

CONTINUAL TRANSFORMERS: REDUNDANCY-FREE ATTENTION FOR ONLINE INFERENCE

Lukas Hedegaard, Arian Bakhtiarnia & Alexandros Iosifidis

Department of Electrical and Computer Engineering

Aarhus University

Aarhus, Denmark

{lhm, arianbakh, ai}@ece.au.dk

ABSTRACT

Transformers in their common form are inherently limited to operate on whole token sequences rather than on one token at a time. Consequently, their use during online inference on time-series data entails considerable redundancy due to the overlap in successive token sequences. In this work, we propose novel formulations of the Scaled Dot-Product Attention, which enable Transformers to perform efficient online token-by-token inference on a continual input stream. Importantly, our modifications are purely to the order of computations, while the outputs and learned weights are identical to those of the original Transformer Encoder. We validate our *Continual* Transformer Encoder with experiments on the THUMOS14, TVSeries and GTZAN datasets with remarkable results: Our *Continual* one- and two-block architectures reduce the floating point operations per prediction by up to $63\times$ and $2.6\times$, respectively, while retaining predictive performance.

1 INTRODUCTION

Many real-life usage scenarios such as the perception in self-driving cars and live monitoring of critical resources process a continual stream of inputs and require near-instantaneous predictions per time-step. This stands in contrast to what many common benchmarks for deep learning evaluate, namely the operation on distinct batches of data with no inter-batch relationships. Consequently, a plethora of methods have been developed (Ji et al., 2013; Carreira & Zisserman, 2017; Varol et al., 2018; Yan et al., 2018; Heidari & Iosifidis, 2021; Vaswani et al., 2017; Arnab et al., 2021; Bakhtiarnia et al., 2021b), which focus on batch-wise processing, but fail to optimise for online operation, where new information (e.g., a video frame / token) arrives at each step from a continual input stream, and future information is not available at the current time-step. We need a class of networks, which operate efficiently on *both batches of data and on continual streams*.

Accordingly, we propose a reformulation of the Transformer Encoder as a Continual Inference Network (CIN, Section 2.1) which accelerates the stream processing on time-series data, while retaining weight-compatibility. Specifically, we derive two variants of Continual Scaled Dot-Product Attention (SDA) for the cases where prior output tokens *should* and *should not* be updated after observing a new input token. Notably, our attention formulations reduce the per-step cost of SDA (Vaswani et al., 2017) from time complexity $\mathcal{O}(n^2d)$ to $\mathcal{O}(nd)$ and memory complexity $\mathcal{O}(n^2)$ to $\mathcal{O}(nd)$ and are readily embedded into Continual Multi-Head Attention (MHA) and Continual Transformer Encoder blocks. Finally, we propose the use of Recycling Positional Encoding to accommodate progressive caching of partial attention results for continual data streams.

Due to the interdependence of SDA outputs, Continual Transformers are most efficient for shallow architectures. Shallow Transformers have many applications such as augmentations of CNNs (Touvron et al., 2021), light-weight Natural Language Processing (Cornia et al., 2020), fusion operations in multi-modal (e.g. audio-visual) settings (Chumachenko et al., 2022) and early exit branches in multi-exit architectures (Bakhtiarnia et al., 2021a;b). In our experiments¹, we validate their exceptional efficiency improvements on common benchmarks in Online Action Detection (Idrees et al., 2017) and Online Audio Classification (Tzanetakis et al., 2001).

¹Source code: <https://github.com/lukashedegaard/continual-transformers>.

Figure 1: Multi-block Continual Transformer Encoder with Recycling Positional Encoding. For $b > 2$ blocks, regular Transformer Encoder blocks can be added between an initial Continual Retroactive block and a final Single-Output block. A class-token may be used after the initial block.

2 RELATED WORK

2.1 CONTINUAL INFERENCE NETWORKS

Definition (Continual Inference Network). A Deep Neural Network, which

- is capable of continual step inference without computational redundancy,
- is capable of batch inference corresponding to a non-continual Neural Network,
- produces identical outputs for batch- and step inference given identical receptive fields,
- uses one set of trainable parameters for both batch and step inference.

These requirements ensure that a Neural Network has broad applicability for both (of ine) batch-wise inference (i.e., most research benchmarks) and online stream processing. While non-CINs can operate on streams of data by caching prior steps in a first-in first-out (FIFO) queue and aggregating them to a full (spatio-)temporal input, which is processed similarly to an of ine batch, this entails computational redundancy in proportion with the sequence length. CINs perform step-wise inference without such caching and repeat computation. Uni-directional Recurrent Neural Networks are an example of Continual Inference Networks. Their default mode of operation is by time-step and they are easily applied to spatio-temporal batches of data by concatenation of the step-wise outputs. Recently, a modification to the spatio-temporal 3D convolution was proposed (Hedegaard & Joshi, 2021), which enables existing 3D CNNs to operate efficiently during continual inference. A similar principle was used to enhance Spatio-temporal Graph Convolutions as well (Hedegaard et al., 2022). In Section 3, we derive a CIN formulation for Transformer Encoders.

2.2 TRANSFORMER ARCHITECTURES

Initially proposed for sequence-to-sequence modelling in Natural Language Processing, the Transformer (Vaswani et al., 2017) has become a canonical building block in many applications of Deep Learning, including Computer Vision (Dosovitskiy et al., 2021; Arnab et al., 2021; Wang et al., 2021; Carion et al., 2020) and Audio Classification (Gong et al., 2021). Their success can be partly attributed to reduced inductive bias compared with CNNs and RNNs, which allows better adaptations when sufficiently large datasets are available; the Scaled Dot-Product Attention (SDA) maps a set of input tokens to a set of outputs without inherent preconceptions. However, this many-to-many attention exhibits quadratic growth in time and space complexity with the token size of the set.

A great deal of research has sought to improve the efficiency of Transformers (Tay et al., 2020). Block-wise or Chunking methods such as Image Transformer (Parmar et al., 2018) and Vision Transformer (Dosovitskiy et al., 2021) group up entities of a local receptive field into a single block, reducing the $O(n^2)$ complexity to $O(n_b^2)$, where $n_b < n$ is the number of blocks. Techniques such as sliding windows, dilation and pooling can be used to achieve a similar effect (Beltagy et al., 2020). The Reformer (Kitaev et al., 2020) reduces the complexity to $O(n \log n)$ by learning group-

ings in a data-driven manner via Locality-Sensitive Hashing (LSH). A different paradigm aims to derive approximations of the self-attention matrix. Methods such as Linformer (Wang et al., 2020), Nyströmformer (Xiong et al., 2021) and Performer (Choromanski et al., 2021) reduce the complexity from $O(n^2)$ to $O(n)$. Unlike these efforts, our approach produces the same computational outputs for temporal sequences as the original Multi-Head Attention.

