

Movie Trailer Genre Classification Using Multimodal Pretrained Features

Serkan Sulun^{a,1,*}, Paula Viana^{a,b}, Matthew E. P. Davies^c

^a*Institute for Systems and Computer Engineering, Technology and Science (INESC TEC), Porto, 4200-465, Portugal*

^b*ISEP, Polytechnic of Porto, School of Engineering, Porto, 4200-072, Portugal*

^c*SiriusXM, New York, 10020, NY, USA*

Abstract

We introduce a novel method for movie genre classification, capitalizing on a diverse set of readily accessible pretrained models. These models extract high-level features related to visual scenery, objects, characters, text, speech, music, and audio effects. To intelligently fuse these pretrained features, we train small classifier models with low time and memory requirements. Employing the transformer model, our approach utilizes all video and audio frames of movie trailers without performing any temporal pooling, efficiently exploiting the correspondence between all elements, as opposed to the fixed and low number of frames typically used by traditional methods. Our approach fuses features originating from different tasks and modalities, with different dimensionalities, different temporal lengths, and complex dependencies as opposed to current approaches. Our method outperforms state-of-the-art movie genre classification models in terms of precision, recall, and mean average precision (mAP). To foster future research, we make the pretrained features for the entire MovieNet dataset, along with our genre classification code and the trained models, publicly available.

*Corresponding author.

Email addresses: serkan.sulun@inesctec.pt (Serkan Sulun), pmv@isep.ipp.pt, paula.viana@inesctec.pt (Paula Viana), matthew.davies@siriusxm.com (Matthew E. P. Davies)

¹Serkan Sulun received the support of fellowship FCT - Fundação para a Ciência e a Tecnologia with the fellowship code 2022.09594.BD.

²This work has been funded by National Funds through the Portuguese funding agency, FCT - Fundação para a Ciência e a Tecnologia, within project LA/P/0063/2020.

Keywords: video classification, multimodal features, cinematic trailer classification, transformers

1. Introduction

The classification of movies into genres serves as a vital navigational tool in the immense catalog of cinematic content, allowing audiences to locate films aligned with their tastes and interests. It also provides a framework for scholars to explore thematic patterns, the evolution of narratives, and cultural shifts reflected in cinematic art. Moreover, for industry stakeholders, accurate genre classification guides key decisions around marketing and distribution strategies. The cinematic world contains an extraordinary blend of visual and auditory stimuli, each film constituting a unique assortment of aesthetics, themes, and narrative structures. However, this incredible diversity makes it challenging to accurately categorize films by genre.

Movie trailers serve as condensed representations of the full-length film, encompassing an array of vital elements within a brief span. Trailers provide a comprehensive overview of a film’s narrative and genre by incorporating pivotal scenes, characters, objects, and vital auditory cues such as dialogue, music, and informative text elements. The availability of movie trailers on online video platforms enables researchers to link movie metadata such as genre to the URL of the movie trailers, creating large datasets for movie trailer genre classification such as MovieNet (Huang et al., 2020) and Movie-scope (Cascante-Bonilla et al., 2019).

In trailer genre classification, as well as in almost all video classification tasks, deep neural networks (DNNs) constitute the state-of-the-art. The majority of these models use raw data such as pixels and audio waveforms as their inputs. In contrast to older methods that use hand-crafted features, DNNs are expected to extract useful high-level features from raw data in an implicit way within the model layers. These models often contain millions of parameters, requiring significant amounts of time, memory, and training data. Specifically in video classification, the uncompressed size of the raw video frame sequence increases the time and memory requirements even further. Previous works deal with this problem by processing only a fixed and small number of input frames or clips (Huang et al., 2020; Wehrmann & Barros, 2017). However, this approach inherently leads to information loss, potentially compromising classification accuracy. Other works tackle this

issue by using lower-dimensional, non-visual, and non-auditory inputs such as movie posters, plots, and metadata, for movie genre classification (Rodríguez Bribiesca et al., 2021; Ak et al., 2023; Miyazawa et al., 2022). These approaches are dataset-specific and rely on auxiliary information, failing to classify solely based on video elements, namely on raw video pixels and audio waveforms.

The computational complexity of processing raw videos increases further with the adaptation of larger models—a trend that continues to expand over the years (Tan & Le, 2019; Bernstein et al., 2021). When the training data is limited, using larger input dimensionality and larger models also increases the chance of overfitting (Defernez & Kemsley, 1999). The common and very similar ways to deal with these issues are transfer learning, fine-tuning, and using pretrained features (Tan et al., 2018; Niu et al., 2020). These terminologies can be distinguished as follows: Transfer learning involves applying a model trained on one task to a different task. Fine-tuning entails further training the transferred model for the new task. Using pretrained features involves performing inference on the data using a pretrained model, and then feeding the resulting outputs or activations (features) into a new model that is trained. These methods are frequently employed in video processing tasks such as human activity recognition, video summarization, and video recommendation (Ray & Kolekar, 2024; Khan et al., 2024; Almeida et al., 2022).

Using pretrained features brings multiple advantages. Firstly, since the pretrained model is used in inference mode, its weights remain frozen, therefore reducing the time and memory complexity by requiring only a forward pass, without an additional backward pass. Secondly, it considerably reduces the size of the input fed to the subsequent model that is being trained. As an example, a very small video frame contains $224 \times 224 \times 3 = 163,968$ values, while the state-of-the-art image analysis model CLIP (Contrastive Language-Image Pre-Training) represents an image using only 512 values (Radford et al., 2021). This size reduction becomes more important as we use trailer videos that average around 2 minutes in length, containing a long sequence of frames. Thirdly, assuming that the pretrained features of multimedia analysis models present valuable and relevant information for the task of video classification, their use removes the need to train new models to extract similar features implicitly from raw data. This reduces the size of the subsequent model that needs to be trained. Finally, since it reduces the input size and the number of weights in training, it also reduces the chance of overfitting (Defernez & Kemsley, 1999).

In this work, we focus on classifying movie trailers into genres using pre-trained features that we extract entirely from video and audio. To obtain the relevant pretrained features, we exploit the availability of open-source pretrained multimedia analysis models. Specifically, we employ image analysis, optical character recognition, automatic speech recognition, audio tagging, and music classification models to incorporate information about visual scenery, objects, text, speech, music, and audio effects. These pretrained models are used in inference mode, meaning that they are not fine-tuned, resulting in much lower time and memory requirements. Another important strength of our methods is the use of the transformer model to fuse and process the pretrained features and predict the genre of the input video (Vaswani et al., 2017). The transformer model is designed to process long input sequences efficiently while handling long-term dependencies. In our case, we feed the pretrained features that stem from many video frames and audio chunks into the transformer. The small size of the pretrained features and the efficiency of the transformer model enables us to use entire videos while maximizing the available information, and not use a fixed and small number of input frames or audio chunks. Furthermore, since the pretrained models extract meaningful features from raw data, using a shallow transformer with one or two layers as the classifier becomes sufficient. We propose and compare multiple architectures to handle the non-trivial task of merging and processing multimodal pretrained feature sequences with complex correspondences and different temporal lengths and dimensionalities.

We implement additional models to reflect some of the works in the literature and to compare against our methods. As done by Huang et al. (2020), we use a temporal averaging module followed by a multi-layer perceptron (MLP) instead of the transformer, to average the pretrained features along the time dimension, and then classify them. Our empirical results show that the transformer outperforms this approach, especially when more input frames are used. We secondly implement a baseline model that works on raw video and audio, to mirror some of the recent models used in trailer genre classification (Rodríguez Bribiesca et al., 2021). While we empirically show that this approach performs worse than our proposed models, our preliminary experiments also show that it is prone to overfitting when trained end-to-end. Lastly, we compare our results with other models addressing trailer genre classification on the MovieNet dataset, demonstrating a notable performance improvement.

