

# HIERARCHICAL ACTIVITY RECOGNITION AND CAPTIONING FROM LONG-FORM AUDIO

Peng Zhang, Qingyu Luo, Philip J.B. Jackson, Wenwu Wang

Centre for Vision, Speech and Signal Processing (CVSSP), University of Surrey, U.K.

## ABSTRACT

Complex activities in real-world audio unfold over extended durations and exhibit hierarchical structure, yet most prior work focuses on short clips and isolated events. To bridge this gap, we introduce *MultiAct*, a new dataset and benchmark for multi-level structured understanding of human activities from long-form audio. *MultiAct* comprises long-duration kitchen recordings annotated at three semantic levels (activities, sub-activities and events) and paired with fine-grained captions and high-level summaries. We further propose a unified hierarchical model that jointly performs classification, detection, sequence prediction and multi-resolution captioning. Experiments on *MultiAct* establish strong baselines and reveal key challenges in modelling hierarchical and compositional structure of long-form audio. A promising direction for future work is the exploration of methods better suited to capturing the complex, long-range relationships in long-form audio.

**Index Terms**— Long-form audio, hierarchical activity recognition, audio captioning, sound events, *MultiAct*

## 1. INTRODUCTION

While audio understanding has seen increasing attention in recent years, the recognition of complex human activities purely from auditory signals remains largely understudied [1, 2, 3]. Despite recent progress in sound event detection and captioning, most existing datasets [4, 5, 6] and models [7, 8, 9] are constrained to short clips (e.g., 10–30 seconds), which typically capture only isolated acoustic events or brief action snippets. As a result, the understanding of long-form, structured human activities remains largely unexplored. However, in real-world settings, meaningful activities often span extended durations and comprise sequences of temporally related events [10]. Effective modelling of this long-form acoustic structure, particularly procedural relationships among events, is essential for deeper semantic understanding but is not adequately addressed in current research [11, 12].

To advance long-form activity understanding, it is essential to model the structural complexity of human behaviour while addressing the limitations of existing audio datasets [13]. Drawing inspiration from the ecological psychology framework [14, 15], human activity can be analysed across three levels: behavior settings, behavior episodes, and

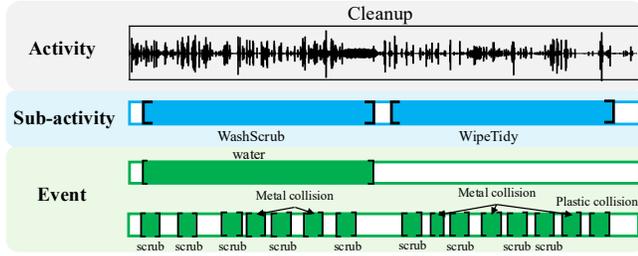
**Table 1.** Comparison of audio datasets by annotation level.

Dataset	Scale	Avg. Dur	Event	Sub-activity	Activity
AudioSet [4]	~5800h	10s	✓	×	×
VGGSound [5]	~550h	10s	✓	×	×
Clotho [6]	~24h	15s–30s	(Captions)	×	×
EPIC-SOUNDS [3]	~100h	~4.9s	✓	×	×
<b>MultiAct</b>	~9h	–	✓	✓	✓

atomic physical units. A behavior setting denotes a stable activity pattern bounded by specific spatial, temporal, and social constraints. In a kitchen scenario, a high-level activity such as `Cleanup` functions as a behavior setting that structures goal-oriented behavior episodes, including sub-activities such as `scrubbing a pan`. These episodes are further composed of fine-grained atomic physical units, manifested in the auditory domain as primitive sound events such as `water running` or `metal collision`. Modelling such a hierarchy is therefore crucial for robust semantic representation and reasoning over complex auditory sequences. However, widely used audio datasets, such as AudioSet [4], VGGSound [5] and Clotho [6], are not well suited for this purpose. As shown in Table 1, these datasets consist of short clips with flat, event-level annotations, lacking the temporal scope, sequential relationships, and hierarchical organisation required to study compositional semantics and high-level activity.

