



PDF Download
3746027.3755041.pdf
17 December 2025
Total Citations: 0
Total Downloads: 68

Latest updates: <https://dl.acm.org/doi/10.1145/3746027.3755041>

RESEARCH-ARTICLE

Character-Centric Understanding of Animated Movies

ZHONGRUI GUI, University of Oxford, Oxford, Oxfordshire, U.K.

JUNYU XIE, University of Oxford, Oxford, Oxfordshire, U.K.

TENGDA HAN, University of Oxford, Oxford, Oxfordshire, U.K.

WEIDI XIE, Shanghai Jiao Tong University, Shanghai, China

ANDREW ZISSERMAN, University of Oxford, Oxford, Oxfordshire, U.K.

Open Access Support provided by:

University of Oxford

Shanghai Jiao Tong University

Published: 27 October 2025

[Citation in BibTeX format](#)

MM '25: The 33rd ACM International
Conference on Multimedia
October 27 - 31, 2025
Dublin, Ireland

Conference Sponsors:
SIGMM

Character-Centric Understanding of Animated Movies

Zhongrui Gui
VGG, University of Oxford
Oxford, United Kingdom
zhongrui@robots.ox.ac.uk

Junyu Xie
VGG, University of Oxford
Oxford, United Kingdom
jyx@robots.ox.ac.uk

Tengda Han
VGG, University of Oxford
Oxford, United Kingdom
htd@robots.ox.ac.uk

Weidi Xie
SAI, Shanghai Jiao Tong University
Shanghai, China
weidi@sjtu.edu.cn

Andrew Zisserman
VGG, University of Oxford
Oxford, United Kingdom
az@robots.ox.ac.uk

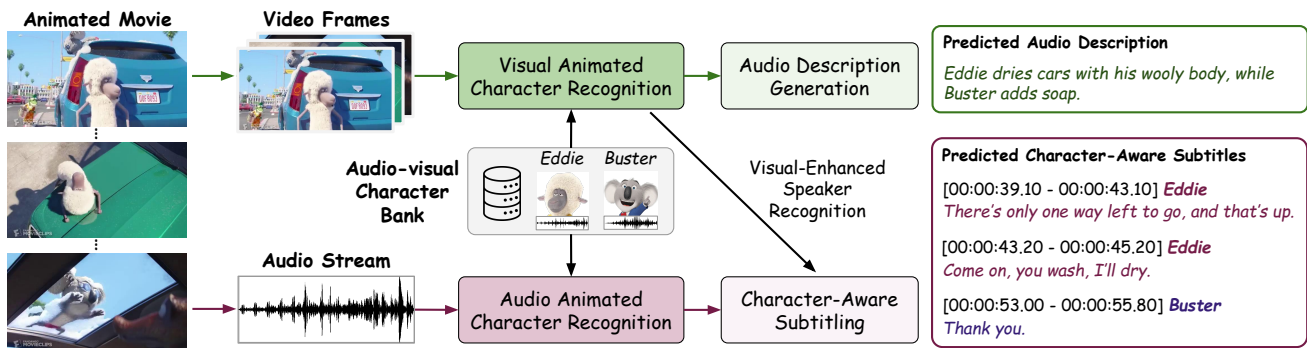


Figure 1: *Character-Centric Understanding of Animated Movies*. We propose a pipeline for recognising characters in animated movies across both visual and audio domains. Our approach involves constructing an audio-visual character bank from online sources, which is then used for video and audio (voice) animated character recognition. We use character recognition for two downstream tasks: Audio Description (AD) generation and character-aware subtitling. The movie clip is from *Sing*.

Abstract

Animated movies are captivating for their unique character designs and imaginative storytelling, yet they pose significant challenges for existing recognition systems. Unlike the consistent visual patterns detected by conventional face recognition methods, animated characters exhibit extreme diversity in their appearance, motion, and deformation. In this work, we propose an audio-visual pipeline to enable automatic and robust animated character recognition, and thereby enhance character-centric understanding of animated movies. Central to our approach is the automatic construction of an audio-visual character bank from online sources. This bank contains both visual exemplars and voice (audio) samples for each character, enabling subsequent multi-modal character recognition despite long-tailed appearance distributions. Building on accurate character recognition, we explore two downstream applications: Audio Description (AD) generation for visually impaired audiences, and character-aware subtitling for the hearing impaired. To support research in this domain, we introduce CMD-AM, a new dataset of 75

animated movies with comprehensive annotations. Our character-centric pipeline demonstrates significant improvements in both accessibility and narrative comprehension for animated content over prior face-detection-based approaches. For the code and dataset, visit https://www.robots.ox.ac.uk/~vgg/research/animated_ad/.

CCS Concepts

• **Computing methodologies** → *Video summarization*; **Computer vision**; *Speech recognition*; *Object identification*; • **Human-centered computing** → *Accessibility technologies*; • **Information systems** → *Multimedia information systems*.

Keywords

Animated Character Recognition, Speaker Diarisation, Audio Description, Character-Aware Subtitling

ACM Reference Format:

Zhongrui Gui, Junyu Xie, Tengda Han, Weidi Xie, and Andrew Zisserman. 2025. Character-Centric Understanding of Animated Movies. In *Proceedings of the 33rd ACM International Conference on Multimedia (MM '25)*, October 27–31, 2025, Dublin, Ireland. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3746027.3755041>

1 Introduction

When we were children, animated movies were more than just entertainment – they introduced us to our favourite characters, sparked our imagination with captivating stories, and left us with



This work is licensed under a Creative Commons Attribution 4.0 International License. *MM '25, Dublin, Ireland*

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2035-2/2025/10
<https://doi.org/10.1145/3746027.3755041>

lasting memories. Today, modern computer vision systems are rapidly advancing towards superhuman capabilities across many domains, from recognising bird species to describing fine-grained actions. Yet, intriguingly, existing technologies often struggle to comprehend animated movies – a medium effortlessly understood by young minds.

Ask any four-year-old child to describe an animation they have just seen, and the reason for this computer vision system failing is obvious: the child names the characters and can recognise them, but current computer vision systems using face recognition cannot. The challenge is that animated characters have *huge* diversity. Although there is typically some anthropomorphic element, their appearance can range from human-like faces, to cars, fish, blobs, skeletons, ... (and that is just Pixar movies). This long-tailed distribution of appearances poses a significant challenge for recognition methods. In contrast, (human) face recognition, which is the workhorse for character recognition in movies, is a far more constrained problem: a face has eyes, a nose, a mouth, maybe some facial decorations (beard, glasses), in a regular configuration.

In this paper, we present a computer vision pipeline for animated movies with a key design to recognise and track *animated characters*, despite the long-tailed distribution of character appearance and exaggerated actions. We achieve this by using *visual* recognition methods to detect, segment, track and encode the appearance of the characters, and also *audio* recognition methods, since the voice of the animated character remains mostly the same throughout the movie. We leverage online movie forums like IMDb and Fandom as external knowledge banks, and build a novel audio-visual character bank for animated movies automatically, including the voice samples of each voice actors. Both the visual and audio cues are then utilised to recognise and track animated characters effectively.

Recognising animated characters enables multiple downstream applications, and we develop two of these in this paper: Audio Description generation and character-aware subtitling, as illustrated in Fig. 1. Both of these applications depend heavily on character recognition.

