

Great Ape Detection in Challenging Jungle Camera Trap Footage via Attention-Based Spatial and Temporal Feature Blending

Xinyu Yang

Dept of Computer Science
University of Bristol, UK

xinyu.yang@bristol.ac.uk

Majid Mirmehdi

Dept of Computer Science
University of Bristol, UK

m.mirmehdi@bristol.ac.uk

Tilo Burghardt

Dept of Computer Science
University of Bristol, UK

tilo@cs.bris.ac.uk

Abstract

We propose the first multi-frame video object detection framework trained to detect great apes. It is applicable to challenging camera trap footage in complex jungle environments and extends a traditional feature pyramid architecture by adding self-attention driven feature blending in both the spatial as well as the temporal domain. We demonstrate that this extension can detect distinctive species appearance and motion signatures despite significant partial occlusion. We evaluate the framework using 500 camera trap videos of great apes from the Pan African Programme containing 180K frames, which we manually annotated with accurate per-frame animal bounding boxes. These clips contain significant partial occlusions, challenging lighting, dynamic backgrounds, and natural camouflage effects. We show that our approach performs highly robustly and significantly outperforms frame-based detectors. We also perform detailed ablation studies and a validation on the full ILSVRC 2015 VID data corpus to demonstrate wider applicability at adequate performance levels. We conclude that the framework is ready to assist human camera trap inspection efforts. We publish key parts of the code as well as network weights and ground truth annotations with this paper.

1. Introduction

The problem of visually identifying the presence and locations of animal species filmed in natural habitats [16] is of central importance for automating the interpretation of large-scale camera trap imagery. This is particularly challenging in scenarios where lighting is difficult, backgrounds are non-static, and major occlusions, image noise, as well as animal camouflage effects occur: filming great apes via camera traps in jungle environments constitutes one such setting. There, animals appear uniformly dark and blend into the forest during eating, playing or moving in groups sometimes behind trees or thicket. An animal's visual presence and location is thus often only determinable by linking selective spatial and temporal information about species-typical appearance and locomotion across video snippets.

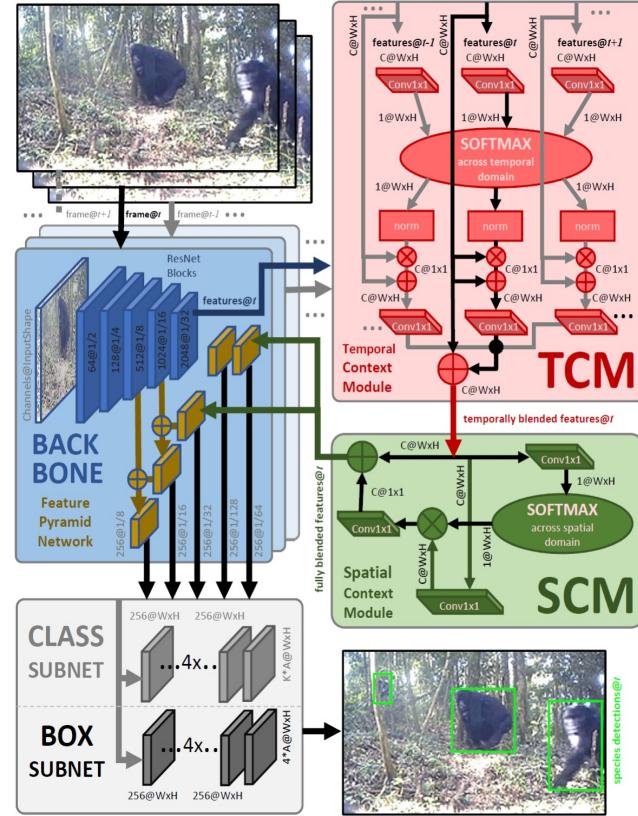


Figure 1: Architectural Overview. Our framework processes video snippets (top left) and robustly detects per-frame great ape locations (bottom right). The architecture extends a traditional detection backbone (blue) by two self-attention components, which select and ‘blend’ most relevant temporal (TCM, red) and spatial (SCM, green) information via SoftMax folding. Blended feature maps are fed back into the feature pyramid network (yellow) before fuelling the detection head, which contains bounding box regression (black) and classification streams (grey).

In this paper, we address this specific video object detection challenge by proposing a blended detection framework integrated into a feature pyramid network (FPN) architecture [11, 18] as illustrated in Figure 1. We introduce two extra components to a traditional detection backbone: a Spa-

tial Context Module (SCM) and a Temporal Context Module (TCM). Each of these modules is driven by a self-attention mechanism tasked to learn how to emphasise most relevant elements of a feature given its context. In particular, these attention components are effective in learning how to ‘blend’ spatially and temporally distributed visual cues in order to reconstruct object locations under dispersed partial information; be that due to occlusion or lighting.

