

Journal Pre-proof

CrimeNet: Neural Structured Learning using Vision Transformer for violence detection

Fernando J. Rendón-Segador, Juan A. Álvarez-García, Jose L. Salazar González, Tatiana Tommasi



PII: S0893-6080(23)00060-6
DOI: <https://doi.org/10.1016/j.neunet.2023.01.048>
Reference: NN 5519

To appear in: *Neural Networks*

Received date : 16 June 2022
Revised date : 20 December 2022
Accepted date : 30 January 2023

Please cite this article as: F.J. Rendón-Segador, J.A. Álvarez-García, J.L.S. González et al., CrimeNet: Neural Structured Learning using Vision Transformer for violence detection. *Neural Networks* (2023), doi: <https://doi.org/10.1016/j.neunet.2023.01.048>.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2023 Published by Elsevier Ltd.

CrimeNet: Neural Structured Learning using Vision Transformer for Violence Detection

Fernando J. Rendón-Segador^{a,*}, Juan A. Álvarez-García^a, Jose L. Salazar González^a,
Tatiana Tommasi^b

^a*Dpto. de Lenguajes y Sistemas Informáticos, Universidad de Sevilla, Spain*

^b*Politecnico di Torino & Italian Institute of Technology, Italy*

Abstract

The state of the art in violence detection in videos has improved in recent years thanks to deep learning models, but it is still below 90% of average precision in the most complex datasets, which may pose a problem of frequent false alarms in video surveillance environments and may cause security guards to disable the artificial intelligence system.

In this study, we propose a new neural network based on Vision Transformer (ViT) and Neural Structured Learning (NSL) with adversarial training. This network, called CrimeNet, outperforms previous works by a large margin and reduces practically to zero the false positives. Our tests on the four most challenging violence-related datasets (binary and multi-class) show the effectiveness of CrimeNet, improving the state of the art from 9.4 to 22.17 percentage points in ROC AUC depending on the dataset. In addition, we present a generalisation study on our model by training and testing it on different datasets. The obtained results show that CrimeNet improves over competing methods with a gain of between 12.39 and 25.22 percentage points, showing remarkable robustness.

Keywords: Deep Learning, Neural Structured Learning, Vision Transformer, Violence Detection, Adversarial Learning

1. Introduction

Violence detection is a very important functionality in public or private security. In smart cities, more and more surveillance cameras are being installed that enable several use cases such as traffic management, infraction or weapon detection [1]. In addition, in schools, hospitals, shopping centres, and other buildings they are also generally used as a dissuasion and in the worst case to identify criminals once a crime has been committed. This growing number of cameras requires sufficient human resources to control the volume of video they

*Corresponding author.

Email addresses: frendon@us.es (Fernando J. Rendón-Segador), jaalvarez@us.es (Juan A. Álvarez-García), jsalazar@us.es (Jose L. Salazar González), tatiana.tommasi@polito.it (Tatiana Tommasi)

generate, however the number of cameras needing attention is greater than the number of staff. Furthermore, after 20 minutes of monitoring a CCTV system, operators' attention spans are considerably reduced [2, 3]. The current difficulties in tackling this problem become even more evident when considering that security is an issue of international concern which scales up from single cities to the World Wide Web. The amount of violent audiovisual content circulating on the network is excessive. To such an extent that the operators in charge of controlling and filtering videos of this type on social networks end up with mental health problems due to over-exposure to violent content¹. Despite this, video surveillance systems are still being manned by humans because the number of false positives is not acceptable in production environments and a human in the loop is still necessary. When analysing the literature, we found that for the most challenging datasets the state of the art does not reach 90% accuracy [4]. To overcome this issue, we aim at designing a robust and sufficiently accurate model to minimise the number of false positives, making it possible to use a video violence detector in real environments.

This work covers the video detection of all types of violent events, both visually intentional actions such as a fight between people, as well as unintentional acts such as an explosion. Moreover, we go beyond standard anomaly detection which differentiates violent from non-violent events: we target a model able to differentiate between different types of violence. This is a complex problem since many types of violent actions can be similar even if their classes are different (there are categories such as abuse, arrest or assault that can be confused with each other) or they can even occur in the same video. To tackle these challenges we propose to combine the powerful self-attention learning paradigm with Neural Structured Learning (NSL). The former is obtained by exploiting the most recent Vision Transformer deep architecture (ViT, [5]). The latter leverages the relation among neighboring samples during training. Specifically, it provides a regularization effect by biasing the network to learn similar hidden representations for close instances.

Our key contributions can be summarized as follows:

- We introduce CrimeNet: a deep model that combines adversarial NSL with ViT for violent activity recognition in videos. Up to our knowledge this is the first time that NSL is used with Transformers rather than with standard convolutional neural networks, and also the first time that it is applied for video recognition.
- CrimeNet improves the state of the art for violence detection on four datasets by reaching an accuracy over 99.98%, with an advantage ranging from 9.4 to 22.17 percentage points in ROC AUC over its competitors. A detailed ablation shows that NSL provides an improvement of 9.55% points in ROC AUC over the use of only ViT as the architecture of the model.
- We present an extensive cross-dataset analysis and show the generalization abilities of CrimeNet. Despite the challenging setting of training and testing on different datasets,

¹<https://www.bloomberg.com/news/articles/2021-12-24/tiktok-sued-by-content-moderator-traumatized-by-graphic-videos>

CrimeNet advances the state of the art from 12.39 to 25.22 percentage points with respect to existing methods that are affected by the domain shift.

This paper is organised as follows: Section 2 provides a brief survey of the state of the art of the problem. Section 3 describes in detail each of the datasets used in this work. Section 4 provides in-depth details on the proposed model architecture. Section 5 summarises the types of experiments proposed in this work. Then, in sections 6 and 7 the results obtained from the experimentation are shown and described. Finally Section 8 summarises the relevance of the results obtained and the possible lines of progress for future development.

2. Related Works

In this section, a review of the state of the art on video detection of violence is carried out. Furthermore, several Neural Graph Learning proposals are shown, where NSL is a specific case, applied to computer vision problems.