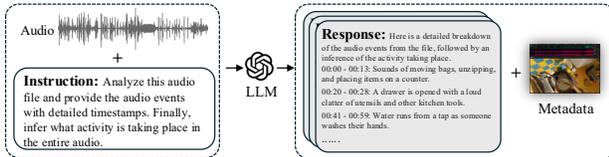
To address this challenge, we introduce the *MultiAct* dataset, a new benchmark for multi-level structured understanding of human activities from long-form audio. *MultiAct* is partially derived from EPIC-SOUNDS [3], where we re-annotate a curated subset and reorganise it into long, contiguous segments with hierarchical temporal labels and paired textual descriptions. The *MultiAct* dataset has three distinguishing features: (1) hierarchical annotations at three semantic levels: high-level activities, sub-activities and events, enabling the study of compositional structure and sequential reasoning; (2) long-form audio recordings of up to 50 minutes, allowing models to capture extended temporal dependencies and activity flow; (3) dual natural-language descriptions for each segment, comprising fine-grained captions and high-level summaries, supporting multi-resolution language generation. To the best of our knowledge, *MultiAct* is the first audio dataset with these features.

Furthermore, we introduce a unified hierarchical frame-



(a) Dataset Overview

Phase 1: LLM-Assisted Event Proposal



Phase 2: LLM-Powered Annotation Refinement



(b) Dataset Annotation Workflow

**Fig. 1.** The hierarchical structure of *MultiAct* and LLM-assisted annotation pipeline.

work for understanding long-form audio activities. We explore tasks such as hierarchical activity recognition, sub-activity sequence prediction and captioning at multiple resolutions. Experiments on the *MultiAct* dataset demonstrate strong performance across tasks and reveal key challenges in modelling structured long-form audio. All resources are available at: [github.com/PennyZhang9/MultiAct](https://github.com/PennyZhang9/MultiAct).

## 2. DATASET

The *MultiAct* dataset is derived from recordings of individuals engaged in cooking activities within a kitchen environment. This dataset is structured in a three-level hierarchy grounded in activity paronomy, providing rich temporal annotations at the levels of activity, sub-activity, and event. As shown in Fig 1(a), each high-level activity (e.g., *Cleanup*) is composed of multiple sub-activities (e.g., *WashScrub*, *WipeTidy*), which in turn consist of a sequence of fine-grained sound events (e.g., *scrub*, *metal collision*, *water*). This hierarchical structure enables the modelling of procedural semantics across different temporal resolutions.

*MultiAct* comprises 8.97 hours of long-form audio from 17 kitchen environments sampled from EPIC-SOUNDS [3], augmented with additional hierarchical annotations. It includes 51 activity instances across 3 coarse activity classes, with durations ranging from 15.7s to 3164s (mean 628s). At the sub-activity level, it contains 472 annotated instances spanning 12 categories, with durations from 0.57s to 572s (mean 63.6s). For event-level modelling, the taxonomy reuses and refines 44 classes, comprising 7312 instances with durations between 0.2s and 100s (mean 4.91s). In addition to temporal annotations, *MultiAct* also includes fine-grained captions aligned with short-term acoustic events

and high-level summaries describing the overall activity, supporting text generation tasks at multiple granularities. Our *MultiAct* annotations are released under the Creative Commons Attribution (CC BY) license. The accompanying audio is sourced from EPIC-SOUNDS [3], a derivative of EPIC-KITCHENS [10], and remains subject to the original non-commercial EPIC-KITCHENS Research License.