In movies and TV series, *Audio Descriptions (ADs)* are narrations that convey essential visual information to complement the original soundtrack. They help ensure a continuous and coherent narrative flow, enabling *visually-impaired* audiences to follow the plot effectively. ADs prioritise the most visually salient and story-relevant content such as character dynamics and significant objects, while omitting redundant details like background figures.

Character-aware subtitling, on the other hand, concentrates on helping *hearing-impaired* people follow dialogues in video material. This actually requires determining three things for each utterance: *what is being said*, *when it is said*, and *who is saying it*. By explicitly attributing dialogue to specific characters, it enhances the clarity of conversation flow, particularly in scenes with overlapping speech or multiple speakers.

In this paper, we make contributions over four areas:

(i) We develop an automatic pipeline for constructing an audio-visual character bank from online sources, consisting of a character appearance bank and a character voice bank (Sec. 3).

(ii) We tackle character recognition by leveraging both visual and audio information. For accurate visual animated character recognition, we propose Track-Guided Region Proposals (TGRP) and adapt

DINOv2 to animated content. Additionally, we introduce visual enhancements to improve audio-only speaker recognition (Sec. 4).

(iii) We use our animated character recognition to solve two downstream tasks – Audio Description generation and character-aware subtitling; where previous face-based methods fail due to the large domain gap with animated content (Sec. 5).

(iv) We introduce a new dataset, CMD-AM, consisting of 75 animated movies with unified annotations, including character bounding boxes, ground truth Audio Descriptions, and sentence-level diarisation labels (Sec. 6).

Overall, our pipeline achieves superior character recognition accuracy compared to conventional human face detection methods. This, in turn, enables the generation of high-quality Audio Descriptions and character-aware subtitles, significantly enhancing the accessibility of animated movies.

2 Related Work

2.1 Animated Character Recognition

Animated character recognition is often considered more challenging than real-world human identification, primarily due to large domain variations in animated media in terms of character geometry and styles. As a result, conventional human facial recognition methods [3, 12, 13, 31] often struggle with this task.

To address this issue, several datasets [2, 49, 52, 63, 72] have been proposed for *image-based* animated character recognition, with follow-up research focusing on animated face detection [15, 41, 69] and recognition [19, 29, 34, 39, 43, 55, 67, 71].

On the other hand, character recognition in animated *videos* remains relatively under-explored. Kim et al. [30] adapted Faster R-CNN [48] for detecting characters in animated movies, while Somandepalli et al. [54] employed an unsupervised clustering approach to “discover” characters. More recently, CAST [40] improved character clustering through self-supervised training of animation representations. In contrast, our method automatically constructs character banks for target animated movies and formulates the recognition process as a few-shot retrieval task, thereby achieving scalable and generalisable animated character recognition.

2.2 Audio Description Generation

Audio Descriptions (ADs) are narrations in movies and TV series that provide visual information complementary to the audio, helping visually impaired individuals better understand the story. Existing AD generation approaches can be broadly divided into fine-tuned models [17, 20–22, 35, 59] and training-free methods [9, 64, 65, 70]. The former typically involve end-to-end architectures trained on ground truth AD sentences, while the latter adopts multi-stage strategies, leveraging pre-trained vision-language models (VLMs) and large language models (LLMs) as core components at different stages.

One of the key challenges in AD generation is accurately identifying the characters to be described. AutoAD-II [20] first addressed this by training character prediction modules on existing movie datasets [26]. AutoAD-Zero [65] further improved accuracy by incorporating more robust off-the-shelf face recognition methods. However, these approaches are not directly applicable to animated movies due to the significant domain gap between real and animated

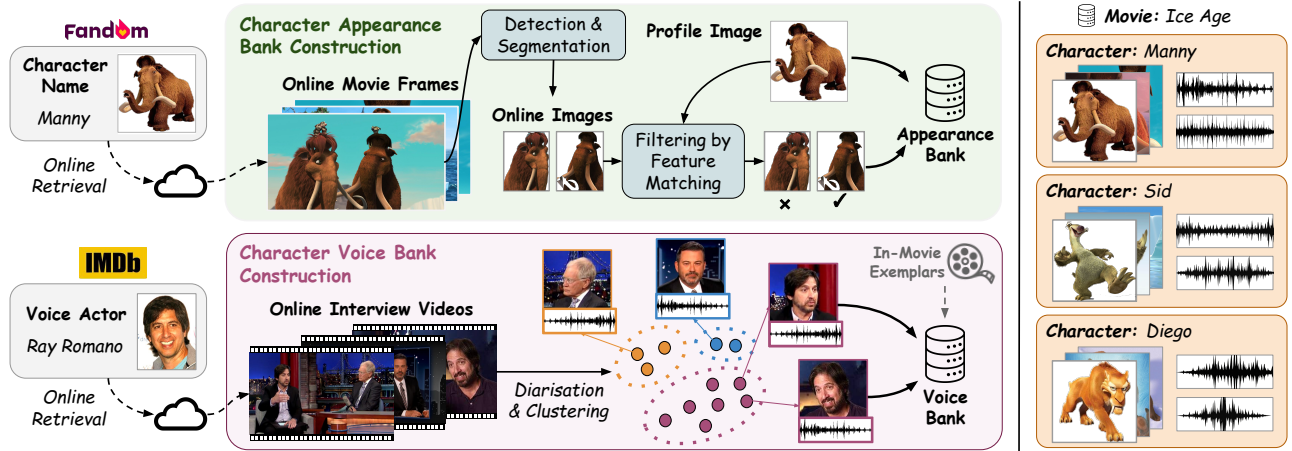


Figure 2: *Audio-Visual Character Bank Construction*. *Left*: The audio-visual character bank is built automatically from online sources. Character names, corresponding voice actor names, and profile images are retrieved from IMDB and Fandom. We then construct a character appearance bank and character voice bank from online sources. *Right*: An example audio-visual character bank for *Ice Age*.

characters. In this work, we propose a scalable and generalisable pipeline that accurately identifies animated characters and supports downstream AD generation.

2.3 Character-Aware Subtitling

Character-aware subtitling [27, 32] is a recently emerging task that aims to generate subtitles for media such as movies and TV series while identifying the corresponding speakers. This significantly enhances the accessibility of video content for hearing-impaired people. The task generally involves two key steps: (i) Automatic speech recognition (ASR) [5, 45, 46, 51], which transcribes spoken dialogue into text; and (ii) Speaker diarisation, which determines “who spoke when” in the audio stream.

In particular, speaker diarisation can be categorised into two main directions. The first focuses on audio-only diarisation [18, 25, 33, 61, 62, 68]. The second direction, benefiting from recent advances in multimodal learning, leverages both audio and visual signals [38] to enhance diarisation performance. For instance, visual cues such as lip movements [23, 24] and character faces [10, 11, 44, 66] have been shown to improve the accuracy of speaker attribution. In this work, we address speaker diarisation for animated movies and demonstrate that accurate visual recognition of animated characters can significantly boost diarisation performance.

3 Animated Character Bank Construction

In animated movies, characters take on various forms; some resemble humans, while others have more abstract or imaginative appearances. Such a long-tailed distribution of character appearances presents a significant challenge to character recognition.

In this section, we introduce an automatic pipeline for constructing an audio-visual character bank that captures both the appearance and corresponding voice of each character.