Whilst the self-attention concept has been used recently in various application contexts [28, 4], we tailor it here to spatio-temporal video object detection and propose a flexible component setup that can be utilized as an add-on to different backbone networks. We show that the approach is beneficial in scenarios where distinctive species appearance and motion signatures are only partly accessible and intermittently dispersed across the spatial and temporal domain. Figure 2 exemplifies such scenarios on the Pan Africa camera trap data used in this work. This dataset contains 500 × 15s video clips with the resolution of 720 × 404 and was collected by the Pan African Programme ‘The Cultured Chimpanzee’ (see Acknowledgements). It was subsequently labelled for this paper with accurate bounding box ground truth for all animals in each frame.

Contributions. (1) FPN-integrated blended detection framework driven by spatial and temporal self-attention components; (2) Cross-dataset ablation study of the framework quantifying details for various backbones and setups; (3) Comprehensive animal bounding box annotations in 180K frames of the Pan Africa dataset.



Figure 2: Pan Africa Great Ape Camera Trap Dataset. Sample frames from four representative videos of the dataset with ground truth animal annotations in red. (*1st row*) clip showing two easily detectable animals walking past a camera trap with a third animal moving behind thicket; (*2nd row*) clip where two great apes hide and emerge from behind bushes; (*3rd row*) clip with frames of extreme close-ups; (*4th row*) clip with challenging lighting.

2. Related Work

Object detection in video is usually implemented via video feature extraction [25, 5, 15], frame feature aggregation [6, 30, 29], or detection post-processing [10, 13, 14]. The task is distinct from object tracking [31, 17, 1] since the object instances to follow are not given – they may appear or disappear at any point and need to be localised and determined as absent or present on a per-frame basis.

Similar to other video tasks, video object detection relies on extracting relevant information from the spatio-temporal corpus formed by frame sequences. For instance, C3D [26] and I3D [5] explore 3D convolutional neural networks (CNNs) to generate mixed motion-RGB features from frame sequences. In contrast, dual-stream ConvNets such as Siamese ConvNet [25] and TSN [27] apply separate networks for RGB and motion processing where optical flow is often used as a pre-computed input. Frame feature aggregation, as used in DFF [30], FGFA [29], and ST-Lattice [6], deals with motion implicitly by training networks that fuse frame-level spatial features to explore the temporal structure of videos. Detection post-processing, such as used in T-CNN [13] and Seq-NMS [10], ignores explicit temporal feature construction altogether. Operating closer to a traditional tracking paradigm, these methods instead link detection results from individual frames into tracklets by optimising overall trajectory scores.

Before focusing on the most relevant details of some of the above methods, we will first review the foundations on which many of the approaches rest, that is: how to extract detection-relevant features from single frames.

2.1. Single Image Object Detection

Object detection fundamentally requires two conceptual tasks to be solved: localisation and classification of content of interest. Region proposal based methods such as F-RCNN [24] or Cascade-RCNN [3] pre-process images first to generate class-agnostic regions of interest (ROIs) before classifying these and regressing associated bounding boxes. F-RCNN, in particular, uses a region proposal network (RPN) built on top of a backbone shared with the classification component. More recently, to improve the quality of detection, Cascade-RCNN appends two more stages based on F-RCNN and replaces ROI pooling with ROI alignment making the framework more robust.

In contrast to all region proposal based methods, single-shot detectors infer class probability and bounding box offsets within a single feed forward network. This approach is usually simpler in design and faster at runtime as impressively shown, for instance, by YOLO [22, 23] and SSD [20]. Darknet-based YOLO [22] in particular regresses anchor boxes and box scores directly as well as class probabilities. Due to its versatility and fast performance, YOLO has formed the detection backbone for successfully con-

structed primate face detectors [2] for single images in the past. Similarly to the YOLO approach, SSD predicts a fixed number of bounding boxes. However, by applying detection at different scales, SSD has been shown to adjust better to different object sizes and aspect ratios. More recently, by addressing the class imbalance problem of all of the previous single-shot methods, RetinaNet [19] replaces the cross-entropy loss with the focal loss for classification, which focuses the training on hard examples.

2.2. Video Object Detection

Whilst single image object detectors are directly applicable to video in a frame-by-frame manner, they ignore – by definition – temporal cues. Yet, these are often vital for detection under challenging conditions as for the case at hand. We will next briefly review the key ideas behind most relevant recent works such as FGFA [29], D&T [9], T-CNN [13], and Seq-NMS [10]. The fundamental task for all these methods is to integrate information from the temporal domain directly into detection decisions.