3 CONTINUAL TRANSFORMERS

In this work, we examine the use of Transformer Encoders for stream-processing, where we receive one token per time-step. Specifically, the query, key and value inputs constitute a continual stream of d -dimensional tokens and we wish to compute the outputs for each step immediately considering $n - 1$ prior tokens. We begin our exposition in Section 3.1 by considering the Scaled Dot-Product Attention (SDA) for this task. To alleviate the inefficiencies of SDA, we propose two alternative computational sequences in Section 3.2 and Section 3.3 and compare them to SDA in Section 3.4. Finally, sections 3.5-3.8 build up the full architecture, and discuss architectural considerations.

3.1 REGULAR SCALED DOT-PRODUCT ATTENTION

Denoting query, key, and value sequence matrices $Q, K, V \in \mathbb{R}^{n \times d}$, the regular Scaled Dot-Product Attention first defined by Vaswani et al. (2017) can be written as:

$$\text{Att}(Q; K; V) = D^{-1}AV \quad A = \exp \left(\frac{QK^T}{d} \right) \quad D = \text{diag} A \mathbf{1}_n^T; \quad (1)$$

where $A \in \mathbb{R}^{n \times n}$ and $\mathbf{1}_n$ is a row-vector of ones. In each time-step, we can update Q , and V by discarding their oldest token and prepending a new one in a FIFO manner. This is a common implementation for step-wise inference, e.g. found in the PyTorch library (Ott et al., 2019).

Each time-step results in $2n^2d + 2nd$ multiplications, $2n^2d - nd - n$ additions, and d^2 exponentiations as accounted for in Appendix A.1, which amounts to a time complexity of $O(d^2)$ and a $O(n^2)$ memory complexity originating from the transient feature-map. Furthermore, a constant-sized cache of size $(n - 1)d$ is needed to store the $n - 1$ latest tokens in Q, K and V . We could avoid considerable redundancy by caching A directly. However, this comes with a memory penalty of $(n - 1)^2$. Fortunately, another computational scheme can be devised.

3.2 CONTINUAL RETROACTIVE SCALED DOT-PRODUCT ATTENTION

We can compute $D^{-1}AV$ in a step-wise manner using the latest query, key, and value steps, $q_{\text{new}}, k_{\text{new}}, v_{\text{new}} \in \mathbb{R}^{1 \times d}$, alongside appropriately cached partial results. The softmax normalisation with D^{-1} can be efficiently implemented via column-aligned element-wise multiplications (denoted by \odot hereafter) of a column-vector $\mathbf{d} = A \mathbf{1}_n^T$. If we cache the $n - 1$ values for the prior step tokens, i.e. $\mathbf{d}_{\text{mem}} = A_{\text{prev}}^{(n+1:)} \mathbf{1}_n^T$, alongside Q and K , we can define the step update as:

$$\mathbf{d}^{(n+1:)} = \mathbf{d}_{\text{mem}}^{(n+2:0)} \odot \exp \left(\frac{Q_{\text{mem}} k_{\text{old}}^T}{d} \right) + \exp \left(\frac{Q_{\text{mem}} k_{\text{new}}^T}{d} \right) \quad (2)$$

$$\mathbf{d}^{(0)} = \exp \left(\frac{q_{\text{new}} (K_{\text{mem}} k_{\text{new}})^T}{d} \right) \mathbf{1}_n^T; \quad (3)$$

where $Q_{\text{mem}}, (K_{\text{mem}})$ are the $n - 1$ prior query (key) tokens, k_{old} is the key from n steps ago, and k denotes concatenation of matrices along the first dimension. Negative indices indicate prior time-steps. An update for AV can likewise be defined as a function of the $n - 1$ prior values AV_{mem} :

$$AV^{(n+1:)} = AV_{\text{mem}}^{(n+2:0)} \odot \exp \left(\frac{Q_{\text{mem}} k_{\text{old}}^T v_{\text{old}}}{d} \right) + \exp \left(\frac{Q_{\text{mem}} k_{\text{new}}^T v_{\text{new}}}{d} \right) \quad (4)$$

$$AV^{(0)} = \exp \left(\frac{q_{\text{new}} (K_{\text{mem}} k_{\text{new}})^T v_{\text{new}}}{d} \right) : \quad (5)$$

Finally, we compute the Continual Retroactive Attention output in the usual manner:

$$\text{CoReAtt}(q_{\text{new}}, k_{\text{new}}, v_{\text{new}}) = \mathbf{d}^{-1} AV; \quad (6)$$

A visual depiction of these update steps is provided in Appendix A.2. A time-step can now be computed with $7nd + 2n - 3d$ multiplications, $6nd + 3n - 6d - 3$ additions, and $dn - 2$ exponentials. This time complexity $O(nd)$ per step and $O(nd)$ memory complexity is a significant improvement over the prior $O(n^2d)$ and $O(n^2)$ complexities in Section 3.1.

3.3 CONTINUAL SINGLE-OUTPUT SCALED DOT-PRODUCT ATTENTION

Both the Regular and Continual Retroactive Dot-Product Attentions produce attention outputs for the current step, as well as 1 retroactively updated steps. In cases where retroactive updates are not needed, we can simplify the computation greatly via a Continual Single-Output Dot-Product Attention (CoSiAtt). In essence, the regular SDA is reused, but prior values and v are cached between steps (as in (Ott et al., 2019)), and only the attention corresponding to a single query token q is computed:

$$\text{CoSiAtt}(q; k_{\text{new}}; v_{\text{new}}) = a(V_{\text{mem}} k v_{\text{new}}) = a \mathbf{1}_n^T; \quad a = \exp \frac{q}{d} (K_{\text{mem}} k k_{\text{new}})^T \quad (7)$$

A step output is computed with $2nd + 2d$ multiplications, $2nd - d - 1$ additions, and n exponentials. The time- and memory complexities remain $\mathcal{O}(nd)$ per step. Using the (leading) query q_{new} as input, the attention is purely causal. Alternatively, prior (lagging) query vectors could be cached and used as query input, though this would introduce a network delay.

3.4 COMPARISON OF SCALED DOT-PRODUCT ATTENTIONS

Assuming $n - 1$ prior q , k and v steps have been calculated by the Continual SDA modules, and that $Q = (Q_{\text{mem}} k q_{\text{new}})$, $K = (K_{\text{mem}} k k_{\text{new}})$, and $V = (V_{\text{mem}} k v_{\text{new}})$, we have the correspondence:

$$\text{Att}(Q; K; V)^{(t)} = \text{CoReAtt}(q_{\text{new}}; k_{\text{new}}; v_{\text{new}})^{(t)} = \text{CoSiAtt}(q_t; k_{\text{new}}; v_{\text{new}}) \quad (8)$$

Here, q_t is the t^{th} row of Q , i.e. $Q^{(t)}$. During stream processing, the complexity of the Continual Retroactive SDA scales significantly more favourably than the regular SDA. For example, the floating point operations (FLOPs) are reduced by 31% when $n = d = 100$ and 30% when $n = d = 1000$. If retroactive output updates are not needed, the Continual Single-Output SDA reduces FLOPs by respectively 100% and 1000%. The scaling properties are detailed in Appendix A.1.