As exemplified in Fig. 2 (right), for each movie, we maintain a character name list $\mathcal{S} = \{C^1, \dots, C^N\}$, where the p -th character

$C^p \in \mathcal{S}$ is associated with a tuple $\{I^p, A^p\}$, where I^p denotes a set of n exemplar images $I^p = \{I_1^p, \dots, I_n^p\}$, and $A^p = \{A_1^p, \dots, A_m^p\}$ represents a set of m exemplar voice (audio) segments. In Sec. 3.1, we elaborate on the construction of the appearance bank, and in Sec. 3.2, we detail the voice counterpart.

3.1 Character Appearance Bank Construction

To construct the appearance bank for each movie, we first retrieve a character name list along with their profile images from online movie databases [16, 28], as illustrated in Fig. 2 (left). However, the set of profile images is often limited in size and typically depicts characters in a single, static pose, whereas animated characters exhibit a wide range of visual morphisms. To address this, we augment the character bank with additional images collected via web searching.

Candidate Image Retrieval. Unlike profile images, web-crawled images often contain complex backgrounds and multiple characters. To isolate the character of interest, we extract all regions likely to contain relevant figures. Using an open-vocabulary detection model (OWLv2 [37]) to generate candidate boxes, we use these boxes as visual prompts for SAM2 [47] to segment potential character regions in the web-crawled images.

Image Filtering. We then apply an additional filtering step to retain only the images corresponding to the target character. Specifically, for each character, we compute the cosine similarity between the DINOv2 [42] feature of the profile image and each candidate character image. If the cosine similarity exceeds a predefined threshold, the candidate image is included in the appearance bank.

DINOv2 Adaptation. Nevertheless, we observe that pre-trained DINOv2 features often struggle to distinguish between visually similar characters, such as the Mammoths (Manny and Ellie) from *Ice Age* in Fig. 2.

To mitigate this, we further adapt the DINOv2 feature extractor on the character bank to deal with the long-tailed distribution of the appearance of animated characters. We individually fine-tune DINOv2 for each movie using our previously crawled images.

The fine-tuning process follows a contrastive training scheme. For each character $C^p \in \mathcal{S}$, we randomly sample a positive pair (I_i^p, I_j^p) from their exemplar images, while negative samples $\{I_j^q\}$ are drawn from the remaining characters. Finally, we apply the standard InfoNCE [57] loss to fine-tune the DINOv2 model. Full details are given in the arXiv version of this paper.

3.2 Character Voice Bank Construction

In addition to the appearance bank, we also construct a voice bank containing multiple speech segments for each character. To obtain these voice exemplars, we consider two major sources: (i) online interview videos, and (ii) in-movie speech segments.

Interview Video Exemplar Extraction. As shown in Fig. 2 (left), we first associate each character with their corresponding voice actor using information available from online movie databases. We then query search engines to retrieve multiple interview videos for each voice actor. To extract the speech segments of the target actor from these videos, we leverage the prior that the actor’s speech (as the interviewee) is likely to be the most frequently occurring voice across the interview videos.

Specifically, we employ a speaker diarisation model [7] to partition the audio streams into multiple speaker clusters. We then select the largest resulting cluster and add its speech segments to the character voice bank.

In-Movie Exemplar Extraction. However, compared to live-action movies, animated movies often exhibit a greater disparity between characters’ voices and the natural voices of their voice actors. To further mitigate this gap, we augment the character voice bank with in-movie speech segments, guided by our visual recognition results (detailed in Sec. 4.1).

To be more specific, for each character, we select visually predicted tracks with the visual confidence scores higher than a defined threshold for the corresponding label. We then use an audio-visual synchronisation model [1] to generate a synchronisation score to assess how likely the character on the track aligns with the concurrent speech. If the synchronisation score exceeds a predefined threshold, the corresponding in-movie speech segments are added to the character voice bank. For more details on the audio-visual synchronisation model, please refer to the arXiv version.

4 Audio-Visual Character Recognition

In this section, we address animated character recognition from two complementary perspectives: visual identification, detailed in Sec. 4.1, and speaker attribution, described in Sec. 4.2.

4.1 Visual Recognition of Animated Characters

Given the constructed character appearance bank (as detailed in Sec. 3.1), we aim to visually recognise the animated character following a two-stage pipeline: (i) region proposal, which localises and tracks candidate character regions, and (ii) character identification, which assigns character IDs to each proposed track.

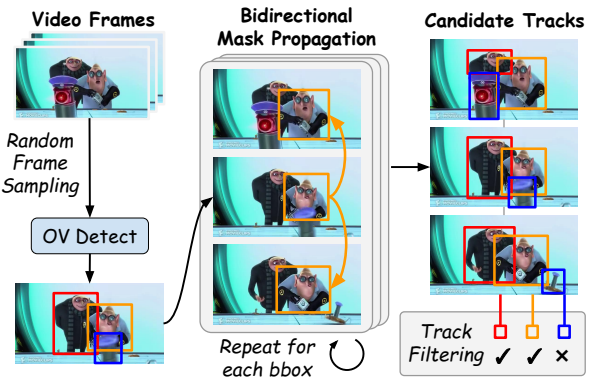


Figure 3: Track-Guided Region Proposals (TGRP). We adopt an open-vocabulary detection (OV Detect) model to generate bounding box (bbox) proposals for sampled movie frames. These proposals are then propagated bidirectionally to produce candidate tracks, which are subsequently refined through a track filtering process.

Track-Guided Region Proposal. To ensure consistency across video clips, we propose a tracking-based approach for region proposal. As a preprocessing step, we apply a shot detection model to partition the video into individual shots, which naturally defines the boundaries for subsequent tracking operations.

As illustrated in Fig. 3, we randomly sample one seed frame from each shot and apply OWLv2 [37], an open-vocabulary detection model, to generate coarse bounding box proposals. These proposals are then used to prompt SAM2 [47], a segmentation and tracking model, which initialises a bi-directional mask propagation process to generate a set of candidate tracks.

We sample $n = 3$ seed frames to repeat the process and generate a set of tracks for each seed. Then, for the track sets, we perform tripartite matching based on overlapping IoUs, considering a match valid if the IoU exceeds a predefined threshold. We discard any unmatched tracks, as they are likely to be outliers.

Character Identification. We use the adapted DINOv2 model (detailed in Sec. 3.1) to assign character IDs to predicted tracks. Specifically, for each track, we uniformly sample 5 frames and compute their features $\{f_j : j = 1, \dots, 5\}$. To determine if the track matches with a character C^p , we calculate the cosine similarity between each query feature f_j and every exemplar feature z_i^p from the visual character bank. The visual matching score s_{vm}^p is then obtained by averaging over the top- k ($k = 3$) cosine similarities, i.e. $s_{vm}^p = \max_j \left(\frac{1}{k} \sum_{\text{top-}k, i} \langle f_j, z_i^p \rangle \right)$. Finally, we assign the track to the character with the highest matching score.

4.2 Audio Recognition of Animated Characters

Beyond visual recognition, we also aim to identify speakers in the audio stream, leveraging the pre-constructed character voice bank (detailed in Sec. 3.2). We start with an audio-only approach, and subsequently explore the incorporation of visual cues to further enhance speaker recognition.

Audio-Only Speaker Recognition. We begin with a straightforward audio-only speaker recognition approach, which matches speech segments in animated movies to voice exemplars in the character bank, thereby predicting speaker labels for each segment.