Linking single frame detections across the temporal dimension as done by T-CNN [13] constitutes possibly the simplest form of temporal domain exploration. T-CNN essentially runs region-based detectors per frame and enforces motion-based propagation to adjacent frames. This classical tracking paradigm thereby extends detections into tubelets, which after re-scoring and suppression of overlaps yield the final set of detected objects. In contrast, D&T [9] interlinks single image detection and tracking in a unified approach using ROI pooling on both detection-based feature maps and tracking-based correlation feature maps, where a specific correlation layer is introduced to produce the latter.

Seq-NMS [10] follows a similar paradigm constructing sequences along temporally close, high confidence bounding boxes in consecutive frames. Their governing metric for sequence association maximises overall confidence and IOU scores, and sequence-based non-maximum suppression is utilised to fuse or filter out overlapping tracklets. Whilst temporal consistency can be extrapolated this way, the spatial distortion effects across the temporal domain are not accounted for.

To address this, FGFA [29] performs optical flow guided spatial warping before aggregating features. The resulting features are subsequently fused temporally by weighted element-wise addition where weights are determined by the optical flow field. These descriptors contain rich spatio-temporal information that have been shown to address problems such as blurred object detection and partially occluded object detection.

Taking closest inspiration from the temporal cue aggregation mechanism used in FGFA, we propose an attention-based spatial *and* temporal feature blending framework that can be used as an add-on to existing (and future) feature

pyramid networks as they are in standard use for object detection today [19]. Attention-based spatial blending was successfully used before in GCNet [4], but attention-based spatio-temporal blending is – to the best of our knowledge – novel as a core concept for video object detection.

We will now describe how such information blending across space and time can be implemented, and how it is beneficial to addressing the problem of great ape detection.

3. Detection via Spatial & Temporal Blending

Task. Our task is to generate frame-level bounding box detections $\mathbf{D} = \{\mathbf{D}_1, \dots, \mathbf{D}_t, \dots\}$ of great apes across a video $\mathbf{X} = \{\mathbf{X}_1, \dots, \mathbf{X}_t, \dots\}$, where \mathbf{D}_t are the detection results for frame \mathbf{X}_t at time step t . To predict the bounding box detections \mathbf{D}_t , the network can utilise both previous *and* future frames, *i.e.* a snippet $\{\mathbf{X}_{t-\tau}, \dots, \mathbf{X}_t, \dots, \mathbf{X}_{t+\tau}\}$ of length T .

Overall Concept. As shown in Figure 1, our integrated architecture extends a standard feature pyramid network by two extra components: a Spatial Context Module (SCM) and a Temporal Context Module (TCM). Each of these modules is driven by a self-attention mechanism that learns how to emphasise the most relevant elements of a feature given its context. Both components follow a principled workflow similar to the one described by Cao et al [4] and visualised in Figure 3. Essentially, after grouping inputs along the dimensions of attentional interest (spatial or temporal), features are embedded into a lower dimensional space and a self-attention map of the feature is created. This map is then applied back onto the features in order to ‘blend’ it and emphasise elements important to the detection process whilst suppressing other content. Critically, these components are trainable as part of the network

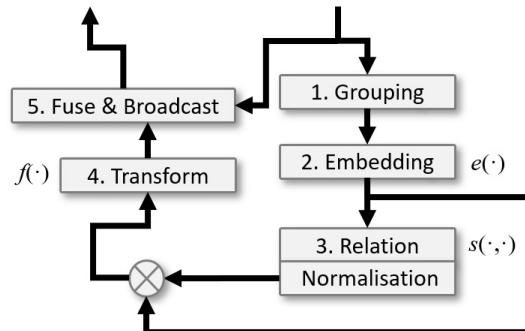


Figure 3: Principled Workflow of the Blending Framework. (1) Input features are grouped by the target dimension; (2) Each location of the feature map is embedded by its position; (3) A correlation map is computed and reduced. This map is then applied back to the embedded feature via matrix multiplication (marked as \otimes) implementing the self-attention principle. (4) The resulting feature is finally transformed, before (5) being fused with the residual to form the blended output. Key functions are shown at the point of application as annotations.

and can be rolled out across space and time so that dispersed species-distinctive information can be selected from within the spatio-temporal volume.

For a given spatial or temporal module and position i in the input feature x_i , the context-enhanced and ‘blended’ output feature z_i can in its simplest form be expressed as:

$$z_i = x_i + f\left(\frac{1}{M} \sum_{j=1}^M s(e_i(x_i), e_j(x_j)) e_i(x_i)\right), \quad (1)$$

where x_i is the i th descriptor of the residual feature map, $e(\cdot)$ is the embedding function, $f(\cdot)$ is the transform function, $s(\cdot)$ is the correlation function, M is the number of positions in the feature, *e.g.* $M = HWT$ for a video sequence (H denotes height, W denotes width, T denotes snippet length), and j enumerates all the positions across the context. The embedding function $e(\cdot)$ and transfer function $f(\cdot)$ are implemented via 1×1 convolution kernels without bias using learnable weights w .