We summarize the shortcomings of the previous works in trailer genre

classification and how we address them as the following:

- The complexity introduced by the large input space due to the use of raw video and audio: By using pretrained features instead of raw pixels and waveforms, we significantly reduce input complexity, dimensionality, and effectively, the chance of overfitting.
- The information loss caused by using a low and fixed number of input frames: Alongside using inputs with lower dimensionality, we utilize the transformer architecture that can process the sequences from the videos in their entirety while handling long-term correspondences.
- Inability to classify solely using videos due to dependence on dataset-specific content such as poster, plot, or metadata: The pretrained features that we used are obtained entirely from video and audio. As a result, our models can theoretically be used for any video classification task.

In summary, our contributions are as follows:

- We propose new strategies for leveraging a variety of pretrained features as inputs and processing them using shallow neural networks for classification capable of handling different input sizes in both channel and temporal dimensions.
- By utilizing the transformer model, we leverage all video keyframes and the entire audio, effectively handling videos of varying lengths without the need for picking out a constant number of frames or average pooling. We also quantitatively demonstrate that the classification performance improves steadily with the inclusion of more frames.
- We improve the state-of-the-art for genre classification significantly, on the largest movie genre classification dataset, namely MovieNet.
- We make available the pretrained features for the MovieNet dataset, as well as our code and trained models.³

³<https://github.com/serkansulun/trailer-genre-classification>
<https://zenodo.org/records/13909366>

2. Related work

In this section, we first explain the cinematic datasets for movie genre classification and further motivate why we have chosen to work with the MovieNet dataset (Huang et al., 2020), and then we discuss the models specifically focused on movie genre classification.

2.1. Cinematic datasets

MovieLens is one of the earliest cinematic datasets, containing movie ratings and metadata like genre and title (Harper & Konstan, 2016). While the latest version includes 86k movies, it does not include direct links to the movie trailers. The trailers can only be obtained through crawling the Internet Movie Database (IMDb) website⁴, which conflicts with their terms of use. The MM-IMDb (Multimodal IMDb) dataset merges movie posters with the metadata obtained from MovieLens, resulting in 26k movies with posters, plots, genres, and other metadata, while lacking trailers or any types of videos (Ovalle et al., 2017). The LMTD (Large Movie Trailer Dataset) is a collection of features and metadata from 3500 of movie trailers, although the project is currently discontinued and the data is unavailable (Simoes et al., 2016). MovieScope is a comprehensive dataset for multi-modal movie analysis, including data like movie trailers, posters, plot synopses, user reviews, and visual-auditory features, belonging to 5k distinct movies (Cascante-Bonilla et al., 2019). The MMTF-14K (Multimodal Movie Trailer Features dataset) provides multimodal features extracted from 14k movie trailers and their metadata such as user reviews and genre (Deldjoo et al., 2018). The Condensed Movies dataset includes full individual scenes, rather than trailers, alongside the plot and the characters, from 4k movies (Bain et al., 2020).

MovieNet is a large-scale, holistic dataset providing movie, trailer, poster, subtitle, plot, tags, and metadata including genre (Huang et al., 2020). While it contains metadata belonging to 375k movies, trailers are available via YouTube links for 33k movies, making MovieNet the largest cinematic dataset in terms of both metadata and trailers. The work also includes the genre classification performance of several state-of-the-art video classification models, providing strong baselines to compare against our work. Overall, because MovieNet is the largest movie trailer dataset, has direct links to trail-

⁴<https://www.imdb.com>

ers, and has quantitative genre classification results to serve as a baseline, we trained and evaluated our models on this particular dataset.

2.2. *Movie genre classification*

One of the earliest works on movie genre classification obtains keyframes using scene detectors, and extracts hand-crafted visual features such as roughness, ruggedness, and openness, on a privately collected dataset (Zhou et al., 2010). Genre classification is achieved by comparing the distance between the feature vectors from the training and testing sets.

One of the first works that use neural networks for movie genre classification also utilizes visual pretrained features (Wehrmann & Barros, 2017), using the LMTD dataset (Simoes et al., 2016). They use pretrained features of classification models that were trained on ImageNet (Deng et al., 2009) and Places365 (Zhou et al., 2018b) datasets, and audio spectrograms. The features from individual frames are fused using a convolution-through-time module, which can be thought of as a standard convolutional neural network (CNN) where the kernel length is equal to the input feature length of each keyframe and the kernel traverses the features belonging to subsequent frames, along the temporal dimension (Wehrmann & Barros, 2017). The kernel length along the temporal dimension is 3, meaning that the model can exploit the correspondence between only 3 frames. Moreover, since the input and output of the CNN have varying numbers of frames, they further take the maximum along the temporal dimension to have a fixed-size vector to feed into the final classification layer. This inevitably leads to a loss of information. Another work exploits the correspondence between facial emotions and cinematic genre, by first extracting human faces from the trailer videos, then classifying their emotions, and finally mapping the emotions to the cinematic genres (Yadav & Vishwakarma, 2020).

The work presenting the MovieNet dataset also introduces a model for movie genre classification, alongside the results from other state-of-the-art video classification models (Huang et al., 2020). While the model they introduce does surpass the performance of the state-of-the-art models, it only uses 8 clips, each with 3 frames from the entire trailer. During the inference phase, predictions are made for each individual clip, and these are subsequently averaged to generate the final prediction. This approach, however, fails to account for the long-term correspondence that exists within the trailers. Finally, the MovieCLIP model first trains a scene classification model

and then uses its final activations to feed into an additional genre classification model, specifically working with the Moviescope dataset (Bose et al., 2023).

Some recent works have used transformers for movie genre classification. Rodríguez Bribiesca et al. (2021) used transformers to individually process raw video frames and raw audio, and later fused the resulting activations with metadata information and poster. Miyazawa et al. (2022) and Ak et al. (2023) used pretrained transformers to process movie posters and plots, excluding any use of videos, for genre classification on the MM-IMDb dataset (Ovalle et al., 2017). Our work differs from these recent approaches by using pretrained features instead of raw inputs; and by working solely on video and audio, namely, only using raw video pixels and audio waveforms. We also implement a baseline model that is similar to Rodríguez Bribiesca et al. (2021) which uses raw inputs and we empirically show that it performs worse than our proposed models.

3. Methodology

Our overall task is to classify each cinematic trailer into its corresponding genres. While the methods in the literature use raw pixels and audio as the input to the classifier model, we exploit learned features created by existing pretrained deep neural networks (DNNs). Since these features are already created by deep and powerful models, we can train shallow models to process these features and predict the final output, enabling us to cut down on the training time and resources, while still making use of semantically relevant features.