3.5 CONTINUAL MULTI-HEAD ATTENTION

Continual Scaled Dot-Product Attentions can replace regular SDA's directly in a Multi-Head Attention (MHA). Given a new query, key, and value q, k, v , the Continual MHA is defined as

$$\text{CoMHA}(q; k; v) = \sum_{i=0}^{h-1} k \text{CoAtt}(qW_Q^i; kW_K^i; vW_V^i) W_O; \quad (9)$$

where k denotes concatenation of heads and $W_Q^i, W_K^i \in \mathbb{R}^{d \times d_k = h}$, $W_V^i \in \mathbb{R}^{d \times d_v = h}$, and $W_O \in \mathbb{R}^{d_v \times d_o}$ are projection matrices of head i . CoAtt can be either CoReAtt or CoSiAtt .

3.6 CONTINUAL TRANSFORMER ENCODER

A Continual MHA block can be integrated in a Continual Transformer Encoder block as follows:

$$z = \text{LayerNorm}(y + \text{FF}(y)); \quad y = \text{LayerNorm}(x + \text{CoMHA}(x; x; x)); \quad (10)$$

where x corresponds to the newest step input and Se selects a single (last) token iff CoMHA is used, or selects all tokens otherwise. FF is a two-layer feed-forward network with weights $W_1; W_2$, biases $w_1; w_2$, and a activation function σ , i.e. $\text{FF}(x) = \sigma(xW_1 + w_1)W_2 + w_2$. Aside from the residual selection, this is identical to common Transformer Encoder implementations (Vaswani et al., 2017; Dosovitskiy et al., 2021).

3.7 RECYCLING POSITIONAL ENCODING

Since a Transformer Encoder does not provide positional bias, it is common to augment a token with a positional encoding, i.e. $x_t = x_i \oplus p_i$, where \oplus could be addition or concatenation. In regular Transformers, the index denotes a position in a sequence rather than a position in time. However, this static positional assignment is problematic in the context of continual inference; the last token at time $t = 0$ will be the next-to-last token at time $t = 1$, and thus in need of a different

positional encoding than in the prior time-step. Instead, CINs require dynamic positions. There have been multiple prior works (Shaw et al., 2018; Huang et al., 2019; Dai et al., 2019) which create relative encodings by augmenting the SDA with positional offsets between query and keys, i.e. $A = \exp(\frac{QK^T}{\sqrt{d}} + P)$. While a similar modification to the continual attentions is possible, it is incompatible with the regular SDA in Eq. (1). Instead, we use Recycling Positional Encoding (RPE), which lets the positional encoding follow each token in time and recycles old encodings:

$$x_t = x_t + p_{t \bmod T}; \quad t = (t_{-1} + 1) \bmod T; \quad (11)$$

where T is the number of encodings. While RPE does not specify relative encodings explicitly, the absolute positional interpretation of each token changes dynamically when a new token arrives. In practice, the network learns relative, shift-invariant positional information by training with random shifts for each batch. Random shifts during training were recently explored in (Kiyono et al., 2021; Likhomanenko et al., 2021; Dehghani et al., 2019) as well. RPE can use either learned or predefined encodings. In the latter case, Cyclic Positional Encoding (Ma et al., 2021), a sinusoidal encoding inspired by Gray code, is a good fit. If we reuse the encoding immediately after an old token has “slided out”, i.e. $T = n$, a token will have the same positional encoding relative to another whether it was m steps older or m steps newer. The positional ambiguity can be avoided by extending the number of positional tokens to $T = 2n - 1$. We explore both options in Section 4.1.2.

3.8 ARCHITECTURAL CONSIDERATIONS

Block count In Section 3.4, we observed an exact correspondence between the results of the continual and regular SDA layers. However, the correspondence does not necessarily hold for stacked layers. Consider the result of stacking two Continual Single-Output Transformer Encoder blocks. While the first block outputs a step that is identical to that in a corresponding regular block, the second block would have been initialised with prior step-wise inputs, which were the result of prior input windows instead of the current one; the correspondence would not hold. Though it is not convertible to/from a regular Transformer Encoder, the stacked Single-Output Transformer Encoder architecture has the merit of efficiency. This was exploited in Transformer-XL (Dai et al., 2019). Given a single step input, the Continual Retroactive Transformer Encoder block produces output tokens corresponding to the entire observed sequence inside the window. Due to this one-to-many input-output mapping, it is not possible to stack multiple such layers. Nevertheless, it can be used in conjunction with a Continual Single-Output Transformer Encoder with optional regular Transformer Encoder blocks in between as illustrated in Fig. 1. The Regular Transformer Encoder blocks in the resulting architecture have a significantly larger computational complexity than the Continual Retroactive and Single-Output blocks. Consequently, we recommend that Continual Transformer Encoders be used primarily in lightweight architectures with one or two blocks unless compatibility with non-continual Transformers is not required and only Single-output blocks are used.

Class token It is common to add a class token as input to transformers (Devlin et al., 2019; Dosovitskiy et al., 2021), which accumulates information from other tokens prior to classification. However, it cannot be used directly with CINs, as this would effectively double the number of input steps. In practice, it can be employed in Continual multi-block Transformer Encoders as input to the second block (see Fig. 1), but this placement limits class token interaction with downstream layers. It can also be used for one-block Transformer Encoders if the value token is omitted as input.

Peak memory reduction trick The FLOPs for $\text{Att}(Q; K; V)$ are exactly n times those of $\text{CoSiAtt}(q; k_{\text{new}}; v_{\text{new}})$. Comparing their memory complexity, the regular SDA is $\mathcal{O}(n^2)$, while the Single-output SDA is $\mathcal{O}(nd)$. In practical applications where system memory is limited, we may thus reduce the maximum memory requirement of the computational device at inference by up to $d=n$ (assuming $n \geq d$) by computing each row of the attention individually. However, this may reduce throughput due to reduced parallelism.

4 EXPERIMENTS

We provide case studies within two perception disciplines, namely Online Action Detection (Section 4.1) and Audio Classification (Section 4.2). In each case, we will start with a brief overview of the field, followed by experiments and results.

4.1 ONLINE ACTION DETECTION

Online Action Detection (OAD) (De Geest et al., 2016) entails the per-frame classification of human actions in a video stream as they happen without the ability to change prior predictions nor to use future information. This is fundamentally more restrictive than Temporal Action Localisation, where the whole video clip is processed before start and end frames of an action are determined (Shou et al., 2016; Xu et al., 2017; Shou et al., 2017; Wu et al., 2019).