Spatial Module. As depicted in Figure 1 (green), we use a simplified non-local attention component with a SoftMax function applied across the embedded feature vector as the SCM. The module output z_i^{SCM} is:

$$z_i^{SCM} = x_i \oplus w^3 \sum_{j=1}^M \left(\frac{\exp(w^1 x_j)}{\sum_{a=1}^M \exp(w^1 x_a)} \otimes (w^2 x_j) \right), \quad (2)$$

where j enumerates all context locations, x_i represents the incoming features, \oplus denotes element-wise broadcast addition, \otimes denotes multiplication of tensor elements, w^1 and w^2 describe the different learnable parameters of linear embedding functions $e(\cdot)$, and w^3 represent the parameters of a linear transform function $f(\cdot)$.

Temporal Module. Figure 1 (red) visualises the temporal module in detail, which follows the principles laid out in Equation (1) in general terms. In particular, the TCM module is constructed to apply self-attention across a short-term temporal context where, for each input frame X_t , a feature x_t from the previous layer is first embedded by a linear function $e(\cdot)$ with weights w_t^4 . Subsequently, its temporal correlation function $s(\cdot)$ is modelled by a global Soft-Max $C(\cdot)$ across a temporal context of $T = 2\tau + 1$ selected nearby frames. For each feature position $i \in M$ the global SoftMax is defined as:

$$C(x_{t,i}; w^4) = \frac{\exp(w_t^4 x_{t,i})}{\sum_{m \in T} \exp(w_m^4 x_{m,i})}, \quad (3)$$

where frame m enumerates all T frame positions, and w_t is the linear embedding parameter for time step t . We then use the mean value of all positional temporal attention to normalise the term yielding a temporal attention map for time step t as:

$$\hat{x}_{t,i} = \frac{1}{HW} C(x_{t,i}; w^4) \sum_{j=1}^{HW} C(x_{t,j}; w^4). \quad (4)$$

In order to visually illustrate this concept, two examples of temporal attention maps projected back into the image domain are depicted in Figure 4. The maps highlight distinctive, dispersed features for context-aware inference revealing target object locations despite heavy partial occlusion.

These maps are subsequently applied back onto the original feature by matrix multiplication \otimes , element-wise broadcast addition \oplus , and two linear transformations. One transform learns adaptive weights (*i.e.* w_m^5) to ‘blend’ between original and attention-adapted features, the other learns weights (*i.e.* w_m^6) to ‘blend’ across the temporal domain. Finally, an additive connection to the original inputs $x_{t,i}$ is made to ensure learning stability yielding:

$$z_{t,i}^{TCM} = x_{t,i} + \sum_{m \in T} w_m^6 (x_{m,i} \oplus w_m^5 \sum_{j=1}^{HW} \hat{x}_{m,j} \otimes x_{m,j}). \quad (5)$$

3.1. Overall Network Architecture

Backbones. We follow the layout of RetinaNet [19] and Cascade [3] as the two major state-of-the-art architectures in our study. Figure 1 (blue) visualises a ResNet50 architecture as an example backbone [11] where we extract FPN levels from network layers P_3 to P_7 . Across different scales of the FPN, anchors cover area sizes of 32^2 to 512^2 distributed over pyramid levels from P_3 to P_7 , respectively. We embed our TCM and SCM after the last layer of the backbone for a maximal receptive field during blending.

Detection Head. The RetinaNet based detection head contains two streams, *i.e.* a localisation stream and classification stream. The classification subnet (coloured grey in

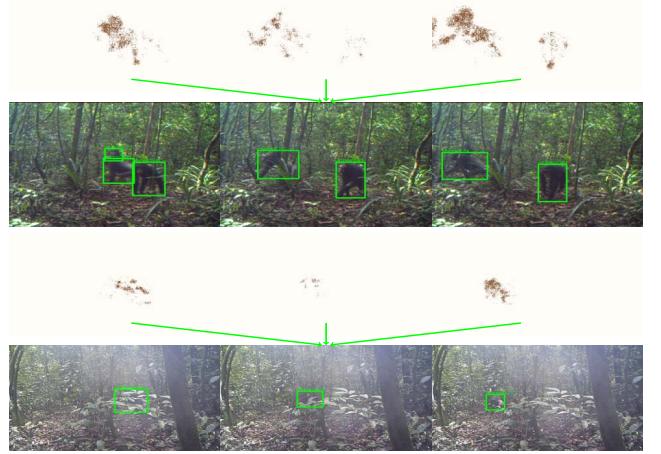


Figure 4: **Visualisation of Temporal Attention Maps for Two Example Scenes.** First and third rows depict attention maps when projected back into the image domain. Second and fourth row show the associated frames with green bounding boxes showing predictions by the network with $T = 3$. Green lines indicate that shown attention maps contribute to detections in the central frame.