3.1. Creating training data

We used the YouTube trailers in the MovieNet dataset (Huang et al., 2020), choosing the videos with the lowest resolution where the height is at least 300 pixels. We re-encoded the videos using the x265 encoder with a constant rate factor of 23. Using FFmpeg (Tomar, 2006), we detected the scene boundaries and extracted the frames that are exactly in the middle of two scene boundaries. We also re-encoded the audio using the Opus codec, with a bitrate of 48k in mono. For each video, we extract various features using pretrained models and store them so that we can simply load them while training the classification model. We use the genre(s) of each movie from the MovieNet dataset as the ground-truth. Importantly, this dataset

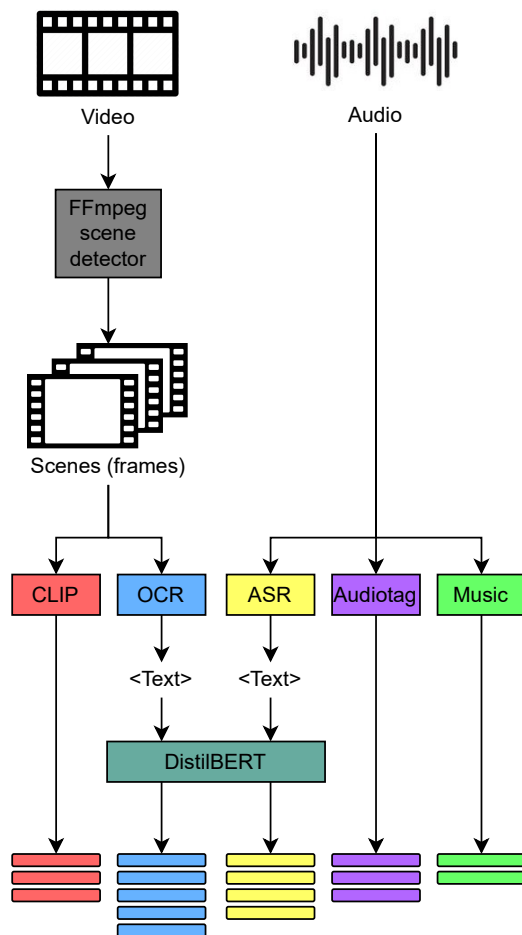


Figure 1: Feature extraction pipeline.

is multi-label, indicating that an individual video can be associated with multiple genres.

3.1.1. Feature extraction

The overall feature extraction pipeline is shown in Figure 1, and each feature is explained below.

CLIP. Contrastive Language-Image Pretraining (CLIP), is the state-of-the-art model for image understanding, that was pretrained using contrastive learning, using a large number of images with captions available on the internet (Radford et al., 2021). The pretrained model first resizes the image

where the longer side has 224 pixels and further takes a 224×224 center crop. Afterward, the pretrained model encodes each frame to a vector with a length of 512.

Audiotag. To extract audio features, we employed a model that was pretrained on the task of audio tagging (Kong et al., 2020). This model operates on chunks of 3 seconds and outputs a vector of probabilities for 527 different labels. Since we are interested in encoding the audio into feature vectors, we don't use the final classification layer and extract the activations prior to that, ending up with vectors with a length of 128, for each audio chunk.

Musicnet. Since we are dealing with cinematic trailers, the soundtrack is an important element that further needs investigation. Hence we extract another audio-related feature, namely the musical feature, using a model that was pretrained on the task of music genre classification (Choi et al., 2016). This model operates on chunks of 22 seconds and outputs a vector of probabilities for 50 different labels. Similarly, we extract the activations prior to the final layer, ending up with vectors with a length of 64, for each audio chunk.

OCR. We also performed optical character recognition (OCR) on each frame using the PaddleOCR model (PaddlePaddle, 2023). We additionally used a pretrained spell correction model on the produced output (Guhr, 2023). The overall output of this stage is in text format.

ASR. We performed automatic speech recognition (ASR) on the audio using OpenAI's Whisper model (Radford et al., 2023). Since this is a learned model that was trained on natural text, there is no need for post-spell correction. However, on the output of Whisper, we used a pretrained language detection model to filter out non-English text (Papariello, 2022). When the output of ASR was in a non-English language, we simply replaced the entire text to denote the language, e.g. changing it to "German language". Similarly, the output of this model is in text format.

DistilBERT. We encode the output text of the OCR and ASR models using a pretrained language model, specifically DistilBERT (Sanh et al., 2019), which is a condensed and compressed variant of the BERT (Bidirectional Encoder Representations from Transformers) model (Devlin et al., 2019), achieved through knowledge distillation (Bucila et al., 2006; Hinton et al.,

2015). DistilBERT utilizes fewer layers than BERT and learns from BERT’s outputs to mimic its behavior. This model converts the input text into tokens of words and sub-words, and encodes each one of them as a vector with a length of 768.

3.2. Classification

We built multiple models to take the previously extracted features of a given video as input and predict the genre of the video. Combining different features is challenging, especially since the encoded vectors have different lengths in both channel and temporal dimensions. For example, considering the audio event and music features, the lengths of each vector are 128 and 64, respectively. Furthermore, since the audio event and music networks operate on chunks of audio with lengths of 3 and 22 seconds respectively, for the same video, there are more embedded vectors for the former. And finally, the number of vectors for the same feature differs between different videos. We addressed these issues using different solutions and different classification models.

We built and trained three distinct models: a *multi-layer perceptron (MLP)*; the *single-transformer* model that integrates features across all modalities; and the *multi-transformer* model, where individual transformers handle features from specific modalities.

The final layer of all our models is a fully-connected (FC) layer with a size of 21, outputting probabilities belonging to 21 different genres. This layer is followed by a sigmoid layer to make sure each output is a probability between 0 and 1. We note that predicted probabilities do not add up to 1 due to the multi-label setting.

3.2.1. MLP

We first implemented a simple MLP classifier to reflect the state-of-the-art method used by Huang et al. (2020). We note that Huang et al. (2020) used raw video and audio from a few small segments of video, while we use pretrained features stemming from the entire video. Our model can be seen in Figure 2. Since MLPs require a fixed-length input, we averaged the feature vectors from each modality along the temporal dimension similar to Huang et al. (2020). We then concatenated the averaged vectors from different modalities and fed them into the MLP.

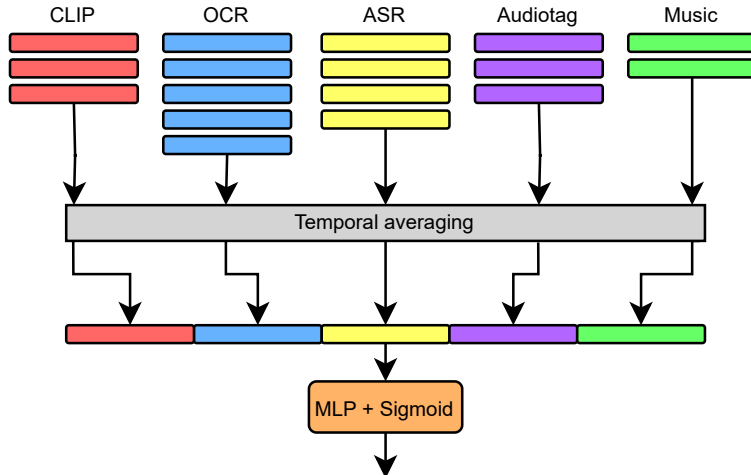


Figure 2: Multilayer perception (MLP). The last layer of the MLP block is the classification layer.

3.2.2. Single-transformer

To address the significant information loss caused by averaging features along the temporal dimension, we utilized a transformer model to leverage short- and long-term correspondences within the video sequence. The transformer is a state-of-the-art sequence processing model, capable of efficiently handling long sequences through its built-in multi-headed attention mechanism (Vaswani et al., 2017). We employed the transformer as a sequence classifier by prepending the sequences with a special learnable vector, called the $\langle \text{CLS} \rangle$ vector. While the transformer generates output vectors for each element of the input sequence, in classification tasks, only the output vector corresponding to the $\langle \text{CLS} \rangle$ vector is forwarded to the next layer, while the others are discarded (Devlin et al., 2019).