2.1. Deep Learning for Violence Behaviour Detection in Videos

Crowd behaviours analysis [6, 7, 8] and detection of violent actions in videos are well-known and well-studied research fields [9, 10], however, since 2014 when the first paper using a deep learning approach [11] appeared, violence detection has advanced by leaps and bounds. Specifically, for some datasets [9, 12], based on short clips and labels for the whole clip, there exist approaches reaching up to 100% accuracy [13].

Given that these datasets were not sufficiently challenging, researchers created new testbeds with hours of CCTV camera recordings: the videos are longer and often annotated at frame-level. In addition, some of them have moved from 2 classes (violence and non-violence) to multiple classes distinguishing the typology of violence in the video.

The first paper to introduce this new family of datasets was [14] in 2018: the UCF Crime collection incorporates the greatest variety of violence classes (14 types). This work also presented a model called DMIL Ranking, in which anomaly detection is approached as a regression problem considering video segments as instances in “multiple instances learning” (MIL) and using a ranking-based loss function to evaluate a fully connected neural network.

The NTU CCTV Fights dataset [15] was presented in 2019: it contains only long violent videos, labelled at frame-level with violent or non-violent classes. This dataset has been evaluated using different known feature extractors such as Two-stream Convolutional Neural Network (CNN) or 3D CNN and different classifiers such as End-to-End CNN, Long-short Time Memory (LSTM) and Support Vector Machine (SVM).

Another method combining the spatio-temporal feature extractor of ResNet 3D [16] and the loss function of Sultani et al. [14] was also published that year, improving the results for the UCF Crime dataset. Zhong et al. [17] evaluated the same dataset using a convolutional graph to clean and refine the classifier based on Temporal Segment Networks [18].

New work focused on anomaly detection emerged in 2020. Degardin [19] provided a new dataset, UBI Fights, and proposed an architecture based on the Gaussian mixture model

(GMM) for the detection of abnormal events in videos applied under the weakly supervised learning paradigm.

Also noteworthy is the work of Kamoona et al. [20] in the weakly supervised setting, using an encoder-decoder architecture (DMIL AutoEncoder). Furthermore in the same year, Wu et al. [21] released the new dataset XD Violence, with 7 classes, the largest number of videos (4754), and hours (217), also incorporating sound. The authors presented a method that exploits visual and audio feature extractors whose output is mixed and provided to graph neural networks to detect short- and long-range temporal relationships.

In 2021 Tian et al. [22] proposed an approach for anomaly detection in videos where a learning function recognises positive instances on the basis of the feature magnitude learning function and using self-attention mechanisms [23]. In the same year, Lv et al. presented a study in which they propose a weakly supervised anomaly localization (WSLA) method that measures variations in both spatial and temporal contexts. Their results mark the current state of the art for the UCF Crime dataset.

In the work by Chang et al. [24], frame anomalies are detected starting from a single binary annotation at the video level. Based on the extracted visual features, attention mechanisms are used to refine the classification of anomalous instances.

Dubey et al. [25] proposed a model that addresses context-dependency by analysing motion and appearance features. They use a 3D ResNet network to extract spatio-temporal and motion feature sets which are then fused and provided as input to a network that learns context-dependencies in a weakly supervised manner using multiple classification measures (MRM).

Another noteworthy study is that of Feng et al. [26] in which they developed a multi-instance self-training framework (MIST) to efficiently refine task-specific discriminative representations with only video-level annotations. The model is composed of a multi-instance pseudo-label generator and a self-guided attention function encoder that aims to automatically focus on anomalous regions in frames while extracting task-specific representations.

Finally, Degardin et al. [27], presented an iterative learning framework composed of two expert systems working in the weakly supervised and self-supervised paradigms. This work combines 3D convolutional neural networks for both paradigms with a Bayesian network that is responsible for performing data augmentation.

To provide some reference background, the next section reviews state of the art of Neural Graph Learning, and general cases for NSL and ViT.

2.2. Applied Neural Graph Learning

Graph-based neural learning has been applied to several tasks related to human action recognition.

In 2019 Shi et al. [28] proposed to represent skeleton data in a directed acyclic graph based on the kinematic dependence between joints and bones. In 2021, Xu et al. [29] presented a convolutional graph neural network to estimate 3-D human pose in which the input data were structured using an hourglass graph. In the same year, a model to classify sports videos was presented in [30]: it consists of a convolutional model designed with

attention mechanisms for assigning weights to neighbouring nodes, combined with a third-order hourglass graph used to structure the features of the videos.

Graph-based learning has been also used in [31] to detect 3D objects in videos. In this case, the data were encoded through a grid message passing network (GMPNet). Considering each grid as a node, the data were structured using a k-NN network and the model was a spatio-temporal Transformer-GRU.

There are other fields of research that have experimented with NSL and adversarial learning with significant results, such as Ren et al. [32] for fake news detection through social networks. It has also been applied to protect against adversarial attacks using perturbed data [33], achieving significantly better performance compared to state of the art defense methods.

2.3. Applied Vision Transformer

ViT [5] profits the Transformer [23] potential, avoiding inductive bias such as translation invariance and locally restricted receptive field in images. To do it, it splits an image in a sequence of patches, flattens them, produces linear embeddings, adds positional embedding to know where is located each patch in the original image, and feeds this sequence as an input to a standard transformer encoder (composed by a multi-head attention layer [23] that allows the model to jointly attend to information from different representation subspaces at different positions) as it can be seen in Figure 4. The success of Transformers, ViT, and their variations [34] is beyond doubt, and have improved the state of the art in many areas such as frame synthesis [35], action recognition [36], or object detection in videos [37].

Our approach differs from previous proposals although being inspired by the use of graph neural networks already leveraged in [21, 17]. As we will describe in the following, we propose a new approach based on supervised NSL [38, 39] and ViT [5]. To our knowledge, the NSL paradigm is used here for the first time for violence recognition in videos.