Annotations are constructed through a two-phase Large Language Model (LLM)-assisted workflow, as shown in Fig. 1(b). In Phase 1, we prompt GPT-4o [16], with long-form audio content and task-specific instructions to generate a structured annotation draft. Although the output may include hallucinations or boundary errors due to the audio complexity and long duration, it serves as a rich initialisation for downstream refinement. In Phase 2, these drafts are iteratively corrected by human annotators, guided by metadata and a knowledge base of predefined categories and labelling rules. This human-in-the-loop process is particularly critical given the dataset’s multi-level annotation scheme and ensures high-quality, consistent labels across all semantic levels.

## 3. METHOD

As illustrated in Fig. 2, we propose a unified hierarchical framework for long-form audio activity recognition and captioning, consisting of three components: Auditory SlowFast (ASF) [17] feature extractor, multi-level hierarchical activity encoder and unified language decoder for caption and summary generation.

A SlowFast-style encoder with synchronised fast and slow paths extracts temporally aligned acoustic features: the fast path captures transient cues at high resolution; the slow path aggregates longer-range spectral structure. The fused outputs form tokenised, frame-level representations for all downstream tasks.

To model semantic structure at multiple temporal resolutions, we design a hierarchical encoder composed of three task-specific components. The **event encoder** performs fine-grained event classification and boundary detection. The **sub-activity encoder** models mid-level temporal structure via sub-activity classification, boundary detection, and sequence prediction, and fuses its embeddings with frame-level ASF features for activity-level prediction. The **activity encoder** predicts high-level procedural categories. All three encoders operate on a shared audio token representation extracted from the ASF backbone and are trained separately to produce level-specific outputs.

To generate natural-language outputs, we adopt a unified BART-based captioning decoder [18]. Audio features from the ASF encoder are first linearly projected and processed by a multi-layer Transformer audio encoder. Optionally, semantic cues from any level of the hierarchical encoder (e.g., sub-activity sequences or activity labels) are encoded by the pre-trained BART text encoder to provide textual conditioning. The audio and text encoder outputs are then concatenated and fed to the BART decoder, which autoregressively generates

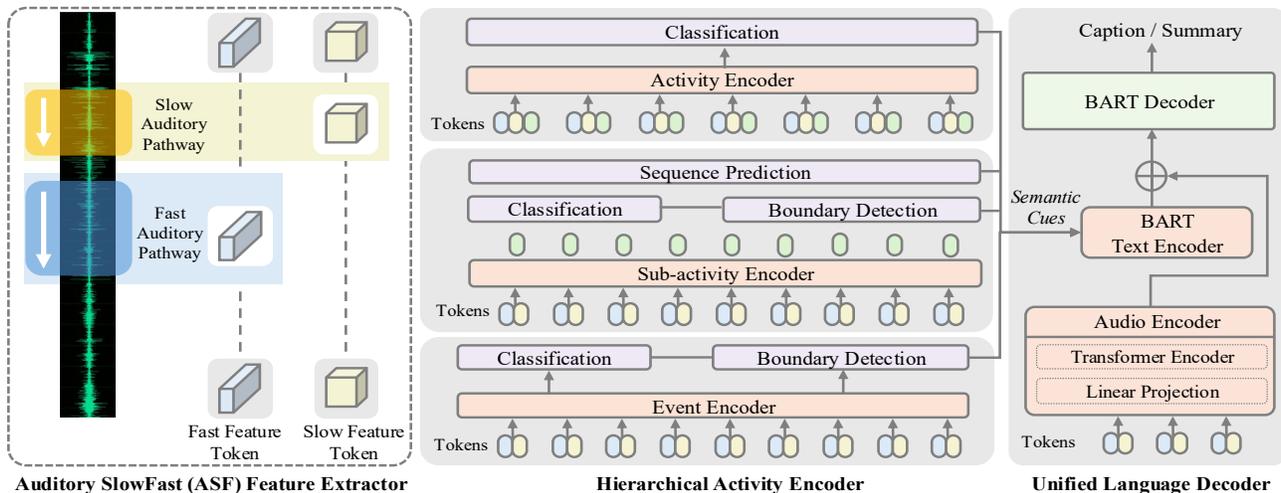


Fig. 2. Overview of the proposed architecture.

either fine-grained captions or high-level summaries, depending on the task.