While the transformer can handle sequences with varying lengths, the vectors in the sequence still need to have equal sizes along the channel dimension. To ensure that, we used fully-connected (FC) layers to transform feature vectors from different modalities into the same size. We note that this layer is applied to each vector in the sequence individually, hence not changing the number of vectors in its input sequence. Next, we concatenated all the vectors along the time dimension before feeding them into a single transformer. This approach aims at exploiting the correspondences

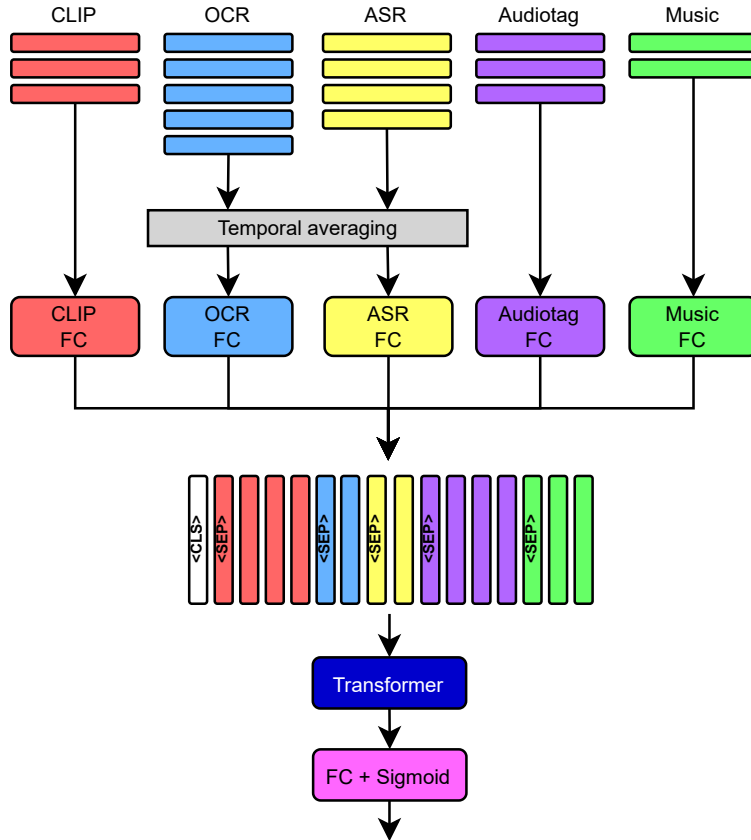


Figure 3: Single-transformer. Regarding the transformer’s output, only the output vector that corresponds to the $\langle \text{CLS} \rangle$ vector is passed to the final classification layer. The $\langle \text{SEP} \rangle$ vectors are used to explicitly separate the elements of each modality.

between different features at the temporal level by processing all of them as a single sequence. However, a trade-off is the complexity of handling different modalities with a single transformer block using a single set of weights. The overall model is shown in Figure 3.

To make sure the transformer can distinguish vectors from different modalities, we used separately learned positional embeddings for sequences from each modality (Wang et al., 2020). We additionally made use of separator ($\langle \text{SEP} \rangle$) vectors, namely, an encoded version of the $\langle \text{SEP} \rangle$ token (Devlin et al., 2019). These are learned vectors that are specific for each modality, prepended to each sequence. Finally, as in almost all classifier transformers,

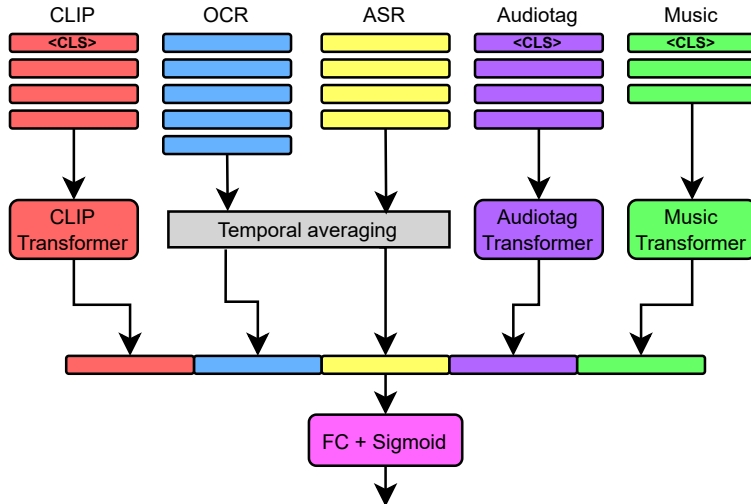


Figure 4: Multi-transformer. Regarding the output of each transformer, only the output vectors that correspond to the $\langle \text{CLS} \rangle$ vectors are used. For inputs that are averaged along the temporal dimensions, no $\langle \text{CLS} \rangle$ vectors are used since temporal averaging already yields a single vector.

we prepended the input sequence to the transformer with a learnable classifying ($\langle \text{CLS} \rangle$) vector, namely, an encoded version of the $\langle \text{CLS} \rangle$ token (Devlin et al., 2019). The output vector corresponding to this classifying vector, i.e. the first vector, is considered as the prediction of the transformer and fed into the final fully-connected layer, while the remaining vectors of the transformer’s output sequence are discarded.

3.2.3. Multi-transformer

With our hypothesis that using a single transformer to process features from multiple modalities can be inefficient due to the increased complexity of the input, we devised our final model that incorporates different transformer models to process features from different modalities, as shown in Figure 4. The inputs of all transformers are prepended with $\langle \text{CLS} \rangle$ token vectors, and the corresponding output vectors are concatenated in channel dimension to be fed into the fully-connected layer to obtain the final probabilities. This approach utilizes the potential correspondences between different features at the global level, though not at the temporal level.

3.2.4. *Optional temporal averaging*

We optionally averaged the text-related features such as OCR and ASR, along the temporal dimension. Non-textual features, namely CLIP, Audio-tag, and Musicnet are obtained from the activations *before* the final prediction layer. Whereas the OCR and ASR features stem from running the DistilBERT model on the predicted text. Here, any error in the predicted text is propagated into the DistilBERT model, corrupting the output features. In order to reduce the resulting noise, we experimented with averaging the textual features, namely OCR and ASR, along the temporal dimension.

3.2.5. *Vision transformer baseline model*

We also implemented a strong baseline model employing state-of-the-art vision transformers (ViT) that work on sequences of 2-dimensional raw visual frames and audio spectrograms, as opposed to the previously introduced models that work with sequences of 1-dimensional pretrained features (Dosovitskiy et al., 2021). ViT segments the image or audio spectrogram into patches, encodes them, and then processes the sequence of encoded vectors using the standard transformer model. While it is purely attention-based, without any convolutional layers, it outperforms CNNs, achieving state-of-the-art results in image (Dosovitskiy et al., 2021), video (Arnab et al., 2021), and audio (Gong et al., 2021) classification.

While our primary objective is video classification, it is important to note that directly applying the video vision transformer (ViViT) (Arnab et al., 2021) to our task is not a suitable approach. There are two main reasons for this. First, ViViT doesn't incorporate audio, which is a crucial aspect of our task. Second, our dataset contains discontinuous frames that depict scenes with no visual continuity between them. As a result, it is more appropriate to handle these frames individually rather than concatenating patches from different frames, as ViViT does.

We utilized a standard pretrained vision transformer (ViT) to process images and a modified version of ViT, known as an audio spectrogram transformer (AST), designed and pretrained specifically for handling audio spectrograms (Gong et al., 2021), to process audio data. Our preliminary experiments showed that training the full model led to overfitting issues that couldn't be mitigated through standard techniques like dropout (Srivastava et al., 2014). To avoid this, we kept the parameters of ViT and AST frozen and replaced their output layers with trainable MLPs, naming them image-MLP and audio-MLP.