Figure 1) predicts a likelihood vector of class membership for each anchor box. This network is fully convolutional and operates on top of the 5 FPN layer outputs P_3 to P_7 . It is composed of 4 cascaded 3×3 convolutional layers, each with 256 channels, and followed by a ReLU activation. The final layer uses 3×3 convolutions with a Sigmoid function as activation for producing the final classification output. In the localisation subnet (coloured black in Figure 1) the final layer yields $4 \times$ channel outputs describing bounding box locations. We assign initial anchor areas ranging from 32^2 to 512^2 elements corresponding to the pyramid levels P_3 to P_7 , respectively. Subsequent rescaling uses the ratios of $[2^0, 2^{1/3}, 2^{2/3}]$ and transforms anchors with three aspect ratios [1:2, 1:1, 2:1]. In total there are 9 anchors per level. Focal loss from [19] is adopted to tackle the foreground and background imbalance problem.

We also implemented a stronger, cascaded detection head [3] for maximum detection quality. The cascaded detection head contains four stages, one for an RPN, which regresses class-agnostic regions of interest, and three for detection with different IoU ground truth assignments parameterized as [0.5, 0.6, 0.7]. Note that the three cascaded detection heads are fully convolutional and ROI alignment is used for ROI feature pooling.

3.2. Training Details

We train the network end-to-end in two separate stages (see Figure 5). First, we pre-train our model on the ImageNet VID dataset for 14 epochs, and then the entire model is trained with synchronised batch normalisation and training data augmentation (sequence-based random brightness, horizontal flip, central crop) for 14 epochs on the 400 video training portion of the Pan Africa dataset. The remaining 100 clips are used for testing.

During the various experiments, we use short-term video snippets and sparse sampling of between $T = 3$ to $T = 8$ frames whose shorter size is resized to at least 512 and up to 800 pixels horizontally keeping the aspect ratio constant by padding if necessary. During training, each mini-batch contains 8 snippets, that is 1 per deployed Nvidia Tesla P100 GPU. The learning rate is warmed up for the first 500 iterations of each training phase from 0.002 to 0.01 and decreased 10 times at epoch 6 and 11, respectively.

We use SGD with the momentum of 0.9 as the optimizer and utilise ImageNet pre-trained weights to initialise the backbone. FPN layers and all other convolutional layers are initialised by normal distributions.

4. Experimental Results

We evaluate the proposed system both quantitatively and qualitatively, and perform detailed ablation studies. We also perform a validation on the full ILSVRC 2015 VID data corpus to demonstrate wider applicability.

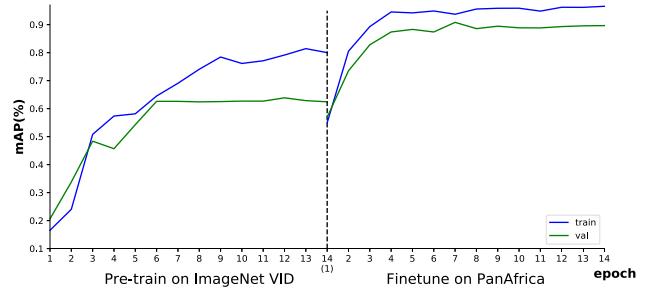


Figure 5: Example Training and Validation Performance. mAP plots for both training (blue) and validation (green) of the ResNet50 RetinaNet SCM+TCM setup along the 28 epochs of optimisation split into two distinct phases, i.e. pre-training on ImageNet VID, and fine-tuning on the Pan Africa dataset.

4.1. Quantitative Evaluation

Using the test portion of the Pan Africa dataset, we first evaluate single frame detectors as our baselines. For evaluation, we compute the average precision (AP) as the area under precision-recall curve and then report the mean of AP (mAP) for the classes in the validation set.

We re-implemented RetinaNet [19] with ResNet50 and ResNet101 baselines and, as shown in Table 1, these two architectures achieve 80.79% and 85.25% mAP, respectively. Whilst adding a SCM component to the ResNet50 RetinaNet setup improves performance slightly to 81.21%, the current state-of-the-art Cascade-RCNN [3] outperforms this setup and older single frame baselines, achieving 88.31% on the Pan Africa test data. The basic addition of a TCM component, on the other hand, produces significant performance improvements by up to 10.02% \uparrow for ResNet50 RetinaNet, and still 2.86% \uparrow for Cascade-RCNN.