Given the audio stream of an animated movie, we first extract all speech segments using WhisperX [5]. We then follow the strategy described in Sec. 3.2, where speech segments are clustered using a diarisation model. For each cluster, we extract its “centroid feature” by averaging ECAPA-TDNN [14] voice embeddings of all contained speech segments.

To associate these clusters with characters in the voice bank, we adopt a feature-based matching process, similar to that introduced in Sec. 4.1. Specifically, we use the cluster’s centroid feature to query the ECAPA-TDNN features of different exemplar speech segments and compute their pairwise cosine similarities. For each character, the top- k ($k = 3$) similarities are then recorded and averaged to obtain an audio matching score. We assign the cluster to the character with the highest score, s_{am} . Furthermore, for each speech segment in the cluster, we estimate a confidence score c_a by scaling s_{am} with the cosine similarity between its feature and the centroid feature of the cluster.

Visual-Enhanced Speaker Recognition. The proposed audio-only recognition, however, is prone to background noise commonly present in movies, such as music and environmental sounds. In addition, the character voice bank may also include noisy exemplars. To address these issues, we incorporate visual character recognition to enhance speaker identification.

To further enhance speaker recognition using visual cues, we focus on speech segments with confidence scores $\{c_a\}$ below a pre-defined threshold, categorising them as low-confidence predictions.

For each such audio segment, we retrieve all visually predicted character tracks that overlap with its time range, as illustrated in Fig. 4. For each retrieved track, we apply the same audio-visual synchronisation model we use to retrieve in-movie voice exemplars in Sec. 3.2 to compute the synchronisation score s_{sync} . This score indicates the likelihood that the track corresponds to the voice, specifically measuring the correlation between the temporal appearance variations of the tracked character and the temporal audio variations.

In addition, we incorporate the visual matching score s_{vm} (introduced in Sec. 4.1) as a complementary measure of visual track reliability. We then combine the synchronisation and visual matching scores into a visual confidence score $c_v = s_{sync} \cdot s_{vm}$. This score is scaled by a hyperparameter λ and compared to the audio-only confidence score c_a . If $\lambda \cdot c_v > c_a$, we interpret this as a more reliable visual-based prediction and update the audio-only prediction with the character ID inferred from the visual track.

5 Applications for Animated Movies

Building on the visual and audio character recognition results, we explore downstream applications that enhance the understanding of animated movies. In particular, we focus on improving accessibility through two tasks: Audio Description (AD) generation (Sec. 5.1) and character-aware subtitling (Sec. 5.2), targeting visually and hearing-impaired audiences, respectively.

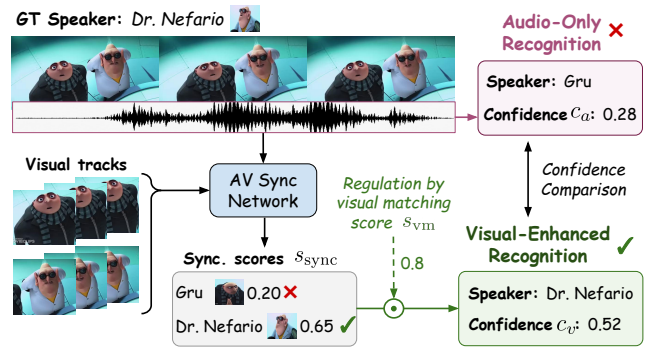


Figure 4: Visual-Enhanced Speaker Recognition. We leverage an audio-visual synchronisation (AV Sync) network to compute synchronisation scores between simultaneous audio and character tracks. These scores are then used to determine whether a character is aligned with the audio signal, indicating that they are the speaker. The synchronisation scores are further regulated using visual matching scores. This results in a visual-enhanced speaker prediction and an associated visual confidence score. If the visual confidence exceeds the audio counterpart, we update the audio-only identified speaker with the visually recognised one.

5.1 Audio Description Generation

Audio Descriptions (ADs) are narrations that capture the most salient visual content in videos, serving as a complement to existing audio information. An AD generation method therefore takes video clips as input and aims to automatically generate AD sentences in textual form.

In this work, we build upon AutoAD-Zero [65], a recent training-free AD generation framework. This method involves two major stages: In Stage I, a video-language model (VideoLLM) is prompted to generate a comprehensive description of the input video frames, covering key visual elements such as characters, actions, etc. In Stage II, a large language model (LLM) refines this dense text description into a concise and narratively coherent AD output.

Prior work [20] has shown that accurate character recognition significantly enhances the quality of ADs. Accordingly, the original AutoAD-Zero [65] employs a robust face detection method (InsightFace [13]) for character identification. However, this approach underperforms in animated content, where conventional face recognition models struggle to detect animated characters.

To address this domain gap, we integrate our animated character bank and corresponding visual character recognition results into the AD generation framework. This is achieved by prompting the Stage I VideoLLM in a training-free manner. Specifically, we visually overlay the predicted character bounding boxes onto each input frame. Each character box is assigned a unique colour, and a text prompt provides the mapping between colours and character names. This allows the VideoLLM to ground the visual entities to their corresponding names and refer to them explicitly during the dense description generation.

Table 1: Overview of the CMD-AM dataset, where “AM” stands for “Animated Movie”. The dataset consists of 75 animated movies, with (partial) annotations including Audio Descriptions (ADs), per-frame character bounding boxes, and speaker labels for each audio segment.

Annotations	No. Movies	Duration (h)	No. Annotations
Audio Descriptions	75	24.5	9,348 ADs
Character Boxes	8	1.5	126,977 frames
Diarisation	8	1.5	1,336 audio segments

5.2 Character-Aware Subtitling

Character-aware subtitling aims to temporally localise spoken dialogues in movies, transcribe the corresponding audio, and assign each utterance to its corresponding speaker. This involves two main components: automatic speech recognition (ASR) and speaker diarisation.

Automatic Speech Recognition. We utilise WhisperX to transcribe the video, producing individual dialogue segments along with their precise timestamps.

Speaker Diarisation. Based on the transcribed timestamps, we extract the corresponding speech segments from the audio stream. These segments are then assigned to characters in the voice bank using the speaker recognition procedure described in Sec. 4.2.

Finally, each speaker label is paired with its corresponding dialogue transcription to form character-aware subtitles.

6 The CMD-AM Dataset

In this section, we introduce CMD-AM (“AM” stands for “Animated Movies”), a curated dataset comprising animated video clips sourced from the Condensed Movie Dataset (CMD) [6]. In total, the dataset contains 565 video clips from 75 animated movies. The statistics are summarised in Tab. 1. CMD-AM provides a wide range of annotations including ground truth Audio Descriptions (Sec. 6.1), grounded bounding boxes for character recognition (Sec. 6.2), and sentence-level speaker labels for diarisation (Sec. 6.3). Note that these annotations are used *for evaluation purposes only*, and our method is not trained on them. The complete list of movies included in the CMD-AM dataset is available in the arXiv version.

6.1 Ground Truth Audio Description

We provide ground truth (GT) Audio Descriptions (ADs) for 565 video clips in 75 animated movies, covering a duration of 24.47 hours. To obtain the GT AD annotations, we first collect AD-blended audio streams narrated by volunteers from AudioVault [4]. We then follow the processing pipeline introduced in [22, 65], which extracts clean, text-form ADs from the downloaded audio streams through a series of transcription, diarisation, and post-filtering steps.