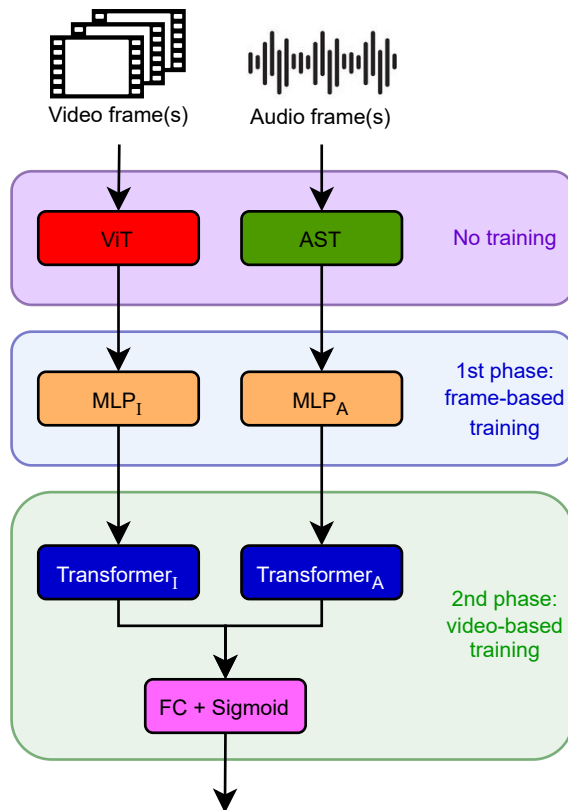


Figure 5: Baseline model that works with 2-dimensional raw video frames and audio (spectrogram) frames.

The subscripts “I” and “A” refer to image and audio, respectively.

The pretrained AST is designed to work with audio spectrograms organized into 10-second segments, which we refer to as “spectrogram frames”. To distinguish them, we use the term “image frames” to describe the visual scenes in our context. Since the trailers in our dataset are longer than 10 seconds, we split the full audio spectrogram into 10-second chunks, with a 50% overlap. Similarly, the ViT works with individual image frames, while our trailers consist of multiple frames. This leads to an output sequence of vectors, where each vector corresponds to an individual image or audio frame. To process these sequences and obtain the final predictions, we developed a fusion module. To this end, we used standard transformer classifiers for both image and audio sequences, yielding a single vector for each modality.

Finally, we concatenated these two vectors and processed them using a linear layer and a sigmoid layer to obtain the final label predictions.

Even when the parameters of pretrained ViT and AST were kept frozen, training the full model in an end-to-end fashion led to overfitting. This is because our training dataset is relatively small with only $33k$ samples, meanwhile, the original authors of ViT and AST trained their models with extensive datasets such as Imagenet with $21M$ samples (Deng et al., 2009) and AudioSet with $2M$ samples (Gemmeke et al., 2017), respectively. To address this overfitting issue while reducing the training dataset size to 1% and 0.1% of the original ones, we employed a two-stage training approach. Initially, we trained the image- and audio-MLPs separately on individual frames. During this phase, the target was the label of the video to which the frames belonged. Following this, we further froze the parameters of the MLPs and trained the fusion module using complete videos. The baseline model and its training scheme can be seen in Figure 5.

3.2.6. Training details and hyperparameters

We applied a filtering process to exclude extremely long and short videos from our dataset. Specifically, we computed the quartiles of video durations and excluded samples that were outside the inner fence, i.e., durations shorter than $Q1 - 1.5 \times IQR$ (corresponding to 19.6 seconds) and longer than $Q3 + 1.5 \times IQR$ (corresponding to 214.4 seconds). This filtering yielded 26412 videos from the original set of 32647. Moreover, following the approach in the MovieNet paper (Huang et al., 2020), we limited the labels to the 21 most frequent genres. These labels are shown in Table 2.

To ensure an unbiased split of the dataset, we arranged the videos alphabetically by their YouTube IDs—given that these IDs are generated randomly. Following the methodology in the original MovieNet paper, we divided the dataset into training, validation, and testing splits with ratios of 0.7, 0.1, and 0.2, corresponding to 18488, 2641, and 5283 samples respectively (Huang et al., 2020). Hyperparameters were determined using a grid search, optimizing for the highest mean average precision on the validation split. The test split was utilized solely for reporting the final results.

While a straightforward way to balance precision and recall is to change the decision threshold, this method does not allow for comparison against the works in the literature where most report their results while setting a fixed decision threshold at 0.5 (Huang et al., 2020). Alternatively, during training, we can balance precision and recall by applying a constant weight

to the loss associated with positive labels. For a single sample, the weighted binary cross-entropy loss becomes:

$$Loss = -\frac{1}{C} \sum_{c=1}^C w \cdot y_c \cdot \log(p_c) + (1 - y_c) \cdot \log(1 - p_c)$$

Here, w represents the weight for positive labels and C is the total number of classes. To align precision and recall values with those observed in MovieNet, we applied a weight of 0.25.

Both the transformer architecture and the averaging module of the MLP architecture support sequences with varying lengths. However, during training, we employed fixed-length inputs to facilitate minibatch processing. Longer sequences were truncated to the desired length, while shorter ones were padded with zero-vectors. The length of feature vectors for each pretrained modality was individually determined through exploratory data analysis, specifically using box plots to analyze sequence lengths across all samples. The maximum sequence length was set to the upper adjacent value (upper whisker of the box plot), which is 1.5 times the interquartile range above the third quartile. Consequently, the sequence lengths were defined as 216 for CLIP, 64 for OCR, 86 for ASR, 140 for Audiotag, and 18 for Musicnet. When using a single-transformer model, the concatenated sequence length totaled 524. It’s important to note that during inference on single samples, our models can handle inputs with varying lengths without the need for padding or truncating.

For the MLP model, the number of layers is 1, the model dimension is 256, and the total number of parameters is $57k$. The single-transformer model has 2 layers, 8 attention heads, and a model dimension of 256, totaling $8.56M$ parameters. The multi-transformer model has individual transformers each with 1 layer, 8 attention heads, and a model dimension of 128, totaling $6.98M$ parameters. Considering the baseline model, both image-MLP and audio-MLP have 2 layers and a dimensionality of 768. The fusing transformers both have 1 layer, 4 attention heads and a dimensionality of 768. All transformers are trained with gradient clipping at a norm of 1. All models are trained with Adam optimizer (Kingma & Ba, 2015), a learning rate of $1e - 5$, a dropout rate of 0.5, and a batch size of 32. We implemented our models using the Pytorch library (Paszke et al., 2019) and trained them on a single NVIDIA Quadro RTX 6000 GPU with 24 GB memory.

4. Experiments and results

We first compare the performance of our models against the baseline model and the models presented in MovieNet paper (Huang et al., 2020). We performed inference on the test set one sample (video) at a time, using the full duration of each video without padding or truncating the sequences. Following the literature, we report macro-averaged precision, recall, and mean average precision (mAP) values, averaged over individual labels (genres).

Adjusting the precision-recall trade-off can be achieved by altering the decision threshold in inference or by modifying the weight assigned to positive labels during training. The MovieNet paper specifies a decision threshold of 0.5 and does not make any reference to label weighting. We kept the decision threshold at 0.5 and aimed for performance similar to the MovieNet paper by assigning a weight of 0.25 to positive (existing) labels.

Table 1 displays the overall performance of our models against the models in the literature. Our baseline model outperforms the models in the literature in terms of recall and mean average precision. Our models using pretrained features outperform all other models across all metrics. Our best-performing model, the multi-transformer improves the state of the art by a large margin. It outperforms all other models in all metrics except precision, where it performs marginally worse than our MLP model.