The training process for a full SCM+TCM setup, as quantified earlier in Figure 5 for instance, furthermore reveals that the generalisation gap of such a model (mAP distance between blue and green curves) narrows significantly for the PanAfrica dataset compared to the ImageNet VID

Backbone and FPN Layout	T_{train}	T_{test}	mAP(%)
baseline[19]	\	\	80.79
Res50 +SCM	\	\	81.21
RetinaNet +TCM	7	21	90.02
+SCM+TCM	7	21	90.81
Res101 baseline[19]	\	\	85.25
RetinaNet +SCM+TCM	5	21	90.21
ResX101 baseline[3]	\	\	88.31
Cascade +SCM+TCM	3	21	91.17

Table 1: Pan Africa Performance Results. Boosts in mean average precision (mAP) on the Pan Africa test dataset when applying the proposed TCM & SCM components with various state-of-the-art FPN architectures.



Figure 6: Examples of Positive Attention Component Impact. Depicted are ground truth labels (red) and detections (green) for a ResNet50 FPN with (rows 2, 4, and 6) and without (rows 1, 3, and 5) SCM+TCM components. Note the clear detection improvements for partially occluded great apes (top two examples) and under challenging lighting (bottom example).

pre-training in late stages. Narrow generalisation gaps indicate that a model is particularly capable of carrying over learned inference strategies to unseen instances of the particular dataset. For the case at hand, this and the improved overall mAP indicate that fine-tuning is indeed successful.

We also validated a basic ResNet50 RetinaNet version that uses SCM+TCM on the full ILSVRC 2015 VID corpus to show that our rather simple and flexible SCM+TCM network extension can achieve strong results also in this general detection setting at 63.85% mAP. This ranks within the top three of the original mAP based competition [12].

4.2. Qualitative Observations on Pan Africa Data

The Pan Africa dataset contains many scenes where illumination, occlusions, noise or animal camouflage effects make it challenging to recognise animals (as seen before in Figure 2). We found that the SCM+TCM setup consistently improves detection robustness compared to baselines in such cases. These improvements make a significant contribution to the overall quantitative results reported before. Figure 6 provides examples of successful cases where per-frame accurate animal detection is achieved by the SCM+TCM components in the presence of partial occlusion or challenging lighting. However, as depicted in the bottom examples of Figure 7, a number of particular animal

appearances remain challenging to detect despite the availability of spatial and temporal context information.

4.3. Ablation Study

We conduct an ablation study on both the Pan Africa and the ILSVRC 2015 VID datasets in order to quantify in detail the impact of the key system parameters. All of the results are reported based on ResNet50+SCM+TCM setup with a RetinaNet detection head.

SCM+TCM Insertion Point. Considering arguments in [4], we acknowledge that the embedding position of our module along the original backbone network is important. To determine the best insertion point, we test three possible positions for insertion into ResNet50: after the 3×3 convolution of the last ResNet block; after the last 1×1 convolution before the residual addition; and after the residual addition. As shown in Table 2(a) and Table 3(a), we find that the last option is superior and can gain 6.01% \uparrow mAP on Pan Africa and 6.42% \uparrow mAP on VID improvement when compared with insertion after the last 1×1 convolution. This indicates that the final residual addition in the base network *does* provide useful extra information.

Temporal Frame Support. Different choices on supporting frames T are ablated in Table 2(b) and 3(b). Results confirm that wider windows for temporal integration do indeed benefit detection, however, particularly during training, GPU sizes limit possible choices of T . During testing, longer exposures are possible. Quantifying the effect of varying test exposures we find that, for $T_{test} = 21$ frames compared to $T_{test} = 5$, there is a 1.96% \uparrow and 3.69% \uparrow improvement for fixed $T_{train} = 7$ on the Pan Africa and ILSVRC 2015 VID datasets, respectively.

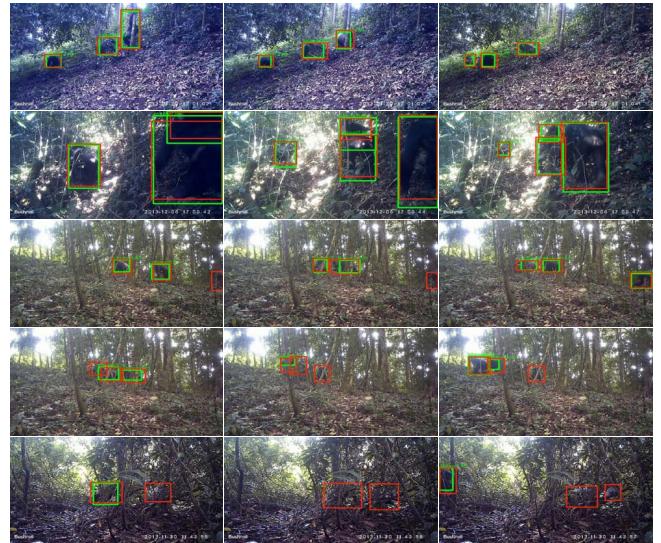


Figure 7: Success and Failure Cases. Depicted are ground truth labels (red) and detections (green) in frames of various example clips using the best performing model ResX101 Cascade with SCM+TCM components.