3. Datasets

For our work, we focus on four video datasets recording violent events. They all contain 1000 videos or more, each one ranging from hundreds to thousands of frames. A summary of the datasets' information is in Table 1, while the following list provides further details:

- NTU CCTV Fights [15] is a dataset containing 1000 videos depicting real-world fights. Of this dataset, 280 videos are recorded from CCTV and 720 from other sources such as mobile cameras or drones, containing different types of fights, ranging from 5 seconds to 12 minutes, with an average duration of 2 minutes.
- UBI Fights [19] is a large-scale dataset of 80 hours of video fully labelled at frame-level. It consists of 1000 videos, where 216 videos contain a fight event and 784 are normal everyday situations. All unnecessary video segments (e.g., video introductions, news, etc.) that could disrupt the learning process were removed. The title of the videos contains indicators related to the type of the respective video. The dataset is divided into binary classes: violence and normal.

- XD Violence [21] is a large-scale dataset with a total duration of 217 hours, containing 4754 untrimmed videos with audio signals and video-level tagging. The dataset is divided into the following seven anomalous categories: abuse, car accident, explosion, fight, shooting, and riot.
- UCF Crime [14] is a dataset consisting of long untrimmed surveillance videos covering 14 real-world violent classes, including abuse, arrest, arson, assault, traffic accident, burglary, explosion, fight, robbery, burglary, shooting, theft, shoplifting, and vandalism.

Dataset	N ^o Items	N ^o Classes	N ^o Hours
NTU CCTV Fights [15]	1000	2	17.68
UBI Fights [19]	1000	2	80
XD Violence [21]	4754	7	217
UCF Crime [14]	1900	14	128

Table 1: Information on datasets used.

4. Model Architecture

This section shows the type of model and architecture used to address the problem. An in-depth definition of the model and the adaptations applied for our use case is provided.

4.1. Pre-processing

4.1.1. Optical flow

Since our model analyses the video frame by frame, including temporal information in each frame is critical. We use optical flow for this purpose [40]: it takes two adjacent frames and represents in an image the amount of pixel variation caused by the observed movements. Of course, the parts of a frame that move together will correspond to pixels with the same intensity as in the example of Figure 1.

The state of the art demonstrates that the use of optical flow as input in violence detection typically improves the use of RGB input [41, 42, 13].

4.1.2. Adversarial neighbours

As it will be seen when describing the model’s details (Section 4.3), NSL is used, a new learning paradigm to train neural networks by using structured signals in a graph. This assumes that the model receives two inputs: the RGB frames, in our case transformed by an optical flow algorithm, and a similarity graph. This graph, or structured signal, is used to represent relationships between samples. The similarity graph regularises the training of a neural network, forcing the model to learn accurate predictions by minimising the supervised loss function while maintaining the input structural similarity by reducing the loss function of the neighbouring node.



Figure 1: Sequence of frames in RGB format and their corresponding optical flow. In this frame sequence, we go from a normal event to a violent event (explosion). Images from the XD-Violence dataset [21].

The similarity graph is generated from the training examples using a graph builder. Each entry is assumed to have an ID and an embedded vector as features. On the one hand, the ID uniquely records and identifies each instance; on the other hand, the embedded vector is assumed to capture the essence of each example by representing it as a list of floating-point values. The graph generator compares the embedded vectors of all input pairs. The degree of similarity between any two samples is calculated as the cosine similarity of their embedded vectors. The brute-force approach to constructing a similarity graph from instance embeddings is $O(n^2)$, which does not scale well to large training sets.

To mitigate that problem, the graph builder uses a well-known randomization technique called locality-sensitive hashing or LSH [43]. To describe this approach, it is assumed that we have a bidimensional embedding vector represented by n points plotted on a cartesian coordinate plane.

The first step of the LSH process is to choose some random hyperplanes through the origin as it can be seen in Figure 2. These hyperplanes divide the space and thus the points into discrete sections that are called LSH buckets. Although there are several buckets, the number of points in each bucket is expected to be much smaller than the complete input set. The quadratic nature of comparing all pairs within each bin has a much smaller impact on performance; comparisons within each bin result in a certain number of graph edges within the bin. This generates a series of connected components each of which corresponds to a bucket, these components are separate and are not part of an overall graph. To complete the similarity graph, the bucketing process is repeated several times, with each round of bucketing choosing to select a different number of random hyperplanes until all connected components are connected.

We highlight that the similarity graph is needed only at training time, since during inference it would not be useful to regenerate this graph by adding only the samples to be inferred as the whole LSH process would be necessary again. The inference workflow will




Figure 2: A representation of the LSH process for the construction of the similarity graph. The dashed red lines corresponding to the labels H1, H2, and H3 are the set of hyperplanes that divide the different points and group them into buckets. α and β correspond to two connected components of the similarity graph. The circles represent instances of the training dataset, in this paper, frames from each training video.

not change, simply providing predictions on new testing video frames.

Obtaining the exact similarity relationship between input data is difficult: it is not trivial to evaluate the likeness of one instance to another without a predefined feature embedding that differentiates the type of violence or separates crimes from normal actions. To overcome this issue we propose to generate synthetic neighbouring instances via Adversarial Learning [44].

In the case of image classification networks, adversarial examples are nothing more than samples to which the pixels have been modified in order to alter the predicted class. Their appearance is similar to that of the original images, but they induce model confusion causing prediction errors. Generally, the adversarial examples are obtained by exploiting the inverse gradient direction of an optimized model. We adopt this logic on the optical flow images.

We start from the hypothesis that normal and violent events have different pixel clustering and variability in the amount of motion measured using optical flow, with the latter being greater in violent events. Likewise, between different types of violent events, there will be variations in the amount of motion and type of pixel clustering. Figure 3 shows the scheme of a structured signal in which the optical flow images in Figure 1 are grouped into two connected components α and β , assuming a binary classifier. The images and small

circles are the new samples produced via adversarial learning.

With this procedure, we obtain adversarial instances, one per image with respect to the total training set. The samples in the same connected components are visually similar but were created on purpose to confuse the original model, thus this data makes it more robust and able to generalise better.