The dominant design in OAD works at the time of writing is to employ a two-stream Convolutional Neural Network as backbone for frame-wise feature extraction with RGB images as inputs in one stream and Optical Flow fields in the other (Gao et al., 2017; Xu et al., 2019; Eun et al., 2020; Wang et al., 2021; Xu et al., 2021). On top of these, OAD methods encode temporal information and perform predictions per time-step, e.g. by means of RNNs (Gao et al., 2017; Xu et al., 2019; Eun et al., 2020) or Transformers (Wang et al., 2021; Xu et al., 2021). Alongside the action detection for the current frame, an action anticipation task may be learned in parallel by means of decoder structures, as this has been found to improve the primary OAD task.

Unlike RNNs, an output update for the regular SDA in a Transformer block cannot be dynamically computed for a single step by feeding successive video frames. Instead, prior step features must be cached, re-loaded and re-processed by the Transformer in each step in correspondence with a predefined window-size of prior steps. As laid out in Section 3.8, Continual Transformers are especially efficient when either one or two Continual Transformer Encoder blocks are used. Accordingly, we start our experiments with a set of ablation studies to simplify a recent transformer-based architecture, the OadTR (Wang et al., 2021). We further investigate the impact of ablating class token position and the use of Recycling Positional Encoding and compare different RPE schemes for Continual Transformers. Finally, we evaluate our configurations on two widely used OAD datasets, THUMOS14 (Idrees et al., 2017) and TVSeries (De Geest et al., 2016).

4.1.1 EXPERIMENTAL SETUP

The THUMOS14 dataset (Idrees et al., 2017) for OAD has 200 and 213 validation and test videos, respectively, with frame-level class annotations across 20 classes. As in prior OAD works, the model is trained on the validation set and evaluated on the test set. Similar to Wang et al. (2021) we use pre-extracted features from a two-stream Temporal Segment Network (TSN) (Wang et al., 2019) trained on ActivityNet v1.3 (Heilbron et al., 2015) or Kinetics-400 (Carreira & Zisserman, 2017).

For TVSeries (De Geest et al., 2016), the network learns on the train and validation sets (20 videos) and evaluates on the test set (7 videos) as in (Wang et al., 2021). RGB and Optical Flow features were extracted using an MMAAction2 (Contributors, 2020) pipeline with ActivityNet v1.3 (Heilbron et al., 2015) and Kinetics-400 (Carreira & Zisserman, 2017) pretrained TSN ResNet-50 (He et al., 2016) backbones. This is similar to the feature extraction process used by LSTR (Xu et al., 2021).

Following Wang et al. (2021), we use a batch size of 128, sequence length 64, initial learning rate 10^{-4} with a factor ten reduction each epoch, alongside weight decay, and dropout with probability 0.1. We report results using two epochs of training on a Nvidia RTX2080 Ti GPU. We track mean Average Precision (mAP) for THUMOS14 and calibrated mean Average Precision (cmAP) (De Geest et al., 2016) for TVSeries, alongside FLOPs per prediction and parameters of the OAD module (feature extraction excluded). We report the mean and standard deviation over five runs.

4.1.2 ABLATION STUDIES

Removing the Decoder As a first step to make an efficient Continual OadTR, we remove the decoder blocks used for action anticipation, which has a large impact on computational efficiency and the ease of transformation to a Continual Inference Network. The first two lines of Table 1a present the results of the removal. Contrary to the observations of Wang et al. (2021), we did not find any drop in accuracy when excluding the decoder. We do, however, gain a large reduction in FLOPs and model size; they were reduced to 58% and 30%, respectively. Given these computational improvements, we exclude the decoder in subsequent experiments.

²The feature extraction commonly used in Online Action Detection (OAD) works is in itself quite computationally costly. We consider the optimisation of the backbone as orthogonal future work and will follow the same feature extraction procedure as other OAD works at this time.

Table 1: Ablation experiments on THUMOS14 with TSN-Anet features. Best metrics are highlighted. '-' indicates that a particular feature was not used.

(a) Class token variations with OadTR. CLS pos. is the encoder block into which CLS is input.

Enc. blocks	Dec.	CLS pos.	mAP (%)	FLOPs (M)	Params (M)
3	X	1	57.0 ^{0.5}	2445.6	74.7
3	-	1	57.0 ^{0.4}	1430.6	22.2
3	-	2	57.3 ^{0.7}	1423.5	22.2
3	-	3	56.7 ^{0.6}	1417.2	22.2
3	-	-	56.8 ^{0.3}	1410.9	22.2
2	-	1	56.5 ^{0.5}	1020.7	15.9
2	-	2	56.7 ^{0.3}	1014.5	15.9
2	-	-	56.6 ^{0.3}	1008.1	15.9
1	-	1	57.1 ^{0.6}	611.7	9.6
1	-	-	56.3 ^{0.2}	605.5	9.6

(b) Positional encodings variations for CoOadTR.

Enc. blocks	Re-cycling	Learn	Pos. tokens	mAP (%)	FLOPs (M)	Params (K)
2	-	X	n	45.3 ^{0.9}	410.9	15832
2	X	X	n	56.4 ^{0.3}	410.9	15832
2	X	X	2n 1	56.0 ^{0.5}	410.9	15897
2	X	-	n	55.8 ^{1.0}	410.9	15767
2	X	-	2n 1	56.8 ^{0.4}	410.9	15767
1	-	X	n	44.0 ^{0.8}	9.6	9535
1	X	X	n	55.6 ^{0.3}	9.6	9535
1	X	X	2n 1	55.6 ^{0.3}	9.6	9599
1	X	-	n	54.4 ^{1.8}	9.6	9469
1	X	-	2n 1	56.1 ^{0.7}	9.6	9469

(Re)moving the Class token Class tokens should not be input naively to the first Transformer Encoder layer of a CIN (see Section 3.8). Accordingly, we ablate its use and position. In cases where it is removed, we predict on the token corresponding to the last input token. The results of varying CLS pos are noted in Table 1a. For the one-block architecture, the removal came with noticeable drop in mAP, while the two-block architecture saw small improvements when removing or introducing the class token later. For the three block model, the use of class tokens in block two achieved the highest mAP. Though it is commonly accepted, that class tokens should be introduced alongside other inputs in the first block, our results indicate that they can accumulate sufficient information with only one or two blocks, and that later stage introduction may work better in some applications. In general, the achieved mAP when varying CLS pos. and number of blocks are very similar to one another, while (re)moving the class token and reducing the block size both reduce computational complexity. This encourages the use of shallow Transformer Encoders over deeper ones as well as the removal of class tokens, as we do in the following experiments.