In Table 2 we compare the performance of our best model multi-transformer, our baseline and the MovieNet model (Huang et al., 2020), across all genres. Out of all 21 genres, our model outperforms the other models in 14 of them in terms of precision, 16 of them in terms of recall, and 20 of them in terms of mean average precision. Outlier values such as very low recall for genres such as Biography, History, and Mystery can be attributed to the inherent imbalance within the MovieNet dataset (Huang et al., 2020). We deliberately avoided balancing the data before training to maintain a fair comparison with the classification model presented in the MovieNet paper. We also believe that this dataset closely reflects the real-world distribution of cinematic genres, accurately portraying the relative rarity of genres such as Biography, History, and Mystery.

In Table 3 we present an ablation study using our best-performing model multi-transformer, demonstrating the gain in terms of mean average precision along with the increase in runtime, due to the addition of each feature while comparing against our baseline model. Runtime is defined as the average duration in seconds, to process a single video, both extracting its pretrained

Table 1: State-of-the-art models vs. ours. P, R, and mAP denote precision, recall, and mean average precision respectively. For precision and recall, 0.5 is used as the decision threshold. Baseline results are obtained from the paper that introduces the MovieNet dataset, where the authors trained different state-of-the-art models, in addition to their own model, on their own dataset (Huang et al., 2020). The full forms of the abbreviated model names are as follows. TSN: Temporal Segment Network (Wang et al., 2016), I3D: Two-Stream Inflated 3D ConvNets (Carreira & Zisserman, 2017), TRN: Temporal Relation Network (Zhou et al., 2018a).

Models	P@0.5	R@0.5	mAP
TSN	78.31	17.95	43.70
I3D	69.58	16.54	35.79
TRN	77.63	21.74	45.23
MovieNet	79.74	24.97	46.88
Baseline	73.78	32.17	59.73
MLP	82.05	33.51	63.16
Single-transformer	81.00	37.11	65.09
Multi-transformer	82.00	38.33	66.02

features and classifying it, during inference. Although the inclusion of textual features such as OCR and ASR initially seems to hurt the performance, possibly due to the text prediction errors, incorporating their averaged version along the temporal dimension reduces the noise and does improve mean average precision. Incorporating textual features also significantly increases runtime. This is primarily because models such as CLIP, Musicnet, and Audiotag are designed to predict a single embedding vector or label, while OCR and ASR models predict text sequences, which can be viewed as predicting many labels corresponding to each word and subword, in an autoregressive way. But most importantly, the settings in which the textual features are excluded yield a better classification performance and faster runtime compared to the baseline. Furthermore, our classification models can seamlessly incorporate features that result from newer and potentially more efficient ASR and OCR models that can be developed in the future, further reducing the overall runtime.

Finally in Figure 6 we attempt to further explain the reasons behind the success of our models. The models in the literature extract a small and fixed number of frames at random locations from each video for classification. The MovieNet baseline (Huang et al., 2020) only uses 8 scenes. They furthermore use MLPs, which cannot process sequences, hence the features from these

Table 2: Results per genre and the macro averages. MovieNet (MN) (Huang et al., 2020) vs baseline model vs our best model multi-transformer. P, R, and mAP denote precision, recall, and mean average precision respectively. For precision and recall, 0.5 is used as the decision threshold.

	P@0.5			R@0.5			mAP		
	MN	Baseline	Ours	MN	Baseline	Ours	MN	Baseline	Ours
Action	73.96	86.59	87.37	22.21	40.48	48.12	54.60	75.64	79.20
Adventure	75.24	74.71	77.51	24.72	22.48	28.67	53.06	56.76	60.93
Animation	93.16	96.38	96.19	74.09	86.59	92.28	86.45	94.65	96.18
Biography	100.00	53.85	70.00	0.04	2.86	5.71	9.13	24.29	36.68
Comedy	68.61	88.55	90.39	48.65	57.94	57.05	68.81	85.47	87.84
Crime	74.12	72.00	80.22	39.30	8.33	22.53	49.25	50.81	60.25
Documentary	85.49	91.30	92.37	4.79	70.87	81.71	21.03	90.29	94.64
Drama	71.16	85.70	89.40	79.42	49.16	55.06	79.95	83.08	86.77
Family	82.55	83.08	86.58	27.11	40.14	48.08	52.19	67.09	73.60
Fantasy	69.83	79.37	87.88	13.51	12.02	20.91	39.12	48.78	59.59
History	82.90	63.16	90.00	12.52	6.49	9.73	34.41	33.63	41.05
Horror	70.03	86.49	88.38	8.76	51.04	55.82	35.51	79.48	84.24
Music	89.04	66.67	68.06	27.24	20.48	29.52	47.13	42.17	48.36
Musical	73.58	70.59	75.00	4.45	11.21	14.02	22.88	32.86	41.27
Mystery	76.42	0.00	57.14	7.76	0.00	1.20	39.70	25.28	30.68
Romance	71.93	76.47	82.32	14.02	13.65	15.75	49.27	51.16	58.59
Sci-Fi	81.35	72.09	83.60	14.51	29.67	37.80	44.14	54.11	68.00
Sport	94.97	69.51	75.86	21.99	42.86	49.62	39.59	60.44	66.80
Thriller	64.98	79.35	78.17	14.50	28.79	37.82	49.80	69.54	71.76
War	86.27	64.71	75.86	12.80	20.89	27.85	34.41	47.66	56.40
Western	88.89	88.89	89.80	51.93	59.70	65.67	73.99	81.14	83.53
AVERAGE	79.74	73.78	82.00	24.97	32.17	38.33	46.88	59.73	66.02

Table 3: The effect of inclusion of pretrained features on mean average precision and runtime. The runtime is defined as the average time, measured in seconds, to classify a single video in inference mode, including the extraction of the pretrained features. The first row shows the performance of the baseline model. Asterisk (*) indicates averaging the features over the sequence (temporal) dimension.

CLIP	Musicnet	Audiotag	OCR	ASR	mAP	Runtime (s.)
					59.73	7.08
✓					64.73	5.32
✓	✓				65.17	5.76
✓	✓	✓			65.31	5.95
✓	✓	✓	✓		63.33	25.85
✓	✓	✓	*		65.46	31.52
✓	✓	✓	✓	✓	64.66	31.44
✓	✓	✓	*	*	66.02	33.59

small number of frames need further averaging along the temporal dimension, causing additional loss of information. Using a simplified experiment, we show the importance of the number of frames fed into the model. Using only the CLIP features, we employ 8, 16, 32, 64, 128 and 256 frames obtained from random locations in each video. We only use the CLIP features since the number of feature vectors differs amongst the pretrained features, and finding an exact number of features that would correspond to the specific number of CLIP features is challenging. We process the features using an MLP and a transformer. Regarding the MLP, the features from different frames are averaged along the temporal dimension as in (Huang et al., 2020). On the contrary, the transformer model is capable of processing sequences hence we feed the features from each frame *as is*.

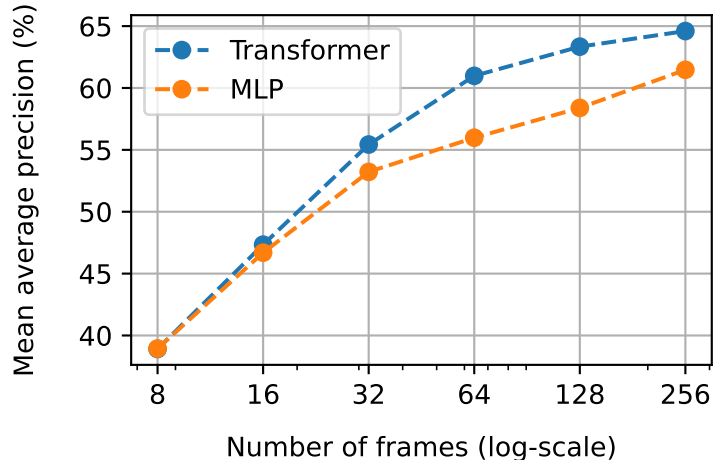


Figure 6: Number of frames vs. mean average precision, only using CLIP features.