Figure 3: Similarity graph of optical flow frames generated using adversarial perturbation of the original input (Figure 1). Component α are images of non-violent events and component β are images of violent events. This similarity graph is used as the second input to the model. The images inside small circles correspond to other adversarial samples

4.2. Vision Transformer

As basic classification model that we expect to be initially fooled by the adversarial examples, we use a deep transformer network.

Specifically, we consider the Vision Transformer (ViT) model [5] which has achieved remarkable results compared to CNN, in addition to reducing the computational cost required for training and exhibiting a generally weaker inductive bias. ViT is a model based on a Transformer architecture [23] initially used for image classification but later adapted to other visual tasks such as next-frame prediction [45] or video classification [46].

It divides an image into a sequence of fixed-size fragments called patches, correctly embeds each of them, and includes their positional encoding as input to the Transformer encoder. The fragments are used similarly to the series of embedded words for text Transformers, and the output is a prediction label for the image.

Although training a ViT model that has good performance requires from 10 million to more than 100 million images and a huge amount of time and resources [5], the authors of [47] released more than 50000 ViT models trained under diverse settings (including patch size of 8, 16 and 32) on various datasets². The one selected in this work is called ViT-S or DeiT-S [48] and it is trained using a different strategy than the original ViT [5], achieving good metrics (83.73% in ImageNet [49] top-1) with reduced size (115MB). It is pre-trained with ImageNet 21K and fine-tuned with ImageNet 1K and the patches used are of dimension 16×16 pixels. Given that the input frames are resized to 224×224 , that means the ViT model ingests 14^2 patches. It is worth considering that smaller patch sizes are computationally more expensive. For our work we keep this standard decomposition cardinality of the video frames in patches since we are mainly interested in how ViT combines with NSL: we are aware that the patch size influences the performance of the ViT models [5, 47, 48], but this is orthogonal to our analysis and we expect that any improved choice of the patch size would inevitably further improve the observed results.

ViT architecture [5, 50], is shown in Figure 4.




Figure 4: ViT Architecture for frame classification, based on Dosovitskiy et al. [5] and Vaswani et al. [23] works

. We split the image into fixed-size patches and linearly embed each of them. We add positional embeddings by generating a vector sequence with which a standard Transformer encoder is fed. For classification, an extra classification element is added to the vector sequence. Classification is performed by a multilayer perceptron.

4.3. Neural Structured Learning

In order to take advantage of the power of the similarity graph generated by means of adversarial instances, we follow the scheme proposed in NSL [38, 39] (See Figure 5). In it we

²https://github.com/google-research/vision_transformer

can see that we have a pair of inputs, the frames of the training set transformed by optical flow, and a similarity graph or structured signal with instances generated by adversarial modified versions of the sample frame to which it is applied with small perturbations. The generated adversarial neighbours form a similarity graph. In the next step, the original instances are combined with their neighbours and serve as input to the ViT. An encoding of examples and their neighbours is obtained as output. The final regularised graph is the sum of the discriminate loss and the regularisation loss of the graph [38].

Using this connection, neural networks learn to maintain the similarities between the sample and the adversarial neighbours while avoiding the confusion resulting from misclassifications, thus improving the quality and accuracy of the overall neural network.




Figure 5: Neural Structured Learning using as processing model ViT (Figure 4). The neural network minimizes two loss functions, the supervised and the adversarial loss function, which is shown in the equation 9. Figure inspired in Juan et al. [51] and Gopalan et al. [39] works.

Being $T = t_1 \dots t_n$ the training dataset and $Y = y_1 \dots y_m$ the set of labels associated with the instances of the training dataset, through the adversarial learning process a set of neighbour nodes n is generated for each instance of the training dataset, see equation 1.

$$\forall t_i \in T \rightarrow t_i : \{n_1 \dots n_k\} \quad (1)$$

For each subset of neighbours associated with an instance t_i a subgraph $H_{t_i} = (V', E')$ is generated, where each of the neighbours will be connected to the instance, see equations 2 and 3.

$$V' = \{t_i, n_1 \dots n_k\} \quad (2)$$

$$E' = \{(t_i, n_1) \dots (t_i, n_k)\} \quad (3)$$

The sum of all subgraphs forms the graph $G(V, E)$ which we call the structured signal (equation 4).

$$G(V, E) = \sum_{t_i \in T} H_{t_i} \quad (4)$$

From the structured signal $S = G(V, E)$ one can distinguish vertices containing the instances (vertices) of the pure training set $V_t = \{t_1 \dots t_n\}$ and those of the neighbouring vertices generated by adversarial learning $V_n = \{n_1 \dots n_l\}$ both of which form the input to the ViT model that finish in a multi-layer perceptron (MLP) generating a feature vector. The model returns an embedding $\Phi(V_t)$ from V_t and an embedding $\Phi(V_n)$ from V_n , the latter only during training. The feature vector $\Phi(F_v)$ is the result of the graph regularisation between the vectors $\Phi(V_t)$ and $\Phi(V_n)$, see equation 5.

$$\Phi(F_v) = \Phi(V_t) + \Phi(V_n) \quad (5)$$

It should be noted that the ViT model takes each of the vertices of the structured signal S and fragments the image it contains into a sequence of patches $P(v) = \{pv_1 \dots pv_n\}$ that embeds a vector Γ and then by the process of structured learning is re-embedded under the vector $\Phi(F_v)$, the actual embedding process is represented mathematically by the equations 6, 7 and 8.

$$V_t = \bigcup_{\forall v_t \in V} P(v_t) \quad (6)$$

$$V_n = \bigcup_{\forall v_n \in V} P(v_n) \quad (7)$$

$$\Phi(F_v) = \Phi(\Gamma(V_t)) + \Phi(\Gamma(V_n)) \quad (8)$$

Once the vectors are embedded, the final step consists of regularising the vector $\Phi(F_v)$ and applying the adversarial loss (sparse categorical cross-entropy) function [52]. Being $y_i \in Y$ the actual label value of a vertex i and $g_\theta(n_j)$ the prediction of the neighbour node j , the adversarial loss function is shown in the equation 9.

$$\sum_{n_j \in V(t_i)} \epsilon(y_i, g_\theta(n_j)) \quad (9)$$

5. Experimental setting

This section shows the guidelines followed during the experimentation. What kind of experimentation has been performed, how it has been performed, and the motivations related to such experimentation.