Positional Encodings We can transfer parameters from the simplified one- and two block OadTR to the corresponding Continual architecture CoOadTR. Here, the one block version CoOadTR-b1) uses CoSMHA, and the two block model CoOadTR-b2) uses CoRMHA in the first block and Single-output MHA in the second. However, a regular positional encoding is not suited for continual inference (see Section 3.7). We evaluate the performance of using non-continual encodings for continual inference, as well as of our proposed Recycling Positional Encodings with fixed or learned parameters. In addition, we explore the impact of extending the number of tokens from n to $2n + 1$ to avoid positional ambiguity. As seen in Table 1b), non-continual encoding used in the continual setting result in severe mAP drop. Recycling Positional Encodings alleviate this. Comparing learned and fixed encodings, we find the learned encodings to work better when the number of encoding tokens corresponds to the sequence length and the fixed encoding to work best when positional ambiguity is alleviated by extending the number of tokens to $2n + 1$. Fixed encoding with $2n + 1$ tokens works best overall and is employed in subsequent experiments unless stated otherwise. There is no difference in FLOPs for either strategy, and the difference in parameter count is negligible.

4.1.3 COMPARISON WITH PRIOR WORKS

We evaluate the CoOadTR architectures on THUMOS14 and TVSeries with two sets of features as described in Section 4.1.1. Since no prior OAD works have reported complexity metrics, we measured the FLOPs for TRN (Xu et al., 2019) based on the publicly available source code to serve as a point of reference. The results of this benchmark are presented in Table 2 and Fig. 2. OadTR and our simplified (continual) one-block (b1) and two-block (b2) versions without decoder and class tokens generally achieve competitive precision in comparison with prior works, surpassing all but OadTR and LSTR. On THUMOS14, our reproduced OadTR results are slightly lower than originally reported (Wang et al., 2021), whereas achieved TVSeries results are higher. The CoOadTR-b#

³The reported 58.3% on THUMOS14 could not be reproduced using their publicly available code.

⁴We attribute our higher mcAP to differences in the feature extraction pipeline.

Table 2: Online Action Detection results. FLOPs per prediction are noted for inference on THUMOS14. The best and next-best metrics are highlighted.

Model	Feat.	THUMOS14 mAP (%)	TVSeries mcAP (%)	FLOPs (M)
RED (Gao et al., 2017)		45.3	79.2	-
TRN (Xu et al., 2019)		47.2	83.7	1387.5
FATS (Kim et al., 2021)		51.6	81.7	-
IDN (Eun et al., 2020)		50.0	84.7	-
TFN (Eun et al., 2021)		55.7	85.0	-
LSTR (Xu et al., 2021)		65.3	88.1	-
OadTR (Wang et al., 2021)	A.Net	58.3	85.4	2445.6
OadTR [†]		57.0 ^{0.5}	88.6 ^{0.1}	2445.6
OadTR-b2		56.6 ^{0.3}	88.3 ^{0.2}	1008.1
OadTR-b [‡]		56.3 ^{0.2}	88.1 ^{0.1}	605.5
CoOadTR-b2 (ours)		56.8 ^{0.4}	87.7 ^{0.6}	410.9
CoOadTR-b1 (ours)		56.1 ^{0.7}	87.6 ^{0.7}	9.6
TRN (Xu et al., 2019)		62.1	86.2	1462.0
FATS (Kim et al., 2021)		59.0	84.6	-
IDN (Eun et al., 2020)		60.3	86.1	-
PKD (Zhao et al., 2020)		64.5	86.4	-
LSTR (Xu et al., 2021)		69.5	89.1	-
OadTR (Wang et al., 2021)	Kin.	65.2	87.2	2513.5
OadTR [†]		64.2 ^{0.3}	88.6 ^{0.1}	2513.5
OadTR-b2		64.5 ^{0.5}	88.3 ^{0.2}	1075.7
OadTR-b [‡]		63.9 ^{0.5}	88.1 ^{0.1}	673.0
CoOadTR-b2 (ours)		64.4 ^{0.1}	87.6 ^{0.7}	411.9
CoOadTR-b1 (ours)		64.2 ^{0.4}	87.7 ^{0.4}	10.6

Figure 2: Visual comparison of OAD methods on THUMOS14 and TVSeries for backbones trained on ActivityNet 1.3 and Kinetics-400.

[†]Using of cial source code or modi cations there-off.

architecture largely retain precision and allow significantly reduced FLOPs per prediction. Our proposed continual variants CoOadTR-b1 and CoOadTR-b2 reduce FLOPs by 95% and 6.1%, respectively, compared to OadTR, while either achieving the same performance or conceding no more than one percentage point. On average, continual and non-continual OadTR-b# models achieve similar mAP on THUMOS14, while OadTR-b# have slightly higher mcAP on TVSeries. We attribute these discrepancies to differences in positional encoding. All in all, CoOadTR-b# models provide far-superior computational efficiency to prior works, achieving state-of-the-art performance/efficiency trade-offs by a large margin.

4.1.4 AUDIO-VISUAL ONLINE ACTION DETECTION

To showcase the validity of our method in audio-visual settings as well, we explore the addition of audio-features to the Online Action Detection task on THUMOS14. As described in Section 4.2, audio-features are extracted using Mel spectrograms and an AudioSet pre-trained VGGish network (Hershey et al., 2017) (output of the penultimate layer) on 1.0 second windows with a step size of 0.2 seconds to match the 5.0 FPS sampling rate of the video features.

The audio-features by themselves do not provide enough signal to reach good Online Action Detection performance (yielding only 6.7% mAP with an OadTR network). When concatenated with RGB and Flow they do provide a modest improvement as seen in Table 3. On average, this amounts to +0.6% mAP when combined with ActivityNet features and +0.5% mAP when used with Kinetics-400 features with shallower models enjoying the largest improvements.

4.2 AUDIO CLASSIFICATION

4.2.1 BACKGROUND

Audio Classification is the categorisation of audio waveforms. Though waveform sequences can be used directly (Lee et al., 2017), it is common to first convert them to spectrograms. Mel spectrograms are obtained by a nonlinear transformation of a frequency scale (Stevens et al., 1937), which is designed based on empirical knowledge about the human auditory system (Choi et al., 2016). By employing spectrograms, audio classification can be approached in the same way as image classification (Palanisamy et al., 2020).

Table 3: Audio-Visual result, THUMOS14.

Model	Feat.	mAP (%)	FLOPs (M)
OadTR		57.6 _{0.6}	2714.9
OadTR-b2	A.Net	57.5 _{0.5}	1277.0
OadTR-b1	+	57.4 _{0.4}	874.1
CoOadTR-b2	AudioSet	56.5 _{0.1}	415.0
CoOadTR-b1		56.8 _{0.5}	13.8
OadTR		64.4 _{0.4}	2781.9
OadTR-b2	Kin.	65.0 _{0.4}	1344.1
OadTR-b1	+	64.5 _{0.4}	941.2
CoOadTR-b2	AudioSet	64.7 _{0.8}	416.0
CoOadTR-b1		64.8 _{0.3}	14.8

Table 4: Audio Classification results for GTZAN.