In Figure 6, we can see that using a higher number of frames improves the performance significantly. Furthermore, as the number of frames increases, the superiority of the transformer over the MLP becomes apparent, due to its ability to process long sequences seamlessly.

In summary, the success of our models is due to using various pretrained features, using *all* scenes as frames and the entire audio, and finally avoiding averaging features along the temporal dimension and feeding them all to the transformer that is the state-of-the-art model for processing long sequences.

5. Conclusion

In conclusion, our research presents a pioneering approach to movie genre classification, significantly enhancing the performance beyond current methodologies. Our strategy leverages a range of pretrained models to extract and intelligently fuse high-level features associated with visual scenery, characters, text, speech, music, and audio effects. A fundamental element of our approach is the employment of the transformer model, which facilitates the efficient handling of sequences of any length. Importantly, our method capitalizes on the entirety of the information present in all frames of movie trailers, contrasting with traditional models that are restricted to a fixed and low number of frames. Finally, our method purely operates on video

and audio, without requiring any dataset-specific auxiliary data, making it potentially applicable to any video classification task.

Upon the acceptance of our paper, we will make the pretrained features for the entire MovieNet dataset publicly available, along with our genre classification code and trained models, to contribute to ongoing advancements in this field. Even though our models seem complex, involving multiple pretrained models, we made sure all the components and their pretrained weights are packed into a single codebase in a “plug-and-play” manner. We also include a script only for inference on single videos, helping non-academic users to analyze individual videos. These resources are expected to advance movie genre classification and multimedia analysis in general while benefiting both researchers and the broader public.

References

- Ak, K. E., Lee, G.-G., Xu, Y., & Shen, M. (2023). Leveraging efficient training and feature fusion in transformers for multimodal classification. In *IEEE International Conference on Image Processing, ICIP 2023, Kuala Lumpur, Malaysia, October 8-11, 2023* (pp. 1420–1424). IEEE. doi:10.1109/ICIP49359.2023.10223098.
- Almeida, A., de Villiers, J. P., De Freitas, A., & Velayudan, M. (2022). The complementarity of a diverse range of deep learning features extracted from video content for video recommendation. *Expert Systems with Applications*, 192, 116335. doi:10.1016/j.eswa.2021.116335.
- Arnab, A., Dehghani, M., Heigold, G., Sun, C., Lucic, M., & Schmid, C. (2021). ViViT: A video vision transformer. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021* (pp. 6816–6826). IEEE. doi:10.1109/ICCV48922.2021.00676.
- Bain, M., Nagrani, A., Brown, A., & Zisserman, A. (2020). Condensed movies: Story based retrieval with contextual embeddings. In H. Ishikawa, C. Liu, T. Pajdla, & J. Shi (Eds.), *Computer Vision - ACCV 2020 - 15th Asian Conference on Computer Vision, Kyoto, Japan, November 30 - December 4, 2020, Revised Selected Papers, Part V* (pp. 460–479). Springer volume 12626 of *Lecture Notes in Computer Science*. doi:10.1007/978-3-030-69541-5_28.

- Bernstein, L., Sludds, A., Hamerly, R., Sze, V., Emer, J., & Englund, D. (2021). Freely scalable and reconfigurable optical hardware for deep learning. *Scientific reports*, *11*, 3144. doi:10.1038/s41598-021-82543-3.
- Bose, D., Hebbar, R., Somandepalli, K., Zhang, H., Cui, Y., Cole-McLaughlin, K., Wang, H., & Narayanan, S. (2023). MovieCLIP: Visual scene recognition in movies. In *IEEE/CVF Winter Conference on Applications of Computer Vision, WACV 2023, Waikoloa, HI, USA, January 2-7, 2023* (pp. 2082–2091). IEEE. doi:10.1109/WACV56688.2023.00212.
- Bucila, C., Caruana, R., & Niculescu-Mizil, A. (2006). Model compression. In T. Eliassi-Rad, L. H. Ungar, M. Craven, & D. Gunopulos (Eds.), *Proceedings of the Twelfth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Philadelphia, PA, USA, August 20-23, 2006* (pp. 535–541). ACM. doi:10.1145/1150402.1150464.
- Carreira, J., & Zisserman, A. (2017). Quo vadis, action recognition? a new model and the kinetics dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017* (pp. 4724–4733). IEEE Computer Society. doi:10.1109/CVPR.2017.502.
- Cascante-Bonilla, P., Sitaraman, K., Luo, M., & Ordonez, V. (2019). Moviescope: Large-scale analysis of movies using multiple modalities. *arXiv preprint arXiv:1908.03180*, . doi:10.48550/arXiv.1908.03180. arXiv:1908.03180.
- Choi, K., Fazekas, G., & Sandler, M. B. (2016). Automatic tagging using deep convolutional neural networks. In M. I. Mandel, J. Devaney, D. Turnbull, & G. Tzanetakis (Eds.), *Proceedings of the 17th International Society for Music Information Retrieval Conference, ISMIR 2016, New York City, United States, August 7-11, 2016* (pp. 805–811). URL: https://wp.nyu.edu/ismir2016/wp-content/uploads/sites/2294/2016/07/009_Paper.pdf.
- Defernez, M., & Kemsley, E. K. (1999). Avoiding overfitting in the analysis of high-dimensional data with artificial neural networks (ANNs). *Analyst*, *124*, 1675–1681. doi:10.1039/A905556H.

- Deldjoo, Y., Constantin, M. G., Ionescu, B., Schedl, M., & Cremonesi, P. (2018). MMTF-14K: a multifaceted movie trailer feature dataset for recommendation and retrieval. In P. César, M. Zink, & N. Murray (Eds.), *Proceedings of the 9th ACM Multimedia Systems Conference, MMSys 2018, Amsterdam, The Netherlands, June 12-15, 2018* (pp. 450–455). ACM. doi:10.1145/3204949.3208141.
- Deng, J., Dong, W., Socher, R., Li, L., Li, K., & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA* (pp. 248–255). IEEE Computer Society. doi:10.1109/CVPR.2009.5206848.
- Devlin, J., Chang, M., Lee, K., & Toutanova, K. (2019). BERT: pre-training of deep bidirectional transformers for language understanding. In J. Burstein, C. Doran, & T. Solorio (Eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)* (pp. 4171–4186). Association for Computational Linguistics. doi:10.18653/v1/n19-1423.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., & Houlsby, N. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net. URL: <https://openreview.net/forum?id=YicbFdNTTy>.
- Gemmeke, J. F., Ellis, D. P. W., Freedman, D., Jansen, A., Lawrence, W., Moore, R. C., Plakal, M., & Ritter, M. (2017). Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2017, New Orleans, LA, USA, March 5-9, 2017* (pp. 776–780). IEEE. doi:10.1109/ICASSP.2017.7952261.
- Gong, Y., Chung, Y.-A., & Glass, J. R. (2021). AST: Audio spectrogram transformer. In H. Hermansky, H. Cernocký, L. Burget, L. Lamel,