Embedding Strategy in the TCM. We found that when applying an embedding strategy in the TCM where only the current main frame features (indicated as black arrows of the TCM in Figure 1) are backpropagated and reference features are used, but not backpropagated, then a marginally improved overall performance can be observed (marked as *Main & Refs* in contrast to *Positional* in Tables 2 and 3).

5. Conclusion and Implications

In this paper we proposed the first multi-frame video object detection framework trained and evaluated for detecting great apes utilising their full body morphology. We demonstrated that the framework is applicable to challenging cam-

era trap footage taken in complex jungle environments. We introduced two self-attention driven feature blending components operating in both the spatial and the temporal domains to facilitate detection under heavy partial occlusions and challenging lighting variations.

We showed that this novel and flexible extension performs robustly at 91.17% mAP on a real world Pan Africa camera trap 500 video dataset, which we labelled accurately with animal ground truth annotations for 180K frames. We conducted detailed ablation studies on our method and showed that the setup significantly outperforms state-of-the-art frame based detectors. For general evaluation beyond the task at hand, we also performed a validation on the ILSVRC 2015 VID data corpus to demonstrate significant performance on non-specialised video object detection.

We note that currently ecological camera trap studies are widely conducted by manual inspection, although great ape face detection [2, 21] has been used for ecological surveys before [7] and DrivenData [8] hosted a recent challenge to classify jungle camera trap clips by species, without detecting animals and their location in frames explicitly.

The presented system, in contrast, provides explicit animal locations and is independent of visibility constraints regarding the animal’s face. It adds a new capability of detection and localisation of animals partly occluded by vegetation at adequate performance levels. Whilst tests against other current video detection frameworks are outstanding and will form part of our future work, we conclude that the presented system is ready to assist human camera trap inspection efforts.

Acknowledgements

We would like to thank the entire team of the Pan African Programme: ‘The Cultured Chimpanzee’ and its collaborators for allowing the use of their data for this paper. Please contact the copyright holder Pan African Programme at <http://panafrican.eva.mpg.de> to obtain the dataset. Particularly, we thank: H Kuehl, C Boesch, M Arandjelovic, and P Dieguez. We would also like to thank: K Zuberbuehler, K Corogenes, E Normand, V Vergnes, A Meier, J Lapuente, D Dowd, S Jones, V Leinert, E Wessling, H Eshuis, K Langergraber, S Angedakin, S Marrocoli, K Dierks, T C Hicks, J Hart, K Lee, and M Murai. Thanks also to the team at <https://www.chimpandsee.org>. The work that allowed for the collection of the dataset was funded by the Max Planck Society, Max Planck Society Innovation Fund, and Heinz L. Krekeler. In this respect we would also like to thank: Foundation Ministre de la Recherche Scientifique, and Ministre des Eaux et Forts in Cote d’Ivoire; Institut Congolais pour la Conservation de la Nature and Ministre de la Recherche Scientifique in DR Congo; Forestry Development Authority in Liberia; Direction des Eaux, Forts Chasses et de la Conservation des Sols, Senegal; and Uganda National Council for Science and Technology, Uganda Wildlife Authority, National Forestry Authority in Uganda.

Table 2: **Pan Africa Dataset Ablation Study.** Key system design choices tested on a ResNet50+SCM+TCM setup with a RetinaNet detection head.

(a) Position					
Embedding	position	T_{train}	T_{test}	mAP(%)	
Positional	after add	8	8	88.21	
Positional	after 1x1	8	8	82.20	
Positional	after 3x3	8	8	87.75	
(b) Supporting Frames					
Embedding	position	T_{train}	T_{test}	mAP(%)	
Main & Refs	after add	7	21	90.81	
Main & Refs	after add	7	5	88.85	
Main & Refs	after add	3	5	87.76	
(c) Embedding Strategy					
Embedding	position	T_{train}	T_{test}	mAP(%)	
Positional	after add	7	21	88.61	
Main & Refs	after add	7	21	90.81	

Table 3: **ILSVRC 2015 VID Dataset Ablation Study.** Key system design choices tested on a ResNet50+SCM+TCM setup with a RetinaNet detection head.