All experiments have been conducted using a computer with an Intel(R) Core(TM) i7-9700F CPU 3.00GHz, 16 GB of RAM and an NVIDIA RTX 2080 Super GPU. Details of experimentation are shown through GitHub³.

5.1. Datasets partitions

In the first batch of experiments, the model is trained and tested with each of the datasets following the predefined partitions for each of them. Each of the datasets opts for a different type of labelling. These types of labelling can be grouped into three. Firstly, frame-level labelling, where each frame has a label associated with it depending on whether it is a frame that harbours violent behaviour or not, this is the case of UBI Fights. Secondly, interval labelling, in which for each instance of the dataset the interval in frames or seconds in which the violent action occurs is provided, this is the case of NTU CCTV Fights (intervals per second) and XD Violence (intervals per frame). And the third and last case is video-level labelling, where the entire instance of the dataset has a single label associated with it that classifies the entire video, this is the case of UCF Crime training set (the testing set is labelled using intervals per frame). We consider the finest labelling to be at frame level so we will use it for our model. For interval-labelled datasets, we take each interval and label each frame belonging to that interval based on interval class. In the case of UCF Crime training set, all frames belonging to an instance will have the same label as the instance.

The datasets provide pre-defined guidelines for their division into training, validation and testing datasets. Rather than creating our own divisions, we strictly follow the guidelines provided by each of the datasets: NTU CCTV Fights uses three randomly selected partitions: 50% training, 25% validation and 25% testing; UBI Fights three fixed subsets (80%, 5% and 15%); XD Violence use 3954 videos for training and 800 for testing; UCF Crime uses 800 normal and 810 anomalous videos for training and 150 normal and 140 anomalous for testing in 4-fold cross-validation. The divisions are made at the instance level (videos) so first, the partition is made into training, validation and test subsets and then the frames are extracted from each instance and labelled.

5.2. Preparation of ablation study data

The second batch of experiments aims to see the contributions of NSL to a deep learning model, in this case, a ViT. For this purpose, an ablation study is performed where NSL is eliminated, including the similarity graph using adversarial learning, and replaced by supervised learning, only standard optical flow frames are used as input of the ViT. In this ablation study, the same type of experimentation is performed as with the first batch (following the predefined partitions), substituting one type of learning (ViT + NSL) for another (only ViT). The previous and this experiments are called In-domain experiments.

5.3. Cross-datasets data preparation

In the last batch of experiments, the objective is to measure the generality of the concept of violent action by means of cross-datasets experiments (do not confuse it with cross-

³<https://github.com/FernandoJRS/CrimeNet-ViT-NSL>

validation). For this purpose, the whole of the source dataset is used as the training set and the whole of the target dataset with which the model is to be evaluated as the test dataset.

On the one hand, the single-class datasets, NTU CCTV Fights, and UBI Fights are crossed with each other by training the model with one set and validating with the other. On the other hand in the multi-class datasets, XD Violence and UCF Crime, cross-datasets experiment with all the classes cannot be performed since they do not have the same amount of classes and not all of them are coincident, for that reason, the model must be trained only with those classes coincident between both datasets. The matched classes between XD Violence and UCF Crime are shown in the table 2. Given the similarity of the instances of the classes ‘Road Accident’ in UCF Crime and ‘Car Accident’ in XD Violence, we consider them matched classes.

UCF Crime Classes	XD Violence Classes	Match Classes
Normal	Normal	✓
Abuse	Abuse	✓
Arrest	-	✗
Arson	-	✗
Assault	-	✗
Burglary	-	✗
Explosion	Explosion	✓
Fighting	Fighting	✓
-	Riot	✗
Road Accident	Car Accident	✓
Robbery	-	✗
Shooting	Shooting	✓
Shoplifting	-	✗
Stealing	-	✗
Vandalism	-	✗

Table 2: Matching classes between the UCF Crime and XD Violence datasets.

Two sequences of frames from the classes Car Accident from the XD Violence dataset and Road Accident from the UCF Crime dataset are shown in Figure 6. The frames from the UCF Crime dataset are of lower quality than those from XD Violence, in this case from a movie. However, both classes capture the same concept of traffic accidents and can be compared and considered the same class despite the difference in quality and camera focus.

5.4. Metrics

The metrics used in the experiments are as follows:

- Confusion Matrix (CM): Allows the display of the number of predictions made by a model that matches the labels.
- Receiver Operating Characteristic Area Under the Curve (ROC AUC): Calculates sensitivity versus specificity for a classifier system as the discrimination threshold is varied.



Figure 6: Comparison between two frame sequences of the classes Car Accident XD Violence (Top) and Road Accident UCF Crime (Bottom).

- Average Precision (AP): Summarises the area under the precision-recall curve (PR AUC) as the weighted average of the accuracy achieved at each threshold.

6. In-domain: Results from each dataset




Figure 7: (a) XD Violence Confusion Matrix, (b) UCF Crime Confusion Matrix.

The results in Table 3 show the effectiveness of CrimeNet with respect to its competitors. More precisely, Figure 7 presents the nearly perfect confusion matrices of CrimeNet for the multi-class XD Violence and UCF Crime datasets. We highlight that the UCF Crime dataset comes with four different train/test dataset splits: for all of them the CrimeNet results are stable with almost zero standard deviation, confirming the robustness of the model. The inference time for these results is about 40 milliseconds.