#### 4. EXPERIMENTS

In this section, we evaluate our approach on four tasks using the MultiAct dataset: (i) **Hierarchical Activity Classification**, predicts one label at each level (activity, sub-activity, event), given a temporally bounded audio segment; (ii) **Audio Activity Detection**, produces class-labelled temporal segments on untrimmed audio; (iii) **Audio Activity Sequence Prediction**, predicts the ordered sub-activity sequence to capture temporal dynamics and intent from a segment; and (iv) **Audio Activity Captioning**, generates either a fine-grained, time-ordered caption or a high-level summary for the long-form audio. In all experiments, the ASF backbone [17] pre-trained on EPIC-SOUNDS [3] serves as a frozen feature extractor, producing 2304-dimensional frame-level audio features that are fed to the task-specific models.

##### 4.1. Hierarchical Activity Classification

The purpose of the activity classification task is to evaluate whether a model offers a comprehensive understanding of activities from auditory information across multiple semantic levels. Given a temporally bounded audio segment (with start and end times), the model predicts one label at each of the activity, sub-activity and event levels. This is formulated as three parallel single-label, multi-class classification tasks.

**Baselines.** We implement three baselines to study architectural choices. Event-level classification is performed by training a linear layer on top of the ASF audio features. For sub-activity and activity classification, we freeze ASF and stack two Bidirectional Gated Recurrent Unit (BiGRU) layers (256 hidden units per direction) on top to encode temporal dynamics, followed by task-specific attention. ASF-Atten applies self-attention pooling [19] to the BiGRU outputs for sub-activity classification. ASF-CrossAtten introduces a cross-attention module [20] and jointly trains sub-activity and activity heads to capture hierarchical dependencies. We report

Table 2. Results of the classification baseline models (% ,  $\uparrow$ ).

Split	Level	Model	Top-1	Top-5	mPCA	mAP	mAUC
Val	Event	ASF	69.0	93.6	30.1	41.6	93.5
		ASF-Atten	50.0	<b>92.6</b>	55.2	<b>64.9</b>	85.5
	Sub-activity	ASF-CrossAtten	<b>51.9</b>	88.9	<b>49.7</b>	65.9	<b>87.5</b>
		ASF-Atten	60.0	N/A	60.7	<b>79.4</b>	87.2
	Activity	ASF-CrossAtten	<b>66.7</b>	N/A	<b>61.9</b>	72.7	84.6
		ASF	67.2	92.4	33.1	41.5	91.9
Eval	Event	ASF	67.2	92.4	33.1	41.5	91.9
		ASF-Atten	<b>51.3</b>	<b>83.3</b>	<b>47.4</b>	<b>41.9</b>	<b>77.5</b>
	Sub-activity	ASF-CrossAtten	42.3	74.4	33.8	37.0	76.4
		ASF-Atten	83.3	N/A	83.3	<b>94.4</b>	<b>95.8</b>
	Activity	ASF-CrossAtten	<b>83.3</b>	N/A	<b>83.3</b>	72.2	70.8
		ASF	67.2	92.4	33.1	41.5	91.9

five standard metrics: Top-1 and Top-5 accuracy, mean Per-Class Accuracy (mPCA), mean Average Precision (mAP), and mean Area Under the Curve (mAUC). Higher values indicate better performance for all metrics. For Activity (3 classes), Top-5 is undefined and therefore omitted (cells marked N/A).

**Results.** Table 2 shows the performance across all three semantic levels on both validation and evaluation splits. ASF-CrossAtten performs best on activity classification (83.3% Eval), while ASF-Atten excels at Sub-activity, benefiting from temporal modelling.