- O. Scharenborg, & P. Motlíček (Eds.), *Interspeech 2021, 22nd Annual Conference of the International Speech Communication Association, Brno, Czechia, 30 August - 3 September 2021* (pp. 571–575). ISCA. doi:10.21437/Interspeech.2021-698.
- Guhr, O. (2023). Spelling-correction-english-base. <https://huggingface.co/oliverguhr/spelling-correction-english-base>. Accessed: 2023-07-03.
- Harper, F. M., & Konstan, J. A. (2016). The MovieLens datasets: History and context. *ACM Trans. Interact. Intell. Syst.*, 5, 19:1–19:19. doi:10.1145/2827872.
- Hinton, G., Vinyals, O., & Dean, J. (2015). Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, . doi:10.48550/arXiv.1503.02531.
- Huang, Q., Xiong, Y., Rao, A., Wang, J., & Lin, D. (2020). MovieNet: A holistic dataset for movie understanding. In A. Vedaldi, H. Bischof, T. Brox, & J.-M. Frahm (Eds.), *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part IV* (pp. 709–727). Springer volume 12349 of *Lecture Notes in Computer Science*. doi:10.1007/978-3-030-58548-8_41.
- Khan, H., Hussain, T., Ullah Khan, S., Ahmad Khan, Z., & Baik, S. W. (2024). Deep multi-scale pyramidal features network for supervised video summarization. *Expert Systems with Applications*, 237, 121288. doi:10.1016/j.eswa.2023.121288.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In Y. Bengio, & Y. LeCun (Eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. URL: <http://arxiv.org/abs/1412.6980>.
- Kong, Q., Cao, Y., Iqbal, T., Wang, Y., Wang, W., & Plumbley, M. D. (2020). PANNs: Large-scale pretrained audio neural networks for audio pattern recognition. *IEEE ACM Trans. Audio Speech Lang. Process.*, 28, 2880–2894. doi:10.1109/TASLP.2020.3030497.

- Miyazawa, K., Kyuragi, Y., & Nagai, T. (2022). Simple and effective multi-modal learning based on pre-trained transformer models. *IEEE Access*, *10*, 29821–29833. doi:10.1109/ACCESS.2022.3159346.
- Niu, S., Liu, Y., Wang, J., & Song, H. (2020). A decade survey of transfer learning (2010–2020). *IEEE Transactions on Artificial Intelligence*, *1*, 151–166. doi:10.1109/TAI.2021.3054609.
- Ovalle, J. E. A., Solorio, T., Montes-y-Gómez, M., & González, F. A. (2017). Gated multimodal units for information fusion. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net. URL: <https://openreview.net/forum?id=B1akgy9xx>.
- PaddlePaddle (2023). PaddleOCR. <https://github.com/PaddlePaddle/PaddleOCR>. Accessed: 2023-07-03.
- Papariello, L. (2022). xlm-roberta-base-language-detection. <https://huggingface.co/papluca/xlm-roberta-base-language-detection>. Accessed: 2023-07-03.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., & Chintala, S. (2019). PyTorch: An imperative style, high-performance deep learning library. In H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. B. Fox, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada* (pp. 8024–8035). URL: https://proceedings.neurips.cc/paper_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., & Sutskever, I. (2021). Learning transferable visual models from natural language supervision. In M. Meila, & T. Zhang (Eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML*

- 2021, 18-24 July 2021, Virtual Event (pp. 8748–8763). PMLR volume 139 of *Proceedings of Machine Learning Research*. URL: <http://proceedings.mlr.press/v139/radford21a.html>.
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., & Sutskever, I. (2023). Robust speech recognition via large-scale weak supervision. In A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, & J. Scarlett (Eds.), *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA* (pp. 28492–28518). PMLR volume 202 of *Proceedings of Machine Learning Research*. URL: <https://proceedings.mlr.press/v202/radford23a.html>.
- Ray, A., & Kolekar, M. H. (2024). Transfer learning and its extensive appositeness in human activity recognition: A survey. *Expert Systems with Applications*, 240, 122538. doi:10.1016/j.eswa.2023.122538.
- Rodríguez Bribiesca, I., López Monroy, A. P., & Montes-y-Gómez, M. (2021). Multimodal weighted fusion of transformers for movie genre classification. In A. Zadeh, L.-P. Morency, P. P. Liang, C. Ross, R. Salakhutdinov, S. Poria, E. Cambria, & K. Shi (Eds.), *Proceedings of the Third Workshop on Multimodal Artificial Intelligence* (pp. 1–5). Mexico City, Mexico: Association for Computational Linguistics. doi:10.18653/v1/2021.maiworkshop-1.1.
- Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, . doi:10.48550/arXiv.1910.01108.
- Simoes, G. S., Wehrmann, J., Barros, R. C., & Ruiz, D. D. (2016). Movie genre classification with convolutional neural networks. In *2016 International Joint Conference on Neural Networks, IJCNN 2016, Vancouver, BC, Canada, July 24-29, 2016* (pp. 259–266). IEEE. doi:10.1109/IJCNN.2016.7727207.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15, 1929–1958. doi:10.5555/2627435.2670313.

- Tan, C., Sun, F., Kong, T., Zhang, W., Yang, C., & Liu, C. (2018). A survey on deep transfer learning. In V. Kůrková, Y. Manolopoulos, B. Hammer, L. Iliadis, & I. Maglogiannis (Eds.), *Artificial Neural Networks and Machine Learning – ICANN 2018* (pp. 270–279). Cham: Springer International Publishing volume 11141. doi:10.1007/978-3-030-01424-7_27.
- Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. In K. Chaudhuri, & R. Salakhutdinov (Eds.), *Proceedings of the 36th International Conference on Machine Learning* (pp. 6105–6114). PMLR volume 97 of *Proceedings of Machine Learning Research*. URL: <https://proceedings.mlr.press/v97/tan19a.html>.
- Tomar, S. (2006). Converting video formats with ffmpeg. *Linux journal*, 2006, 10. URL: <https://dl.acm.org/doi/abs/10.5555/1134782.1134792>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems* (pp. 5998–6008). URL: <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- Wang, B., Zhao, D., Lioma, C., Li, Q., Zhang, P., & Simonsen, J. G. (2020). Encoding word order in complex embeddings. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net. URL: <https://openreview.net/forum?id=Hke-WTVtwr>.
- Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., & Gool, L. V. (2016). Temporal segment networks: Towards good practices for deep action recognition. In B. Leibe, J. Matas, N. Sebe, & M. Welling (Eds.), *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VIII* (pp. 20–36). Springer volume 9912 of *Lecture Notes in Computer Science*. doi:10.1007/978-3-319-46484-8_2.

- Wehrmann, J., & Barros, R. C. (2017). Movie genre classification: A multi-label approach based on convolutions through time. *Appl. Soft Comput.*, *61*, 973–982. doi:10.1016/j.asoc.2017.08.029.
- Yadav, A., & Vishwakarma, D. K. (2020). A unified framework of deep networks for genre classification using movie trailer. *Appl. Soft Comput.*, *96*, 106624. doi:10.1016/j.asoc.2020.106624.
- Zhou, B., Andonian, A., Oliva, A., & Torralba, A. (2018a). Temporal relational reasoning in videos. In V. Ferrari, M. Hebert, C. Sminchisescu, & Y. Weiss (Eds.), *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part I* (pp. 831–846). Springer volume 11205 of *Lecture Notes in Computer Science*. doi:10.1007/978-3-030-01246-5_49.
- Zhou, B., Lapedriza, À., Khosla, A., Oliva, A., & Torralba, A. (2018b). Places: A 10 million image database for scene recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, *40*, 1452–1464. doi:10.1109/TPAMI.2017.2723009.
- Zhou, H., Hermans, T., Karandikar, A. V., & Rehg, J. M. (2010). Movie genre classification via scene categorization. In A. D. Bimbo, S. Chang, & A. W. M. Smeulders (Eds.), *Proceedings of the 18th International Conference on Multimedia 2010, Firenze, Italy, October 25-29, 2010* (pp. 747–750). ACM. doi:10.1145/1873951.1874068.