Method	Pos. Enc.	Acc. (%)	FLOPs (M)	Par. (K)
Maj. Voting	-	92.0	-	0
Trans-b2	learned	95.0 _{0.6}	47.4	509
Trans-b1	learned	93.8 _{0.8}	15.2	286
CoTrans-b2	xed	94.4 _{1.0}	27.0	509
CoTrans-b1	learned	93.2 ₁	0.3	286

4.2.2 EXPERIMENTS

We conduct experiments on the Music Genre Classification dataset GTZAN (Tzanetakis & Cook, 2002). It consists of 100 30-second clips for each of ten music genres. Each audio clip is sampled at 22,050 Hz. Since there are no predefined splits for GTZAN, we randomly select 10% of the data for validation and 10% for testing. The input is transformed to a temporal sequence by sliding a one-second window over each 30-second clip with a slide step size of 250ms, leading to 120 one-second clips. These are subsequently converted to Mel spectrograms. We then re-tune a VGGish network, pre-trained on AudioSet (Hershey et al., 2017) and use the penultimate layer for feature extraction. A batch size of 64 and the Adam optimizer (Kingma & Ba, 2015) are used with an initial learning rate of 10^{-4} . The learning rate is reduced by a factor of 0.5 on plateau with a tolerance of two epochs, and an early stopping mechanism, where a maximum of 100 epochs are allowed. The VGGish base-network attains an accuracy of 86.1% on the dataset of one-second clips with 72.1M parameters and 864.7M FLOPs. Subsequently, the audio features are passed to a (Continual) Transformer Encoder which has 16 attention heads, an embedding dimension of 192 and an MLP dimension of 384. The Transformer Encoder is trained on the whole temporal sequence using a batch size of 32 and the AdamW optimizer (Loshchilov & Hutter, 2019) with a learning rate of 10^{-4} and a weight decay of 10^{-4} for 50 epochs. Since the Transformer Encoder is trained on entire 30-second clips, there are less data points available for this training. Accordingly, the size of the validation set is increased to 18%. All audio classification training procedures were carried out on a single Nvidia RTX 2080 Ti GPU. Table 4 presents the accuracy and efficiency of regular and Continual Transformers during online inference. As a baseline, we also include the result of majority voting among the clips to classify the entire sequence. The Continual Transformers obtain similar accuracy as regular Transformers while consuming 76% less FLOPs when using two blocks and 51% less FLOPs when using one Transformer Encoder block.

5 CONCLUSION

In this work, we presented Continual Transformers, a redundancy-free reformulation of Transformers tailored for online inference. Central to the Continual Transformer are the Continual Retroactive and Single-Output Attention operations, which produce outputs identical to the original Scaled Dot-Product Attention for continual input sequences, while greatly reducing the time and memory complexity per prediction. The applicability of Continual Transformer architectures was experimentally validated in Online Action Detection and Online Audio Classification settings, observing upwards of multiple orders of magnitude reduction in time complexity for lightweight architectures at modest accuracy concessions. Continual Transformers constitute an algorithmic innovation, which could make possible hitherto unseen precision, speed, and power efficiency in online inference use-cases. With applications spanning enhanced perception and reactivity of robots and autonomous vehicles, weather forecasting, price prediction and surveillance, we hope it will be used for the common good.

ACKNOWLEDGMENTS

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 871449 (OpenDR).

REFERENCES

- Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lucic, and Cordelia Schmid. Vivit: A video vision transformer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- Arian Bakhtiarnia, Qi Zhang, and Alexandros Iosifidis. Single-layer vision transformers for more accurate early exits with less overhead. *arXiv preprint arXiv:2105.09121*, 2021a.
- Arian Bakhtiarnia, Qi Zhang, and Alexandros Iosifidis. Multi-exit vision transformer for dynamic inference. *British Machine Vision Conference (BMVC)*, 2021b.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv:2004.05150*, 2020.
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (eds.), European Conference on Computer Vision (ECCV)*, pp. 213–229, 2020.
- J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4724–4733, 2017.
- Keunwoo Choi, George Fazekas, and Mark Sandler. Automatic tagging using deep convolutional neural networks. *International Society of Music Information Retrieval Conference (ISMIR)*, 2016.
- Krzysztof Marcin Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser, David Benjamin Belanger, Lucy J Colwell, and Adrian Weller. Rethinking attention with performers. In *International Conference on Learning Representations (ICLR)*, 2021.
- Kateryna Chumachenko, Alexandros Iosifidis, and Moncef Gabbouj. Self-attention fusion for audiovisual emotion recognition with incomplete data. *arXiv preprint arXiv:2201.11095*, 2022.
- MMAAction2 Contributors. Openmmlab's next generation video understanding toolbox and benchmark. <https://github.com/open-mmlab/mmaaction2>. Apache 2.0 License., 2020.
- Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. Smart: Training shallow memory-aware transformers for robotic explainability. *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1128–1134, 2020. doi: 10.1109/ICRA40945.2020.9196653.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. Transformer-XL: Attentive language models beyond a fixed-length context. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 2978–2988. Association for Computational Linguistics, July 2019. doi: 10.18653/v1/P19-1285.
- Roeland De Geest, Efstratios Gavves, Amir Ghodrati, Zhenyang Li, Cees Snoek, and Tinne Tuytelaars. Online action detection. *European Conference on Computer Vision (ECCV)*, pp. 269–284, 2016.
- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Lukasz Kaiser. Universal transformers. In *International Conference on Learning Representations (ICLR)*, 2019. URL: <https://openreview.net/forum?id=HyzdRiR9Y7>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186. Association for Computational Linguistics, June 2019. doi: 10.18653/v1/N19-1423.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*, 2021.