##### 4.2. Audio Activity Detection

In this experiment, we further evaluate the temporal detection task, which aims to detect activity segments along the temporal axis and assign each segment a class label (e.g., sub-activity or event). This is formulated as a multi-class temporal detection problem that jointly predicts segment boundaries (start and end) and the class of each detected segment.

**Baselines.** We adapt ActionFormer [21], a strong video action detection model, to the audio domain. During training, we extract ASF features using 2-s windows with a 200-ms stride, and feed them to ActionFormer for temporal proposal generation and classification. Performance is reported as mAP across multiple Intersection-over-Union (IoU) thresh-

**Table 3.** Results of the detection baseline models (% ,  $\uparrow$ ).

Split	Level	AP at Different IoU Thresholds					Mean
		@0.1	@0.2	@0.3	@0.4	@0.5	
Val	Event	17.0	14.7	12.8	11.2	9.8	13.1
	Sub-activity	44.3	41.0	30.8	25.3	24.3	33.1
Eval	Event	16.5	15.7	15.0	13.6	12.5	14.6
	Sub-activity	41.8	37.3	32.2	26.0	22.0	31.9

**Table 4.** Results of the sequence prediction models for sub-activities (% ,  $\downarrow$ ).

Split	AER on Different Training Context Length					
	2	3	4	6	8	Full
Val	66.7	72.2	75.9	81.5	88.9	79.6
Eval	69.2	74.4	75.6	87.2	87.2	80.8

olds.

**Results.** As shown in Table 3, sub-activity detection consistently outperforms event detection, with higher AP across both splits. However, performance drops notably at higher IoU thresholds. For instance, on the evaluation set, the sub-activity AP falls from 41.8% at an IoU of 0.1 to 22.0% at an IoU of 0.5, suggesting challenges in precise boundary localization.

### 4.3. Audio Activity Sequence Prediction

In audio datasets, activity sequence prediction has been rarely defined or systematically studied in prior work. Leveraging MultiAct’s dense, time-aligned annotations and rich sub-activity taxonomy, we formalise the task as sequence prediction: given a bounded audio segment, predict the ordered sub-activity sequence. The model infers both which sub-activities appear and their order without explicit segment boundaries, emphasising procedural structure and long-range dependencies rather than local boundary detection.

**Baselines.** We use ASF audio features as input to a Conformer-based encoder [22]. The features are first projected with a linear layer. Adopting the convolution kernel size of 31 from the original study, our encoder is configured with 8 layers and processes these features to produce rich contextual representations. A final linear layer maps these representations to label logits, and the model is trained with the Connectionist Temporal Classification (CTC) objective [23] to predict sub-activity sequences. During inference, we apply sliding window decoding to handle long-form audio. We evaluate the Activity Error Rate (AER) as a function of training context length, varying the context from 2, 3, 4, 6, to 8 sub-activities, and also including the entire sequence (“Full”).

**Results.** Table 4 reveals that shorter contexts (2–4) produce the lowest AER, while error increases with longer contexts (6–8) and only partially improves with Full. These results highlight the importance of local context and the difficulty of modelling long-range dependencies, underscoring sequence prediction as a challenging open problem.

**Table 5.** Results of the fine-grained captioning (C) and high-level summarization (S) baseline models (% ,  $\uparrow$ ).

Split	Task	Method	BLEU1	BLEU4	METEOR	ROUGE-L	CIDEr
Val	C	Rule	24.0	4.9	11.4	20.0	3.8
		Hierarchical	<b>39.2</b>	<b>10.8</b>	<b>15.8</b>	<b>28.5</b>	<b>16.2</b>
	S	Rule	23.7	7.99	11.7	26.9	9.3
		Hierarchical	<b>28.4</b>	<b>9.8</b>	<b>13.5</b>	<b>32.5</b>	<b>24.1</b>
Eval	C	Rule	16.8	2.7	<b>13.0</b>	20.2	2.8
		Hierarchical	<b>17.8</b>	<b>3.4</b>	9.6	<b>23.0</b>	<b>20.1</b>
	S	Rule	17.3	<b>7.5</b>	8.8	20.7	2.2
		Hierarchical	<b>21.8</b>	7.3	<b>11.0</b>	<b>28.3</b>	<b>11.1</b>

### 4.4. Audio Activity Captioning

In this experiment, we investigate whether explicitly modelling hierarchical relations among events, sub-activities and activities improves the coherence and semantic fidelity of captions for long-form audio. This task focuses on generating natural language descriptions of long-form audio segments at two levels: fine-grained captioning and high-level summarisation.