6.2 Frame-Level Character Box Annotation

We provide frame-level bounding box annotations on characters for 40 video clips across 8 animated movies (listed in the arXiv version). In each clip, the video is sampled and labelled at 23.98 fps, resulting

in a total of 126,977 annotated frames, covering approximately 1.5 hours of video footage.

To obtain these frame-level bounding box annotations, we employ a three-stage process: (i) We first partition each video clip into multiple shots using an off-the-shelf shot detection model [53]. Within each shot, we randomly sample a few frames and manually annotate them with precise bounding boxes and corresponding character labels; (ii) We accelerate the annotation process by leveraging SAM2 to bidirectionally propagate the annotated boxes across each shot; (iii) Finally, we manually inspect every frame to correct any missing or erroneous boxes and labels, ensuring the overall accuracy of the annotations.

6.3 Sentence-Level Diarisation Annotation

We provide sentence-level speaker labels for the *same* 40 video clips across 8 movies, covering approximately 1.3k speech segments.

During the curation process, we begin by manually identifying the speech segments for each character. These segments are then further divided at the sentence level, i.e. to correspond to individual sentences. Specifically, we leverage WhisperX to detect sentence boundaries and localise their corresponding timestamps. The segments are then split at these boundaries. Finally, we manually review and correct any timestamp errors to obtain high-quality sentence-level segments with corresponding speaker labels.

7 Experiments

In this section, we first provide implementation details of our method in Sec. 7.1, followed by a comprehensive evaluation of its effectiveness. We assess performance on both visual and audio character recognition (i.e. character diarisation) in Sec. 7.2 and Sec. 7.3, respectively. Additionally, we evaluate the effectiveness of our approach on downstream tasks, including Audio Description generation (Sec. 7.4) and character-aware subtitling (Sec. 7.5).

7.1 Implementation Details

More details are given in the arXiv version of this paper.

Character Appearance Bank. For each movie, we select the top-10 characters from the cast list and collect 8 exemplar images per character from online sources. Character regions are detected using OWLv2 [37] with the prompt “a photo of animated character” and segmented via SAM 2.1 [47]. Manual correction is performed when web-crawled images are incorrect.

To adapt DINOv2 to animated content, we fine-tune DINOv2 ViT-g/14 (w/ 4 register tokens) [42] by updating its final linear layer. Training lasts for 75 epochs with a cosine annealing learning rate schedule from 6×10^{-4} to 5×10^{-6} , using InfoNCE loss with temperature $\tau = 0.07$.

Character Voice Bank. To build voice exemplars, we collect 5 interview videos per voice actor. In-movie audio segments are selected when the visual match score $s_{vm} > 0.6$ and the sync score $s_{sync} > 0.3$, computed using an adapted version of LWTNet [1]. In our character voice bank, each character has 15 voice segments; if a character has fewer than 15 in-movie samples, we supplement them with segments derived from interview audio.

Table 2: Visual Character Recognition on CMD-AM. “DINOv2 adap.” denotes DINOv2 adaptation. Note, “ \times ”: our bounding box annotations are for entire character bodies (as the facial boundaries of animated characters are often ambiguous), while InsightFace detects face boxes. This mismatch results in low IoUs between the two sources, leading to the reported 0.0 mAP value.

Exp.	Method	DINOv2 adap.	Tracking	Char. Name AP \uparrow	Char. Bbox mAP \uparrow
A	InsightFace [13]	–	–	45.3	0.0*
B	Ours	\times	\times	75.3	30.9
C	Ours	\checkmark	\times	81.8	35.2
D	Ours	\checkmark	\checkmark	80.1	45.3

Visual Character Recognition. We detect shot boundaries using SceneDetect [53], then generate region proposals with OWLv2 and track them using SAM 2.1. Tracks are matched if their IoU is at least 0.5. Frame-level predictions are refined using non-maximum suppression (NMS) with IoU threshold 0.5.

Audio Character Recognition. Audio-only predictions with confidence $c_a < 0.35$ are refined using synchronised visual tracks. We compute a composite score (Sec. 4.2) with $\lambda = 1.0$ and update the predicted label to that of the highest-scoring concurrent track.

Audio Description Generation. Following AutoAD-Zero, we sample 8 frames from the AD interval. Character predictions with $s_{vm} > 0.45$ are retained. Bounding boxes around the characters’ bodies are overlaid to indicate character identity, using unique colours and names in the prompt.

7.2 Visual Animated Character Recognition

We evaluate visual animated character recognition on the CMD-AM dataset from two perspectives: (i) whether the character names in a given clip are correctly predicted, and (ii) whether the character positions are accurately detected. The former is measured using Average Precision (AP) between the predicted and ground truth name lists for 3-shot clips. For character detection, we evaluate the predicted bounding boxes and report the mean Average Precision (mAP), calculated by averaging AP scores at IoU thresholds ranging from 0.50 to 0.95 in increments of 0.05. We use full-body bounding boxes for characters instead of face boxes, because animated characters often have ambiguous facial boundaries, as seen with Lightning McQueen in *Cars*.

We report the quantitative results in Tab. 2, where the baseline InsightFace method is found to struggle with recognising animated characters. In contrast, our visual recognition approach yields significant improvements. In particular, we observe a notable performance boost when adapting DINOv2 on the animated character bank, as seen by comparing Exp. B and Exp. C. Furthermore, incorporating the Track-Guided Region Proposal (TGRP) scheme (Exp. D) leads to comparable character name identification results with per-frame detection (Exp. C), while achieving significantly higher accuracy in grounded bounding box predictions.

Table 3: Audio Character Recognition on CMD-AM. We report Average Precision (AP) for various methods while predicting and evaluating speaker labels using ground truth timestamps.

In-Movie Exemplar	Visual Enhanced	Speaker AP \uparrow
\times	\times	73.8
\checkmark	\times	77.3
\checkmark	\checkmark	77.9

7.3 Audio Animated Character Recognition

We evaluate our method’s performance on audio-only animated character recognition. The evaluation is conducted on 8 animated movies from CMD-AM, each annotated with sentence-level speaker diarisation labels. We report the Average Precision (AP) between the predicted and ground truth speaker labels at the sentence level.

As shown in Tab. 3, augmenting the character voice bank with in-movie exemplars improves recognition accuracy compared to using voice exemplars sourced solely from interview videos. Furthermore, visual-enhanced recognition corrects errors made by the audio-only approach, leading to more accurate speaker identification.

7.4 Audio Description Generation

To assess the quality of generated ADs, we follow previous works [22] and report the performance on CRITIC, CIDEr, and LLM-AD-Eval. Specifically, CRITIC [22] measures the accuracy of character identification in predicted ADs. CIDEr [58] evaluates text description quality based on the TF-IDF [50]-weighted word matching between predicted and GT ADs. LLM-AD-Eval [22] leverages large language models (LLaMA2-7B [56] | LLaMA3-8B [36]) to assess the quality of generated ADs, scoring them from 1 (lowest) to 5 (highest).

We compare our approach against AutoAD-Zero under two configurations: a baseline version without any character label information, and a variant that uses InsightFace combined with our constructed character bank for character recognition. We exclude other training-dependent methods, particularly those explicitly tuned on the CMD-AD [22] training set, since our CMD-AM dataset partially overlaps with it. We evaluate our method on all 75 animated movies in the CMD-AM dataset.