- H. Eun, Jinyoung Moon, Jongyoul Park, Chanho Jung, and Changick Kim. Learning to discriminate information for online action detection. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 806–815, 2020.
- Hyunjun Eun, Jinyoung Moon, Jongyoul Park, Chanho Jung, and Changick Kim. Temporal ltering networks for online action detection. *Pattern Recognition* 11:107695, 2021.
- Jiyang Gao, Zhenheng Yang, and Ram Nevatia. RED: reinforced encoder-decoder networks for action anticipation. In *British Machine Vision Conference (BMV)*, 2017.
- Yuan Gong, Yu-An Chung, and James Glass. AST: Audio Spectrogram Transformer. *Interspeech 2021*, pp. 571–575, 2021. doi: 10.21437/Interspeech.2021-698.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. doi: 10.1109/cvpr.2016.90.
- Lukas Hedegaard and Alexandros Iosiadis. Continual 3d convolutional neural networks for real-time processing of videos. *European Conference on Computer Vision (ECCV)*, 2021. Apache 2.0 Licence.
- Lukas Hedegaard, Negar Heidari, and Alexandros Iosiadis. Online skeleton-based action recognition with continual spatio-temporal graph convolutional networks. preprint, arXiv: 2203.11009 2022. Apache 2.0 Licence.
- Negar Heidari and Alexandros Iosiadis. Progressive spatio-temporal graph convolutional network for skeleton-based human action recognition. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* 3220–3224, 2021.
- Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 961–970, 2015. doi: 10.1109/CVPR.2015.7298698.
- Shawn Hershey, Sourish Chaudhuri, Daniel P. W. Ellis, Jort F. Gemmeke, Aren Jansen, R. Channing Moore, Manoj Plakal, Devin Platt, Rif A. Saurous, Bryan Seybold, Malcolm Slaney, Ron J. Weiss, and Kevin Wilson. CNN architectures for large-scale audio classification. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2017. doi: 10.1109/icassp.2017.7952132.
- Cheng-Zhi Anna Huang, Ashish Vaswani, Jakob Uszkoreit, Ian Simon, Curtis Hawthorne, Noam M. Shazeer, Andrew M. Dai, Matthew D. Hoffman, Monica Dinulescu, and Douglas Eck. Music transformer: Generating music with long-term structure. *International Conference on Learning Representations (ICLR)*, 2019.
- Haroon Idrees, Amir R. Zamir, Yu-Gang Jiang, Alex Gorban, Ivan Laptev, Rahul Sukthankar, and Mubarak Shah. The thumos challenge on action recognition for videos “in the wild”. *Computer Vision and Image Understanding* 155:1–23, 2017. ISSN 1077-3142.
- Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35(1):221–231, 2013. doi: 10.1109/TPAMI.2012.59.
- Young Hwi Kim, Seonghyeon Nam, and Seon Joo Kim. Temporally smooth online action detection using cycle-consistent future anticipation. *Pattern Recognition* 116:107954, 2021. ISSN 0031-3203.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), *International Conference on Learning Representations (ICLR)*, 2015.
- Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. In *International Conference on Learning Representations (ICLR)*, 2020.

- Shun Kiyono, Sosuke Kobayashi, Jun Suzuki, and Kentaro Inui. SHAPE: Shifted absolute position embedding for transformers. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 3309–3321, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.266.
- Jongpil Lee, Taejun Kim, Jiyoung Park, and Juhan Nam. Raw waveform-based audio classification using sample-level cnn architecture. *ALP, Machine Learning for Audio Signal Processing Workshop (ML4Audio)* 2017.
- Tatiana Likhomanenko, Qiantong Xu, Gabriel Synnaeve, Ronan Collobert, and Alex Rogozhnikov. Cape: Encoding relative positions with continuous augmented positional embeddings. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds), *Advances in Neural Information Processing Systems*, volume 34, pp. 16079–16092. Curran Associates, Inc., 2021.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *International Conference on Learning Representations (ICLR)* 2019.
- Yining Ma, Jingwen Li, Zhiguang Cao, Wen Song, Le Zhang, Zhenghua Chen, and Jing Tang. Learning to iteratively solve routing problems with dual-aspect collaborative transformer. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds), *Advances in Neural Information Processing Systems* 2021.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. fairseq: A fast, extensible toolkit for sequence modeling. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pp. 48–53, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-4009.
- Kamalesh Palanisamy, Dipika Singhania, and Angela Yao. Rethinking cnn models for audio classification. *arXiv:2007.11154* 2020.
- Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In Jennifer Dy and Andreas Krause (eds), *International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4055–4064. PMLR, 7 2018.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. Self-attention with relative position representations. In *North American Chapter of the Association for Computational Linguistics (NAACL)* 2018.
- Zheng Shou, Dongang Wang, and Shih-Fu Chang. Temporal action localization in untrimmed videos via multi-stage cnns. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1049–1058, 2016. doi: 10.1109/CVPR.2016.119.
- Zheng Shou, Jonathan Chan, Alireza Zareian, Kazuyuki Miyazawa, and Shih-Fu Chang. Cdc: Convolutional-de-convolutional networks for precise temporal action localization in untrimmed videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1417–1426, 2017. doi: 10.1109/CVPR.2017.155.
- S. S. Stevens, J. Volkman, and E. B. Newman. A scale for the measurement of the psychological magnitude pitch. *The Journal of the Acoustical Society of America*, 63(1):185–190, 1937. doi: 10.1121/1.1915893.
- Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey. *arXiv:2009.06732* 2020.
- Hugo Touvron, Matthieu Cord, Alaeldin El-Nouby, Piotr Bojanowski, Armand Joulin, Gabriel Synnaeve, and Hervé Jégou. Augmenting convolutional networks with attention-based aggregation. preprint, *arXiv:2112.13692* abs/2112.13692, 2021.
- G. Tzanetakis and P. Cook. Musical genre classification of audio signals. *IEEE Transactions on Speech and Audio Processing*, 10(5):293–302, July 2002. doi: 10.1109/tsa.2002.800560.

- George Tzanetakis, Georg Essl, and Perry Cook. Automatic musical genre classification of audio signals, 2001. URL <http://ismir2001.ismir.net/pdf/tzanetakis.pdf>.
- Gül Varol, Ivan Laptev, and Cordelia Schmid. Long-term temporal convolutions for action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (TPAMI), 40(6):1510–1517, 2018. doi: 10.1109/TPAMI.2017.2712608.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, pp. 5998–6008, 2017.
- Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks for action recognition in video. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (TPAMI), 41(11):2740–2755, 2019. doi: 10.1109/TPAMI.2018.2868668.
- Sinong Wang, Belinda Z. Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. [arXiv:2006.04768](https://arxiv.org/abs/2006.04768)2020.
- Xiang Wang, Shiwei Zhang, Zhiwu Qing, Yuanjie Shao, Zhengrong Zuo, Changxin Gao, and Nong Sang. Oadtr: Online action detection with transformers. *International Conference on Computer Vision (ICCV) 2021*. URL <https://github.com/wangxiang1230/OadTR>. MIT License.
- Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philip H. S. Tang, and Ross Girshick. Long-term feature banks for detailed video understanding. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 284–293, 2019. doi: 10.1109/CVPR.2019.00037.
- Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. Nystromformer: A nystrom-based algorithm for approximating self-attention. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
- Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. In *IEEE International Conference on Computer Vision (ICCV)*, pp. 5794–5803, 2017. doi: 10.1109/ICCV.2017.617.
- Mingze Xu, Mingfei Gao, Yi-Ting Chen, Larry Davis, and David Crandall. Temporal recurrent networks for online action detection. *IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 5531–5540, 2019. doi: 10.1109/ICCV.2019.00563. URL <https://github.com/xumingze0308/TRN.pytorch>. MIT License.
- Mingze Xu, Yuanjun Xiong, Hao Chen, Xinyu Li, Wei Xia, Zhuowen Tu, and Stefano Soatto. Long short-term transformer for online action detection. *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. *AAAI Conference on Artificial Intelligence*, pp. 7444–7452, 2018.
- Peisen Zhao, Jiajie Wang, Lingxi Xie, Ya Zhang, Yanfeng Wang, and Qi Tian. Privileged knowledge distillation for online action detection. [preprint, arXiv:2011.09158](https://arxiv.org/abs/2011.09158), 2020.

