs to be: time efficient, so it can inspect bags in only a few seconds, accurate, for detecting threats, while having a low rate of false alarms, and capable to be trained with limited datasets. In this research, we propose two hybrid data pipelines involving both deterministic and deep learning methods, that are capable to detect efficiently prohibited items inside 3D images of CT scanners, while meeting all of the above requirements.

For both pipelines, during the first stage the 3D images of the bags are segmented, for isolating the individual items, using deterministic image processing methods. During the second stage the segmented 3D objects are pre-processed and passed to a deep learning model that estimates the probabilities of them belonging to a class of threats (e.g. firearms, knives, etc.).

For addressing the limitations that occur due to the lack of big enough training datasets, we propose two different methods. The first one extracts three 2D orthogonal projections of the isolated 3D object, which are preprocessed by convolutional neural networks pre-trained on large datasets, for identifying their high-level characteristics. Then, the extracted information from these projections, is processed by a small neural network, estimating the probability of the inspected object, being a threat.

The second method incorporates a 3D convolutional neural network, which processes the whole volume of the segmented object. For being able to train a 3D model with only a limited amount of training data, we utilise a method called self-supervised learning, training part of the network using unlabelled data acquired from airports. Then, only a small part of the network is fine-tuned on our labelled dataset for estimating the probability of an object being a threat. To the best of our knowledge, this is the first time this method is applied in the field of aviation security, making it possible exploiting the large volumes of data collected at the airports, for developing AI solutions.

It is important to acknowledge that this project has only been possible thanks to the support received from the Department for Transport, the Home Office, and the innovation accelerator Connected Places Catapult, having funded this research, and supporting me for understanding the requirements and needs of the aviation security. I would also like to acknowledge the contribution of my industrial partner Smiths Detection, for sharing with me both their high quality labelled dataset they developed, as long as the data they were able to acquire from airports. Besides of the dataset, the inputs, and feedback received from Smiths Detection AI team, based on their extensive experience in this field, were highly valuable, for shaping the direction of this study.