We report the performance of the AD generation framework leveraging two different base VLM models, namely VideoLLaMA2-7B [8] and Qwen2-VL-7B [60]. As shown in Tab. 4, incorporating our character appearance bank and tracking-based character recognition method improves the quality of predicted ADs, especially CRITIC scores, indicating more accurate character references in the generated descriptions.

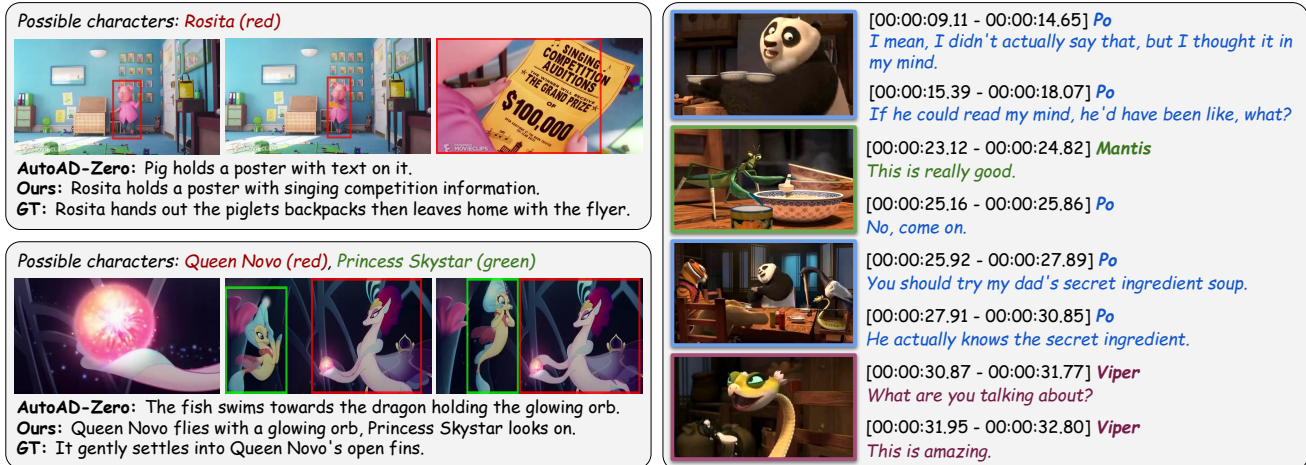
Fig. 5 (left) presents two examples of our predicted ADs. With our method, each AD correctly identifies the character names and exhibits improved description quality compared to AutoAD-Zero. More visualisations are given on the paper website.

7.5 Character-Aware Subtitling

For character-aware subtitling, we measure speech recognition performance using Diarisation Error Rate (DER), which is calculated

Table 4: Audio Description Generation on CMD-AM. “TGRP” is short for Track-Guided Region Proposal, and “DINOv2 adap.” denotes DINOv2 adaptation. The LLM-AD-Eval scores are evaluated using LLaMA2-7B (left) and LLaMA3-8B (right).

Method	Appearance Bank	Character Recognition	VLM	LLM	CRITIC \uparrow	CIDEr \uparrow	LLM-AD-Eval \uparrow
AutoAD-Zero	\times	No Label	VideoLLaMA2-7B [8]	LLaMA3-8B [36]	2.3	9.4	2.64 1.46
AutoAD-Zero	\checkmark	InsightFace	VideoLLaMA2-7B [8]	LLaMA3-8B [36]	15.0	11.3	2.69 1.46
Ours	\checkmark	TGRP + DINOv2 adap.	VideoLLaMA2-7B [8]	LLaMA3-8B [36]	28.6	12.2	2.68 1.43
AutoAD-Zero	\times	No Label	Qwen2-VL-7B [60]	LLaMA3-8B [36]	10.0	13.0	2.91 1.84
AutoAD-Zero	\checkmark	InsightFace	Qwen2-VL-7B [60]	LLaMA3-8B [36]	18.9	14.3	2.94 1.83
Ours	\checkmark	TGRP + DINOv2 adap.	Qwen2-VL-7B [60]	LLaMA3-8B [36]	30.1	15.5	2.96 1.84

**Figure 5: Qualitative Visualisations on CMD-AM.** Left: Predicted Audio Descriptions, where recognised characters are highlighted with coloured bounding boxes, and their names along with corresponding colour codes are provided in the text prompts. The default VLM used is Qwen2-VL-7B. Examples are taken from *Sing* (top) and *My Little Pony: The Movie* (bottom). Right: Predicted character-aware subtitles for a segment from *Kung Fu Panda*.**Table 5: Character-Aware Subtitling on CMD-AM.** Diarisation Error Rates (DER) (in %) are reported. “DER (O)” indicates DER including overlapping speech, whereas “DER” excludes overlapping speech.

Visual Enhanced	DER (O) \downarrow	DER \downarrow
\times	40.3	40.0
\checkmark	35.8	35.7

as the sum of missed speech, false alarms, and speaker confusion errors, divided by the total reference speech time. We report DER both with (DER(O)) and without (DER) overlapping speech. Moreover, we evaluate closed-set performance by excluding segments with ground truth speakers not present in the character bank, as well as segments corresponding to characters singing in *Sing*.

Tab. 5 presents our character-aware subtitling results. By incorporating visual-enhanced speaker recognition, the error rates get further reduced. Fig. 5 (Right) illustrates character-aware subtitles for a segment from *Kung Fu Panda*. Our method produces accurate timestamps and correct character labels, resulting in high-quality subtitles. Further examples are given on the paper website.

8 Discussion

Computer vision systems have been developed over decades for human facial recognition, and are trained on millions of identities/images, while the challenging task of recognising free-form animated characters remains under-explored. Our method achieves accurate audio-visual character recognition with the aid of automatically constructed character banks, enabling character-centric understanding of animated movies for applications such as Audio Description generation and character-aware subtitling.

A limitation of our method is when there are multiple identical animated characters in the same shot. For instance, when there are numerous Minions presenting simultaneously, distinguishing between individual instances becomes challenging. Other than that, the current limitations of our pipeline primarily lie in its downstream tasks. For instance, AutoAD-Zero does not account for temporal context, making it difficult to incorporate story-level plot elements into AD generation. Additionally, when it comes to character-aware subtitling, existing methods often struggle to accurately associate “special” speech segments with the correct speakers, such as those containing singing, crying, or screaming – which are particularly common in animated videos.

Acknowledgements

This research is supported by UK EPSRC Programme Grant VisualAI (EP/T028572/1), a Royal Society Research Professorship RSRP\R\241003, and a Clarendon Scholarship.