A APPENDIX

A.1 SCALING PROPERTIES OF CONTINUAL AND REGULAR MULTI-HEAD ATTENTION

A detailed account for the scaling point operations involved in computing Regular-, Continual Retroactive-, and Single-output Scaled Product Attentions is given in Tables 5, 6, and 7.

Fig. 3 illustrates the scaling of FLOPs and memory footprint with increasing sequence length and embedding dimension. Here, the Continual Retroactive and Single-Output SDAs spend significantly less FLOPs than the Regular SDA, which scales as $O(d^2)$ as opposed to $O(d)$ (and) the continual variants. The Continual Single-Output SDA reduces memory footprint for all value combinations, and the Continual Retroactive SDA does so when d .

Table 5: Floating Point Operations for the Scaled Dot-Product Attention in Eq. (1). $\mathbf{Q}^T \mathbf{K}$ can be efficiently computed as element-wise multiplication with \mathbf{W} .

	Mul.	Add	Exp
Eq. (1.1)	$n^2d + nd$	$nd(n-1)$	0
Eq. (1.2)	$n^2d + nd$	$n^2(d-1)$	n^2
Eq. (1.3)	0	$n(n-1)$	0

Table 6: Floating Point Operations for the Continual Retroactive Dot-Product Attention in Eqs. (2) to (6). The outputs of the exponentials in Eq. (2) and Eq. (3) can be reused in Eq. (4) and Eq. (5) respectively, and are omitted in the count.

	Mul.	Add	Exp
Eq. (2)	$2(n-1)d$	$2(n-2)d + 2(n-1)$	$2(n-1)$
Eq. (3)	$nd + n + d$	$nd + (n-1) + d$	n
Eq. (4)	$2(n-1)d$	$2(n-1)d$	0
Eq. (5)	nd	$(n-1)d$	0
Eq. (6)	$nd + n$	0	0

Table 7: Floating Point Operations for the Continual Single-Output SDA in Eq. (7).

	Mul.	Add	Exp
Eq. (7.1)	$nd + d$	$(n-1)d + n - 1$	0
Eq. (7.2)	$nd + d$	$n(d-1)$	n

(a) $d = 256$

(b) $n = 256$

Figure 3: FLOPs/step and memory footprint for Regular, Continual Retroactive, and Continual Single-Output Scaled Dot-Product Attention at varying sequence length and embedding dimension d . Column (a) has n fixed to 256; Column (b) has d fixed to 256.

A.2 SUPPLEMENTAL VISUALISATIONS

For the visually inclined, we supply a complementary graphical depictions of the Continual Retroactive SDA corresponding to Eqs. (2) to (6) in Fig. 4 and the Single-Output SDA in Eq. (7) in Fig. 5.

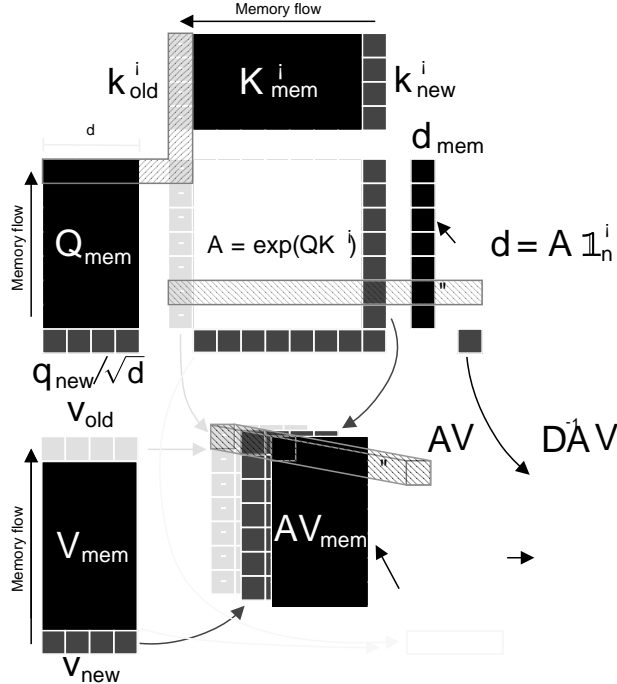


Figure 4: **Continual Retroactive Dot-Product Attention.** The query (Q), key (K), and value (V) matrices are aggregated over time by caching the step vectors q_{new} , k_{new} , and v_{new} in each their FIFO queue (denoted by $_{mem}$). During each step, only the entries of A associated with q_{new} , k_{new} and the oldest K step, k_{old} are computed. The diagonal entries of the row-normalisation matrix D as well as the AV can be updated retroactively by subtracting features corresponding to k_{old} and adding features related to k_{new} to the cached outputs of the previous step, D_{mem} and AV_{mem} , respectively.




Figure 5: **Continual Single-Output Dot-Product Attention.** The key (K) and value (V) matrices are aggregated over time by caching the step vectors k_{new} and v_{new} in a FIFO queue. During each step, only the attention output associated with q is computed.

A schematic illustration of the Audio Classification experiments architecture is depicted in Fig. 6.

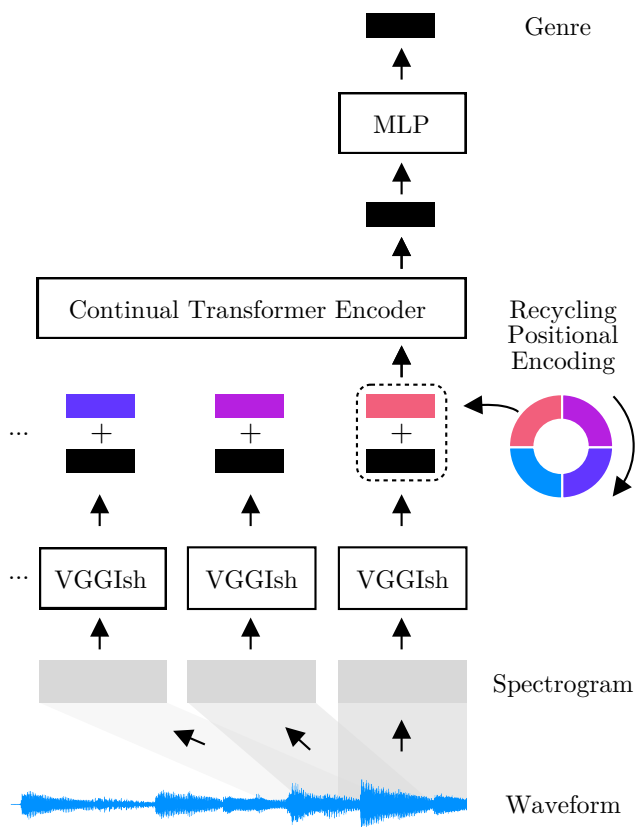


Figure 6: **Audio Classification Architecture.**