**Baselines.** To establish a captioning baseline, we adopt a dual-encoder, BART-decoder [18] architecture. ASF-derived audio features are linearly projected to the BART hidden size and encoded by a Transformer audio encoder (10 layers, 8 attention heads). In parallel, the BART text encoder processes structured cues from multi-level activity encoders. We then concatenate the audio and text encoder outputs and pass them to the BART decoder to generate captions. Training uses teacher forcing with token-level cross-entropy; inference uses beam search (beam size 4) with a length penalty and trigram blocking. We evaluate the performance with BLEU-1/4 [24], METEOR [25], ROUGE-L [26], and CIDEr [27].

**Results.** Table 5 reports the results for fine-grained captioning (C) and high-level summarisation (S) on the validation and evaluation splits. The hierarchical model outperforms the rule-based baseline on most metrics, with the largest gains in CIDEr and ROUGE-L. Summarisation generally achieves higher BLEU, METEOR, and ROUGE than captioning, whereas captioning attains the highest CIDEr on the evaluation split. Overall, performance is higher on summarisation than on fine-grained captioning, underscoring the challenge of generating temporally precise and semantically rich descriptions for long-form audio.

## 5. CONCLUSION

We present MultiAct, a first-of-its-kind benchmark for structured long-form audio, offering hierarchical annotations (activities, sub-activities, events) together with paired fine-grained captions and high-level summaries. Alongside the dataset, we introduce a unified hierarchical model for joint recognition and captioning. Experiments on MultiAct validate the approach and highlight persistent challenges in modelling semantic and procedural structure in long-form audio. We hope this work catalyses sustained progress in high-level activity understanding from long-form audio.

## 6. ACKNOWLEDGMENT

We thank Pablo Martínez-Nuevo, Sven Ewan Shepstone and Jon Francombe (Bang & Olufsen A/S) for valuable discussions. This research was supported by Bang & Olufsen A/S.