Dataset	Method	ROC AUC	AP
NTU CCTV Fights [15]	3D CNN [15]	-	79.50%
	ViT	90.45%	90.40%
	CrimeNet	100%	100%
UBI Fights [19]	BayesianNet[27]	84.60%	-
	ViT	89.76%	89.73%
	CrimeNet	100%	100%
XD Violence [21]	HL-Net [21]	-	78.64%
	Contrastive Attention [24]	-	76.90%
	RTFM [22]	77.81%	-
	CrimeNet	99.98%	99.95%
UCF Crime [14]	DMIL Ranking [14]	75.41%	-
	GMM [19]	75.90%	-
	3D ResNet [16]	76.67%	-
	DMIL AutoEncoder [20]	80.10%	-
	DMIL-MRM [25]	81.91%	-
	GCN [17]	82.12%	-
	MIST [26]	82.30%	-
	RTFM [22]	84.30%	-
	Contrastive Attention [24]	84.62%	-
	WSAL [4]	85.38%	-
	CrimeNet	99.98%	99.97%

Table 3: Comparison of CrimeNet results with the state of the art for the case study datasets. Our ablation study using only ViT is also included.

By reporting the results of ViT we provide an ablation on the role of NSL: indeed CrimeNet builds over ViT and further exploits the sample neighbour graph via NSL. The results indicate that CrimeNet has an advantage over ViT of around 10% points confirming that NSL generates a more robust model by establishing stronger relationships between similar image features. These surprisingly good results are in agreement with those obtained in other works using NSL such as [53].

We note that ViT is already surpassing all the state of the art models except in the case of UBI Fights where results are very close in ROC AUC to Sultani et al. [14] (89.76% vs 90.60%). The approximate inference time of ViT remains at 40 milliseconds, showing the inclusion of NSL in the model does not cause a prediction delay.

Finally, CrimeNet outperforms the current state of the art by far. For the less studied datasets such as NTU CCTV Fights and UBI Fights, perfect results are obtained, which is 20.5% points higher in AP for NTU CCTV Fights and 9.4% points higher ROC AUC in UBI Fights. For the case of the multi-class datasets, the results obtained are again significantly higher than the state of the art. For the XD Violence dataset, the improvement is 22.17% in ROC AUC while for the UCF Crime dataset is 14.6%.

7. Cross-domain: Results across datasets

Given the results of our model CrimeNet using ViT and NSL so close to 100% of the metrics, this experiment is essential to check if overfitting is occurring. The results of this experiment show a good performance of the model in generalising the concept of violent action. Figures 8 and 9 show the confusion matrices for the cross-datasets experiment between NTU CCTV Fights - UBI Fights and XD Violence-UCF Crime datasets (retrained using the common classes shown in Table 2).




Figure 8: (a) Crossvalidation NTU CCTV Fights - UBI Fights, (b) Crossvalidation UBI Fight - NTU CCTV Fight.

Table 4 shows the accuracy results for each of the cross-datasets experiments. These results show that the single-class datasets generalise the concept of violence better, achieving results of around 80% accuracy. The multi-class datasets show significantly lower performance, but higher than 70% in both cases. It should be noted that none of the state of the art works studied used this experiment to check the robustness of their models, so we carried out it with the state of the art models that have obtained the best results for each dataset in Table 3, that is, GMM [19] for UBI Fights, RTFM [22] for XD Violence and WSAL [4] for UCF Crime and whose code is available to reproduce. ViT (without NSL) is also included to check the performance of CrimeNet over the state of the art and ViT. As it can be seen, CrimeNet far outperforms the cross-datasets results of the other models, improving in 25.22% points ROC AUC for UBI Fights, 18.73% for XD Violence and 12.39% for UCF Crime respectively. For the NTU CCTV Fights dataset, there is no reproducible code is available with which to obtain results for comparison.

When XD Violence is used as the training dataset, better generalisation can be observed than when UCF Crime is used. We consider that the size of the first dataset, with more




Figure 9: (a) Crossvalidation UCF Crime - XD Violence, (b) Crossvalidation XD Violence - UCF Crime.

than twice as many videos as the second one, allows better training and ROC AUC than the second one (73.5% vs. 70.20%).

The ablation study on the role of NSL again shows that CrimeNet has a significant advantage over ViT, around 10% (+9.47%, +11.06%, +10.89%, and +9.47%) in the cross-validation experiments. This difference is in agreement with the results obtained in the In-domain experiments and shows that ViT applied to video overpass state of the art but NSL includes much more robustness.

7.1. Recommendations

The practically perfect in-domain results motivated us to carry out the cross-dataset experiment, showing that CrimeNet maintains its advantage over the competitors even in this challenging setting.

Interestingly we noticed that our result in the cross-dataset experiment training in UBI Fights and testing in NTU CCTV Fights, achieves a better performance than the current state of the art (81.35% vs 79.50% [15]).

We also consider that the cross-dataset experiment is a good practice that must be used in these benchmarks, which normally is not present in the related works.

On the basis of the results we showed, it is clear that this challenging setting is still in need of further research.

Finally, we showed how NSL approach using Adversarial learning for generating the similarity graph is a very good approach to obtaining better results and making the model more robust.

Dataset Train	Dataset Test	Method	ROC AUC	AP
NTU CCTV Fights* [15]	UBI Fights [19]	ViT	69.43%	69.50%
		CrimeNet	78.90%	78.87%
		GMM [19]	56.13%	-
UBI Fights [19]	NTU CCTV Fights [15]	ViT	70.29%	70.26%
		CrimeNet	81.35%	81.35%
		RTFM [22]	54.77%	-
XD Violence [21]	UCF Crime [14]	ViT	62.59%	62.55%
		CrimeNet	73.50%	73.44%
		WSAL [4]	57.81%	-
UCF Crime [14]	XD Violence [21]	ViT	60.73%	60.73%
		CrimeNet	70.20%	70.10%

Table 4: Cross-datasets experiments between the four datasets. Binary: NTU CCTV Fights and UBI Fights; Multiclass: UCF Crime and XD Violence. Our ablation study using only ViT is also included. *Code not available to reproduce.

8. Conclusions and future work

In this paper, we present a case study on the detection of violent events in videos and the generalisation of this concept. To address the case study we proposed an innovative deep learning model that combines NSL with ViT architectures called CrimeNet.