References

- [1] Luca Bertinetto, Jack Valmadre, Joo F. Henriques, Andrea Vedaldi, and Philip H.S. Torr. Fully-convolutional siamese networks for object tracking. In *ECCV*, pages 850–865, 2016.
- [2] Clemens-Alexander Brust, Tilo Burghardt, Milou Groenewberg, Christoph Käding, Hjalmar S Kühl, Marie L Manguette, and Joachim Denzler. Towards Automated Visual Monitoring of Individual Gorillas in the Wild. In *ICCVW*, pages 2820–2830, 2017.
- [3] Zhaowei Cai and Nuno Vasconcelos. Cascade R-CNN: Delving into High Quality Object Detection. In *CVPR*, pages 6154–6162, 2018.
- [4] Yue Cao, Jiaru Xu, Stephen Lin, Fangyun Wei, and Han Hu. GCNet: Non-local Networks Meet Squeeze-Excitation Networks and Beyond. *arXiv:1904.11492 [cs]*, 2019.
- [5] Joo Carreira and Andrew Zisserman. Quo Vadis, action recognition? A new model and the kinetics dataset. In *CVPR*, pages 4724–4733, 2017.
- [6] Kai Chen, Jiaqi Wang, Shuo Yang, Xingcheng Zhang, Yuanjun Xiong, Chen Change Loy, and Dahua Lin. Optimizing Video Object Detection via a Scale-Time Lattice. In *CVPR*, pages 7814–7823, 2018.
- [7] Anne-Sophie Crunchant, Monika Egerer, Alexander Loos, Tilo Burghardt, Klaus Zuberbühler, Katherine Corogenes, Vera Leinert, Lars Kulik, and Hjalmar S Kühl. Automated face detection for occurrence and occupancy estimation in chimpanzees. *American journal of primatology*, 79(3):e22627, 2017.
- [8] DrivenData. Competition: Pri-matrix Factorization [online]. <https://www.drivendata.org/competitions/49/deep-learning-camera-trap-animals/>. Accessed: 2019-07-30.
- [9] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Detect to Track and Track to Detect. In *ICCV*, pages 3038–3046, 2017.
- [10] Wei Han, Pooya Khorrami, Tom Le Paine, Prajit Ramachandran, Mohammad Babaieizadeh, Honghui Shi, Jianan Li, Shuicheng Yan, and Thomas S. Huang. Seq-NMS for Video Object Detection. Technical report, 2016.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [12] ImageNet. ImageNet Challenge 2015: Object Detection from Video [online]. <http://image-net.org/challenges/LSVRC/2015/index#vid>. Accessed: 2019-07-30.
- [13] Kai Kang, Hongsheng Li, Junjie Yan, Xingyu Zeng, Bin Yang, Tong Xiao, Cong Zhang, Ruohui Wang, Xiaogang Wang, and Wanli Ouyang. T-CNN: Tubelets with Convolutional Neural Networks for Object Detection from Videos. *IEEE T-CSVT*, 28(10):2896–2907, 2018.
- [14] Kai Kang, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. Object detection from video tubelets with convolutional neural networks. In *CVPR*, pages 817–825, 2016.
- [15] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Fei Fei Li. Large-scale video classification with convolutional neural networks. In *CVPR*, pages 1725–1732, 2014.
- [16] Hjalmar S. Kuehl and Tilo Burghardt. Animal biometrics: quantifying and detecting phenotypic appearance. *TREE*, 28(7):432 – 441, 2013.
- [17] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. SiamRPN++: Evolution of Siamese Visual Tracking with Very Deep Networks. *arXiv:1812.11703*, 2018.
- [18] Tsung Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 936–944, 2017.
- [19] Tsung Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal Loss for Dense Object Detection. In *ICCV*, pages 2999–3007, 2017.
- [20] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng Yang Fu, and Alexander C. Berg. SSD: Single shot multibox detector. In *ECCV*, volume 9905 LNCS, pages 21–37, 2016.
- [21] Alexander Loos and Andreas Ernst. An automated chimpanzee identification system using face detection and recognition. *EURASIP Journal on Image and Video Processing*, 2013(1):49, 2013.
- [22] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, pages 779–788, 2016.
- [23] Joseph Redmon and Ali Farhadi. YOLO9000: Better, faster, stronger. In *CVPR*, pages 6517–6525, 2017.
- [24] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *PAMI*, 39(6):1137–1149, 2017.
- [25] Karen Simonyan and Andrew Zisserman. Two-Stream Convolutional Networks for Action Recognition in Videos. In *NIPS*, pages 568–576, 2014.
- [26] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3D convolutional networks. In *ICCV*, pages 4489–4497, 2015.
- [27] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, pages 20–36, 2016.
- [28] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local Neural Networks. In *CVPR*, pages 7794–7803, 2018.
- [29] Xizhou Zhu, Yujie Wang, Jifeng Dai, Lu Yuan, and Yichen Wei. Flow-Guided Feature Aggregation for Video Object Detection. In *ICCV*, pages 408–417, 2017.
- [30] Xizhou Zhu, Yuwen Xiong, Jifeng Dai, Lu Yuan, and Yichen Wei. Deep feature flow for video recognition. In *CVPR*, pages 4141–4150, 2017.
- [31] Zheng Zhu, Qiang Wang, Bo Li, Wei Wu, Junjie Yan, and Weiming Hu. Distractor-aware Siamese Networks for Visual Object Tracking. In *ECCV*, pages 101–117, 2018.