## 7. REFERENCES

- [1] J. Shin, N. Hassan, A. S. M. Miah, and S. Nishimura, “A comprehensive methodological survey of human activity recognition across diverse data modalities,” *Sensors*, vol. 25, no. 13, p. 4028, 2025.
- [2] G. Wijngaard, E. Formisano, M. Esposito, and M. Dumontier, “Audio-language datasets of scenes and events: A survey,” *IEEE Access*, 2025.
- [3] J. Huh, J. Chalk, E. Kazakos, D. Damen, and A. Zisserman, “EPIC-SOUNDS: A Large-Scale Dataset of Actions That Sound,” *IEEE TPAMI*, 2025.
- [4] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, “Audio Set: An ontology and human-labeled dataset for audio events,” in *ICASSP*, 2017, pp. 776–780.
- [5] H. Chen, W. Xie, A. Vedaldi, and A. Zisserman, “VGGSound: A large-scale audio-visual dataset,” in *ICASSP*, 2020, pp. 721–725.
- [6] K. Drossos, S. Lipping, and T. Virtanen, “Clotho: an Audio Captioning Dataset,” in *ICASSP*, 2020, pp. 736–740.
- [7] Y. Hou, B. Kang, A. Mitchell, W. Wang, J. Kang, and D. Botteldooren, “Cooperative scene-event modelling for acoustic scene classification,” *IEEE/ACM TASLP*, vol. 32, pp. 68–82, 2023.
- [8] X. Mei, C. Meng, H. Liu, Q. Kong, T. Ko, C. Zhao, M. D. Plumbley, Y. Zou, and W. Wang, “WavCaps: A chatgpt-assisted weakly-labelled audio captioning dataset for audio-language multimodal research,” *IEEE/ACM TASLP*, vol. 32, pp. 3339–3354, 2024.
- [9] J. Hu, Y. Cao, M. Wu, F. Kang, F. Yang, W. Wang, M. D. Plumbley, and J. Yang, “PSELDNets: Pre-trained neural networks on a large-scale synthetic dataset for sound event localization and detection,” *IEEE/ACM TASLP*, 2025.
- [10] D. Damen, H. Doughty, G. M. Farinella, S. Fidler, A. Furnari, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price *et al.*, “The EPIC-KITCHENS Dataset: Collection, Challenges and Baselines,” *IEEE TPAMI*, vol. 43, no. 11, pp. 4125–4141, 2020.
- [11] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, “PANNs: Large-scale pretrained audio neural networks for audio pattern recognition,” *IEEE/ACM TASLP*, vol. 28, pp. 2880–2894, 2020.
- [12] Y. Hou, S. Song, C. Yu, W. Wang, and D. Botteldooren, “Audio event-relational graph representation learning for acoustic scene classification,” *IEEE Signal Processing Letters*, vol. 30, pp. 1382–1386, 2023.
- [13] X. Mei, X. Liu, M. D. Plumbley, and W. Wang, “Automated audio captioning: an overview of recent progress and new challenges,” *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2022, no. 1, p. 26, 2022.
- [14] R. G. Barker and H. F. Wright, “Midwest and its children: The psychological ecology of an american town,” *Row, Peterson and Company*, 1955.
- [15] Z. Luo, W. Xie, S. Kapoor, Y. Liang, M. Cooper, J. C. Niebles, E. Adeli, and F.-F. Li, “MOMA: Multi-Object Multi-Actor Activity Parsing,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 17 939–17 955, 2021.
- [16] A. Hurst, A. Lerer, A. P. Goucher, A. Perelman, A. Ramesh, A. Clark, A. Ostrow, A. Welihinda, A. Hayes, A. Radford *et al.*, “GPT-4o System Card,” *arXiv preprint arXiv:2410.21276*, 2024.
- [17] E. Kazakos, A. Nagrani, A. Zisserman, and D. Damen, “Slow-Fast Auditory Streams for Audio Recognition,” in *ICASSP*, 2021, pp. 855–859.
- [18] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, “BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension,” in *ACL*, 2020, pp. 7871–7880.
- [19] Y. Zhu, T. Ko, D. Snyder, B. Mak, and D. Povey, “Self-attentive speaker embeddings for text-independent speaker verification,” in *Interspeech*, 2018, pp. 2–6.
- [20] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in Neural Information Processing systems*, vol. 30, 2017.
- [21] C.-L. Zhang, J. Wu, and Y. Li, “ActionFormer: Localizing Moments of Actions with Transformers,” in *ECCV*, 2022, pp. 492–510.
- [22] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu *et al.*, “Conformer: Convolution-augmented Transformer for Speech Recognition,” in *Interspeech*, 2020, pp. 5036–5040.
- [23] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks,” in *ICML*, 2006, pp. 369–376.
- [24] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “BLEU: a Method for Automatic Evaluation of Machine Translation,” in *ACL*, 2002, pp. 311–318.
- [25] S. Banerjee and A. Lavie, “METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments,” in *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, 2005, pp. 65–72.
- [26] C.-Y. Lin, “ROUGE: A Package for Automatic Evaluation of Summaries,” in *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, 2004, pp. 74–81.
- [27] R. Vedantam, C. L. Zitnick, and D. Parikh, “CIDER: Consensus-based Image Description Evaluation,” in *CVPR*, 2015, pp. 4566–4575.