This model has far surpassed the current state of the art in violence detection after analysing four datasets and outperforming each of the best performing works by 9.4 to 22.17% points in ROC AUC.

In view of the results, we can state that the combination of ViT with NSL and Adversarial learning for the similarity graph input generates a unique model with high efficiency and robustness in the task of video detection of violence. After the ablation study, we checked the importance of using NSL instead of supervised learning directly. It is shown that applying NSL improves the results by about 10% points compared to supervised learning (from 9.55 to 12.75% points in ROC AUC).

The model also shows a state of the art performance when carrying out cross-dataset experiments to check the generalisation of the concept of violent action. The results reach values above 70% ROC AUC for multiclass (73.5% and 70.2%) and close to 80% for binary-class datasets (81.35% and 78.9%), meaning an improvement with respect to previous models from 12.39 to 25.22 points.

These results again manifest the robustness of the model and the possibility to use this approach in a number of lines of future work such as video classification, detection, classification of human actions, tracking of elements in a video, etc. Another study that we want to carry out is the comparison of different backbones, substituting for instance ViT by 3D ResNet, and the generation of different types of data augmentation such as mixup [54].

9. Acknowledgements

This research is partially supported by the DISARM project - Grant n. PDC2021-121197, and the HORUS project - Grant n. PID2021-126359OB-I00 funded by MCIN/AEI/10.13039/501100011033 and by the “European Union NextGenerationEU/PRTR”.

References

- [1] J. L. Salazar González, C. Zaccaro, J. A. Álvarez-García, L. M. Soria Morillo, F. Sancho Caparrini, Real-time gun detection in CCTV: an open problem, *Neural Networks* 132 (2020) 297–308.
- [2] S. A. Velastin, B. A. Boghossian, M. A. Vicencio-Silva, A motion-based image processing system for detecting potentially dangerous situations in underground railway stations, *Transportation Research Part C: Emerging Technologies* 14 (2) (2006) 96–113.
- [3] T. Ainsworth, Buyer beware, *Security Oz* 19 (2002) 18–26.
- [4] H. Lv, C. Zhou, Z. Cui, C. Xu, Y. Li, J. Yang, Localizing anomalies from weakly-labeled videos, *IEEE Transactions on Image Processing* 30 (2021) 4505–4515.
- [5] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al., An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale, in: *International Conference on Learning Representations*, 2021.
- [6] A. Sadeghian, A. Alahi, S. Savarese, Tracking the untrackable: Learning to track multiple cues with long-term dependencies, in: *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 300–311.
- [7] X. Li, M. Chen, F. Nie, Q. Wang, A multiview-based parameter free framework for group detection, in: *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [8] X. Li, M. Chen, F. Nie, Q. Wang, Locality adaptive discriminant analysis., in: *IJCAI*, Vol. 2201, 2017.
- [9] E. Bermejo Nievas, O. Deniz Suarez, G. Bueno García, R. Sukthankar, Violence detection in video using computer vision techniques, in: *International conference on Computer analysis of images and patterns*, Springer, 2011, pp. 332–339.
- [10] O. Deniz, I. Serrano, G. Bueno, T.-K. Kim, Fast violence detection in video, in: *2014 International Conference on Computer Vision Theory and Applications (VISAPP)*, Vol. 2, IEEE, 2014, pp. 478–485.
- [11] C. Ding, S. Fan, M. Zhu, W. Feng, B. Jia, Violence detection in video by using 3D convolutional neural networks, in: *International Symposium on Visual Computing*, Springer, 2014, pp. 551–558.
- [12] T. Hassner, Y. Itcher, O. Kliper-Gross, Violent Flows: Real-time detection of violent crowd behavior, in: *2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, IEEE, 2012, pp. 1–6.
- [13] F. J. Rendón-Segador, J. A. Álvarez-García, F. Enríquez, O. Deniz, Violencenet: Dense multi-head self-attention with bidirectional convolutional LSTM for detecting violence, *Electronics* 10 (13) (2021) 1601.
- [14] W. Sultani, C. Chen, M. Shah, Real-World Anomaly Detection in Surveillance Videos, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [15] M. Perez, A. C. Kot, A. Rocha, Detection of real-world fights in surveillance videos, in: *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2019, pp. 2662–2666.
- [16] S. Dubey, A. Boragule, M. Jeon, 3D ResNet with Ranking Loss Function for Abnormal Activity Detection in Videos, in: *2019 International Conference on Control, Automation and Information Sciences (ICCAIS)*, IEEE, 2019, pp. 1–6.
- [17] J.-X. Zhong, N. Li, W. Kong, S. Liu, T. H. Li, G. Li, Graph convolutional label noise cleaner: Train a plug-and-play action classifier for anomaly detection, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1237–1246.