References

- [1] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman. 2020. Self-Supervised Learning of Audio-Visual Objects from Video. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- [2] Kiyoharu Aizawa, Azuma Fujimoto, Atsushi Otsubo, Toru Ogawa, Yusuke Matsui, Koki Tsubota, and Hikaru Ikuta. 2020. Building a Manga Dataset “Manga109” With Annotations for Multimedia Applications. *IEEE MultiMedia* (2020).
- [3] Mohamad Alansari, Oussama Abdul Hay, Sajid Javed, Abdulhadi Shoufan, Yahya Zweiri, and Naoufel Werghi. 2023. GhostFaceNets: Lightweight Face Recognition Model From Cheap Operations. *IEEE Access* (2023).
- [4] AudioVault. 2025. AudioVault. <https://audiovault.net/>.
- [5] Max Bain, Jaesung Huh, Tengda Han, and Andrew Zisserman. 2023. WhisperX: Time-Accurate Speech Transcription of Long-Form Audio. In *Proceedings of the 24th Annual Conference of the International Speech Communication Association (INTERSPEECH)*.
- [6] Max Bain, Arsha Nagrani, Andrew Brown, and Andrew Zisserman. 2020. Condensed Movies: Story Based Retrieval with Contextual Embeddings. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*.
- [7] Hervé Bredin. 2023. pyannote.audio 2.1 speaker diarization pipeline: principle, benchmark, and recipe. In *Proceedings of the 24th Annual Conference of the International Speech Communication Association (INTERSPEECH)*.
- [8] Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, and Lidong Bing. 2024. VideoLLaMA 2: Advancing Spatial-Temporal Modeling and Audio Understanding in Video-LLMs. *arXiv preprint arXiv:2406.07476* (2024).
- [9] Peng Chu, Jiang Wang, and Andre Abrantes. 2024. LLM-AD: Large Language Model based Audio Description System. *arXiv preprint arXiv:2405.00983* (2024).
- [10] Joon Son Chung, Jaesung Huh, Arsha Nagrani, Triantafyllos Afouras, and Andrew Zisserman. 2020. Spot the conversation: speaker diarisation in the wild. In *Proceedings of the 21st Annual Conference of the International Speech Communication Association (INTERSPEECH)*.
- [11] Joon Son Chung, Bong-Jin Lee, and Icksang Han. 2019. Who said that?: Audio-visual speaker diarisation of real-world meetings. In *Proceedings of the 20th Annual Conference of the International Speech Communication Association (INTERSPEECH)*.
- [12] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. 2020. RetinaFace: Single-Shot Multi-Level Face Localisation in the Wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*.
- [13] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. 2019. ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*.
- [14] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuyne. 2020. Ecapadtnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification. In *Proceedings of the 21st Annual Conference of the International Speech Communication Association*.
- [15] Arpita Dutta and Samit Biswas. 2019. CNN Based Extraction of Panels/Characters from Bengali Comic Book Page Images. In *International Conference on Document Analysis and Recognition Workshops (ICDARW)*.
- [16] Fandom. 2025. Fandom. <https://www.fandom.com/>.
- [17] Bo Fang, Wenhao Wu, Qiangqiang Wu, Yuxin Song, and Antoni B. Chan. 2024. DistinctAD: Distinctive Audio Description Generation in Contexts. *arXiv preprint arXiv:2411.18180* (2024).
- [18] Yusuke Fujita, Naoyuki Kanda, Shota Horiguchi, Kenji Nagamatsu, and Shinji Watanabe. 2019. End-to-end neural speaker diarization with permutation-free objectives. In *Proceedings of the 20th Annual Conference of the International Speech Communication Association (INTERSPEECH)*.
- [19] Yue Gong and Weifeng Zhong. 2024. A Survey of Faical Detection and Recognition methods for Cartoon Characters. In *Proceedings of the 10th International Conference on Communication and Information Processing (ICCP)*.
- [20] Tengda Han, Max Bain, Arsha Nagrani, Gül Varol, Weidi Xie, and Andrew Zisserman. 2023. AutoAD II: The Sequel – Who, When, and What in Movie Audio Description. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- [21] Tengda Han, Max Bain, Arsha Nagrani, Gül Varol, Weidi Xie, and Andrew Zisserman. 2023. AutoAD: Movie Description in Context. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*.
- [22] Tengda Han, Max Bain, Arsha Nagrani, Gül Varol, Weidi Xie, and Andrew Zisserman. 2024. AutoAD III: The Prequel – Back to the Pixels. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*.
- [23] Ming-Kai He, Jun Du, and Chin-Hui Lee. 2022. End-to-End Audio-Visual Neural Speaker Diarization. In *Proceedings of the 23rd Annual Conference of the International Speech Communication Association (INTERSPEECH)*.
- [24] Mao-Kui He, Jun Du, Shu-Tong Niu, Qing-Feng Liu, and Chin-Hui Lee. 2024. Quality-Aware End-to-End Audio-Visual Neural Speaker Diarization. *arXiv preprint arXiv:2410.22350* (2024).
- [25] Shota Horiguchi, Yusuke Fujita, Shinji Watanabe, Yawen Xue, and Kenji Nagamatsu. 2020. End-to-end speaker diarization for an unknown number of speakers with encoder-decoder based attractors. In *Proceedings of the 21st Annual Conference of the International Speech Communication Association (INTERSPEECH)*.
- [26] Qinghui Huang, Yu Xiong, Anyi Rao, Jiaze Wang, and Dahua Lin. 2020. MovieNet: A Holistic Dataset for Movie Understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- [27] Jaesung Huh and Andrew Zisserman. 2024. Character-aware Audio-visual Subtitling in Context. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*.
- [28] IMDb. 2025. IMDb: Internet Movie Database. <https://www.imdb.com/>.
- [29] Jianhui Jiang, Xiao Qin, Biquan Chen, Molan Tang, Xiao Huang, and Li Miao. 2024. Anime Face Recognition Method Based on Multi-scale Channel Attention. In *Proceedings of the IEEE First International Conference on Data Intelligence and Innovative Application (DIIA)*.
- [30] Hayeon Kim, Eun-Cheol Lee, Yongseok Seo, Dong-Hyuck Im, and In-Kwon Lee. 2021. Character Detection in Animated Movies Using Multi-Style Adaptation and Visual Attention. *IEEE Transactions on Multimedia* (2021).
- [31] Minchul Kim, Anil K. Jain, and Xiaoming Liu. 2022. AdaFace: Quality Adaptive Margin for Face Recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*.
- [32] Bruno Korbar, Jaesung Huh, and Andrew Zisserman. 2024. Look, Listen and Recognise: character-aware audio-visual subtitling. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- [33] Youngki Kwon, Hee Soo Heo, Jaesung Huh, Bong-Jin Lee, and Joon Son Chung. 2021. Look who’s not talking. In *Proceedings of the IEEE Spoken Language Technology Workshop (SLT)*.
- [34] Yong Li, Lingjie Lao, Zhen Cui, Shiguang Shan, and Jian Yang. 2022. Graph Jigsaw Learning for Cartoon Face Recognition. *IEEE Transactions on Image Processing* (2022).
- [35] Kevin Qinghong Lin, Pengchuan Zhang, Difei Gao, Xide Xia, Joya Chen, Ziteng Gao, Jinheng Xie, Xuhong Xiao, and Mike Zheng Shou. 2024. Learning Video Context as Interleaved Multimodal Sequences. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- [36] Meta. 2024. The Llama 3 Herd of Models. *arXiv preprint arXiv: 2407.21783* (2024).
- [37] Matthias Minderer, Alexey Gritsenko, and Neil Houlsby. 2023. Scaling Open-Vocabulary Object Detection. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*.
- [38] Victoria Mingote, Alfonso Ortega, Antonio Miguel, and Eduardo Lleida. 2024. Audio-Visual Speaker Diarization: Current Databases, Approaches and Challenges. *arXiv preprint arXiv:2409.05659* (2024).
- [39] Martinus Grady Naftali, Jason Sebastian Sulistyawan, and Kelvin Julian. 2023. AniWho : A Quick and Accurate Way to Classify Anime Character Faces in Images. *arXiv preprint arXiv:2208.11012* (2023).
- [40] Oron Nir, Gal Rapoport, and Ariel Shamir. 2022. CAST: Character labeling in Animation using Self-supervision by Tracking. *Computer Graphics Forum* (2022).
- [41] Toru Ogawa, Atsushi Otsubo, Rei Narita, Yusuke Matsui, Toshihiko Yamasaki, and Kiyoharu Aizawa. 2018. Object Detection for Comics using Manga109 Annotations. *arXiv preprint arXiv:1803.08670* (2018).
- [42] Maxime Quab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. 2024. DINOv2: Learning Robust Visual Features without Supervision. *Transactions on Machine Learning Research (TMLR)* (2024).
- [43] Xiaoran Qin, Yafeng Zhou, Yonggang Li, Siwei Wang, Yongtao Wang, and Zhi Tang. 2019. Progressive deep feature learning for manga character recognition via unlabeled training data. In *Proceedings of the ACM Turing Celebration Conference - China*.
- [44] Qiyue Qiu, Tong Xu, and Enhong Chen. 2022. Visual-Enhanced End-to-End Neural Diarization. In *Proceedings of the 4th International Conference on Image, Video and Signal Processing (IVSP)*.
- [45] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- [46] Sriji Radhakrishnan, Chao-Han Huck Yang, Sumeer Ahmad Khan, Rohit Kumar, Narsis A. Kiani, David Gomez-Cabrero, and Jesper N. Tegner. 2023. Whispering LLaMA: A Cross-Modal Generative Error Correction Framework for Speech Recognition. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [47] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric

- Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. 2025. SAM 2: Segment Anything in Images and Videos. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- [48] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*.
- [49] Edwin Arkel Rios, Wen-Huang Cheng, and Bo-Cheng Lai. 2021. DAF:re: A Challenging, Crowd-Sourced, Large-Scale, Long-Tailed Dataset For Anime Character Recognition. *arXiv preprint arXiv:2101.08674* (2021).
- [50] Stephen Robertson. 2004. Understanding inverse document frequency: on theoretical arguments for IDF. *Journal of Documentation* (2004).
- [51] Andrew Rouditchenko, Yuan Gong, Samuel Thomas, Leonid Karlinsky, Hilde Kuhne, Rogerio Feris, and James Glass. 2024. Whisper-Flamingo: Integrating Visual Features into Whisper for Audio-Visual Speech Recognition and Translation. In *Proceedings of the 25th Annual Conference of the International Speech Communication Association (INTERSPEECH)*.
- [52] Ragav Sachdeva, Gyungin Shin, and Andrew Zisserman. 2024. Tails Tell Tales: Chapter-Wide Manga Transcriptions with Character Names. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*.
- [53] SceneDetect. 2025. SceneDetect. <https://www.sceneinspect.com/>.
- [54] Krishna Somandepalli, Naveen Kumar, Tanaya Guha, and Shrikanth S. Narayanan. 2018. Unsupervised Discovery of Character Dictionaries in Animation Movies. *IEEE Transactions on Multimedia* (2018).
- [55] Gürkan Soykan, Deniz Yuret, and Tevfik Metin Sezgin. 2023. Identity-Aware Semi-Supervised Learning for Comic Character Re-Identification. *arXiv preprint arXiv:2308.09096* (2023).
- [56] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. *arXiv preprint arXiv: 2307.09288* (2023).
- [57] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation Learning with Contrastive Predictive Coding. *arXiv preprint arXiv:1807.03748* (2018).
- [58] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. CIDER: Consensus-Based Image Description Evaluation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*.
- [59] Hanlin Wang, Zhan Tong, Kecheng Zheng, Yujun Shen, and Limin Wang. 2024. Contextual AD Narration with Interleaved Multimodal Sequence. *arXiv preprint arXiv:2403.12922* (2024).
- [60] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. Qwen2-VL: Enhancing Vision-Language Model's Perception of the World at Any Resolution. *arXiv preprint arXiv:2409.12191* (2024).
- [61] Quan Wang, Carlton Downey, Li Wan, Philip Andrew Mansfield, and Ignacio Lopez Moreno. 2018. Speaker Diarization with LSTM. In *Proceedings of the 43rd IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- [62] Quan Wang, Yiling Huang, Guanlong Zhao, Evan Clark, Wei Xia, and Hank Liao. 2024. DiarizationLM: Speaker Diarization Post-Processing with Large Language Models. In *Proceedings of the 25th Annual Conference of the International Speech Communication Association (INTERSPEECH)*.
- [63] Yan Wang. 2019. Danbooru 2018 Anime Character Recognition Dataset. <https://github.com/grapeot/Danbooru2018AnimeCharacterRecognitionDataset>
- [64] Junyu Xie, Tengda Han, Max Bain, Arsha Nagrani, Eshika Khandelwal, Gül Varol, Weidi Xie, and Andrew Zisserman. 2025. Shot-by-Shot: Film-Grammar-Aware Training-Free Audio Description Generation. *arXiv preprint arXiv:2504.01020* (2025).
- [65] Junyu Xie, Tengda Han, Max Bain, Arsha Nagrani, Gül Varol, Weidi Xie, and Andrew Zisserman. 2024. AutoAD-Zero: A Training-Free Framework for Zero-Shot Audio Description. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*.
- [66] Eric Zhongcong Xu, Zeyang Song, Satoshi Tsutsui, Chao Feng, Mang Ye, and Mike Zheng Shou. 2022. AVA-AVD: Audio-visual Speaker Diarization in the Wild. In *Proceedings of the 30th ACM International Conference on Multimedia (ACM MM)*.
- [67] Fan Yi, Jiaxiang Wu, Minyi Zhao, and Shuigeng Zhou. 2023. Anime Character Identification and Tag Prediction by Multimodality Modeling: Dataset and Model. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*.
- [68] Aonan Zhang, Quan Wang, Zhenyao Zhu, John Paisley, and Chong Wang. 2019. Fully supervised speaker diarization. In *Proceedings of the 44th IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*.
- [69] Bin Zhang, Jian Li, Yabiao Wang, Zhipeng Cui, Yili Xia, Chengjie Wang, Jilin Li, and Feiyue Huang. 2020. ACFD: Asymmetric Cartoon Face Detector. *arXiv preprint arXiv:2007.00899* (2020).
- [70] Chaoyi Zhang, Kevin Lin, Zhengyuan Yang, Jianfeng Wang, Linjie Li, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. 2024. MM-Narrator: Narrating Long-form Videos with Multimodal In-Context Learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*.
- [71] Wenbo Zheng, Lan Yan, Fei-Yue Wang, and Chao Gou. 2020. Learning from the Past: Meta-Continual Learning with Knowledge Embedding for Jointly Sketch, Cartoon, and Caricature Face Recognition. In *Proceedings of the 28th ACM International Conference on Multimedia (ACM MM)*.
- [72] Yi Zheng, Yifan Zhao, Mengyuan Ren, He Yan, Xiangju Lu, Junhui Liu, and Jia Li. 2020. Cartoon Face Recognition: A Benchmark Dataset. In *Proceedings of the 28th ACM International Conference on Multimedia (ACM MM)*.