- [18] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, L. Van Gool, Temporal segment networks for action recognition in videos, *IEEE transactions on pattern analysis and machine intelligence* 41 (11) (2018) 2740–2755.
- [19] B. M. Degardin, Weakly and Partially Supervised Learning Frameworks for Anomaly Detection, Ph.D. thesis, Universidade da Beira Interior (Portugal) (2020).
- [20] A. M. Kamoon, A. K. Gosta, A. Bab-Hadiashar, R. Hoseinnezhad, Multiple instance-based video anomaly detection using deep temporal encoding-decoding, arXiv preprint arXiv:2007.01548.
- [21] P. Wu, J. Liu, Y. Shi, Y. Sun, F. Shao, Z. Wu, Z. Yang, Not only look, but also listen: Learning multimodal violence detection under weak supervision, in: *European Conference on Computer Vision*, Springer, 2020, pp. 322–339.
- [22] Y. Tian, G. Pang, Y. Chen, R. Singh, J. W. Verjans, G. Carneiro, Weakly-supervised video anomaly detection with robust temporal feature magnitude learning, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 4975–4986.
- [23] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, *Advances in Neural Information Processing Systems* 30.
- [24] S. Chang, Y. Li, J. S. Shen, J. Feng, Z. Zhou, Contrastive Attention for Video Anomaly Detection, *IEEE Transactions on Multimedia*.
- [25] S. Dubey, A. Boragule, J. Gwak, M. Jeon, Anomalous Event Recognition in Videos Based on Joint Learning of Motion and Appearance with Multiple Ranking Measures, *Applied Sciences* 11 (3) (2021) 1344.
- [26] J.-C. Feng, F.-T. Hong, W.-S. Zheng, Mist: Multiple instance self-training framework for video anomaly detection, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 14009–14018.
- [27] B. Degardin, H. Proença, Iterative weak/self-supervised classification framework for abnormal events detection, *Pattern Recognition Letters* 145 (2021) 50–57.
- [28] L. Shi, Y. Zhang, J. Cheng, H. Lu, Skeleton-based action recognition with directed graph neural networks, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 7912–7921.
- [29] T. Xu, W. Takano, Graph stacked hourglass networks for 3D human pose estimation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 16105–16114.
- [30] M. Gao, W. Cai, R. Liu, AGTH-Net: attention-based graph convolution-guided third-order hourglass network for sports video classification, *Journal of Healthcare Engineering* 2021.
- [31] J. Yin, J. Shen, X. Gao, D. Crandall, R. Yang, Graph Neural Network and Spatiotemporal Transformer Attention for 3D Video Object Detection from Point Clouds, *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- [32] Y. Ren, B. Wang, J. Zhang, Y. Chang, Adversarial active learning based heterogeneous graph neural network for fake news detection, in: *2020 IEEE International Conference on Data Mining (ICDM)*, IEEE, 2020, pp. 452–461.
- [33] W. Jin, Y. Ma, X. Liu, X. Tang, S. Wang, J. Tang, Graph structure learning for robust graph neural networks, in: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 66–74.
- [34] S. Khan, M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, M. Shah, Transformers in vision: A survey, *ACM Comput. Surv.* 54 (10s). doi:10.1145/3505244. URL <https://doi.org/10.1145/3505244>
- [35] Z. Liu, S. Luo, W. Li, J. Lu, Y. Wu, S. Sun, C. Li, L. Yang, Convtransformer: A convolutional transformer network for video frame synthesis, arXiv preprint arXiv:2011.10185.
- [36] R. Girdhar, J. Carreira, C. Doersch, A. Zisserman, Video action transformer network, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 244–253.
- [37] Y. Chen, Y. Cao, H. Hu, L. Wang, Memory enhanced global-local aggregation for video object detection, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 10337–10346.

- [38] T. D. Bui, S. Ravi, V. Ramavajjala, Neural graph learning: Training neural networks using graphs, in: Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, 2018, pp. 64–71.
- [39] A. Gopalan, D.-C. Juan, C. I. Magalhaes, C.-S. Ferng, A. Heydon, C.-T. Lu, P. Pham, G. Yu, Y. Fan, Y. Wang, Neural Structured Learning: training neural networks with structured signals, in: Proceedings of the 14th ACM International Conference on Web Search and Data Mining, 2021, pp. 1150–1153.
- [40] G. Farneback, Two-frame motion estimation based on polynomial expansion, in: Scandinavian Conference on Image analysis, Springer, 2003, pp. 363–370.
- [41] P. Zhou, Q. Ding, H. Luo, X. Hou, Violence detection in surveillance video using low-level features, PLoS one 13 (10) (2018) e0203668.
- [42] J. Mahmoodi, A. Salajeghe, A classification method based on optical flow for violence detection, Expert systems with applications 127 (2019) 121–127.
- [43] M. S. Charikar, Similarity estimation techniques from rounding algorithms, in: Proceedings of the thirty-fourth annual ACM symposium on Theory of computing, 2002, pp. 380–388.
- [44] B. H. Zhang, B. Lemoine, M. Mitchell, Mitigating unwanted biases with adversarial learning, in: Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society, 2018, pp. 335–340.
- [45] M. Jahanbakht, W. Xiang, M. R. Azghadi, Sediment Prediction in the Great Barrier Reef using Vision Transformer with finite element analysis, Neural Networks 152 (2022) 311–321.
- [46] A. Arnab, M. Deghani, G. Heigold, C. Sun, M. Lučić, C. Schmid, Vivit: A video vision transformer, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 6836–6846.
- [47] A. P. Steiner, A. Kolesnikov, X. Zhai, R. Wightman, J. Uszkoreit, L. Beyer, How to train your ViT? Data, Augmentation, and Regularization in Vision Transformers, Transactions on Machine Learning Research.
URL <https://openreview.net/forum?id=4nPswr1KcP>
- [48] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, H. Jégou, Training data-efficient image transformers & distillation through attention, in: International Conference on Machine Learning, PMLR, 2021, pp. 10347–10357.
- [49] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, L. Fei-Fei, ImageNet Large Scale Visual Recognition Challenge, International Journal of Computer Vision (IJCV) 115 (3) (2015) 211–252. doi:10.1007/s11263-015-0816-y.
- [50] S. Paul, P.-Y. Chen, Vision Transformers are Robust Learners, in: AAAI Conference on Artificial Intelligence, 2022.
- [51] D.-C. Juan, C.-T. Lu, Z. Li, F. Peng, A. Timofeev, Y.-T. Chen, Y. Gao, T. Duerig, A. Tomkins, S. Ravi, Ultra fine-grained image semantic embedding, in: Proceedings of the 13th International Conference on Web Search and Data Mining, 2020, pp. 277–285.
- [52] I. Goodfellow, J. Shlens, C. Szegedy, Explaining and Harnessing Adversarial Examples, in: International Conference on Learning Representations, 2015.
URL <http://arxiv.org/abs/1412.6572>
- [53] M. Z. Uddin, A. Soylu, Human activity recognition using wearable sensors, discriminant analysis, and long short-term memory-based neural structured learning, Scientific Reports 11 (1) (2021) 1–15.
- [54] H. Zhang, M. Cisse, Y. N. Dauphin, D. Lopez-Paz, mixup: Beyond Empirical Risk Minimization, in: International Conference on Learning Representations, 2018.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Journal Pre-proof