

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

6-2023

GroundNLQ @ Ego4D natural language queries challenge 2023

Zhijian HOU

Lei JI

Difei GAO

Wanjun ZHONG

Kun YAN

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

HOU, Zhijian; JI, Lei; GAO, Difei; ZHONG, Wanjun; YAN, Kun; NGO, Chong-wah; CHAN, Wing-Kwong; NGO, Chong-Wah; DUAN, Nan; and SHOU, Mike Zheng. GroundNLQ @ Ego4D natural language queries challenge 2023. (2023). *3rd International Ego4D workshop, Vancouver, 2023 June 19*.

Available at: https://ink.library.smu.edu.sg/sis_research/8416

This Conference Paper is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Author

Zhijian HOU, Lei JI, Difei GAO, Wanjun ZHONG, Kun YAN, Chong-wah NGO, Wing-Kwong CHAN, Chong-Wah NGO, Nan DUAN, and Mike Zheng SHOU

GroundNLQ @ Ego4D Natural Language Queries Challenge 2023

Zhijian Hou^{1*}, Lei Ji², Difei Gao³, Wanjun Zhong⁴, Kun Yan⁵, Chao Li²,
Wing-Kwong Chan¹, Chong-Wah Ngo⁶, Nan Duan², Mike Zheng Shou³

¹ City University of Hong Kong ² Microsoft ³ National University of Singapore
⁴ Sun Yat-sen University ⁵ Beihang University ⁶ Singapore Management University

{zjhou3-c@my., wkchan}@cityu.edu.hk; {leiji, clao nanduan}@microsoft.com;
{daniel.difei.gao, mike.zheng.shou}@gmail.com; zhongwj25@mail2.sysu.edu.cn;
kunyan@buaa.edu.cn; cwngo@smu.edu.sg

Abstract

In this report, we present our champion solution for Ego4D Natural Language Queries (NLQ) Challenge in CVPR 2023. Essentially, to accurately ground in a video, an effective egocentric feature extractor and a powerful grounding model are required. Motivated by this, we leverage a two-stage pre-training strategy to train egocentric feature extractors and the grounding model on video narrations, and further fine-tune the model on annotated data. In addition, we introduce a novel grounding model GroundNLQ, which employs a multi-modal multi-scale grounding module for effective video and text fusion and various temporal intervals, especially for long videos. On the blind test set, GroundNLQ achieves 25.67 and 18.18 for $R1@IoU=0.3$ and $R1@IoU=0.5$, respectively, and surpasses all other teams by a noticeable margin. Our code will be released at <https://github.com/houzhijian/GroundNLQ>.

1. Introduction

The Ego4D [1] NLQ challenge aims to localize a temporal window within a long-form first-person video corresponding to a natural language (NL) question. Existing approaches (summarized in Table 1) primarily explore two research directions: 1) pre-training a representative egocentric feature extractor using the Ego4D video-narration dataset [2, 3], or 2) deploying a powerful grounding model that processes the interaction between video and text features to predict the relevant temporal interval [4, 5, 6].

To develop a discriminative egocentric video representation, EgoVLP [2] and InternVideo [3] have been pre-trained on the Ego4D video-narration dataset, thus becoming go-to egocentric feature extractors. Recently, NaQ [7] proposes

an effective data augmentation strategy to mitigate the data scarcity problem for egocentric video grounding, achieving substantial performance gains. Inspired by this, we employ a two-stage pre-training pipeline that includes both the feature extractor and model pre-training, followed by fine-tuning the grounding model on annotation data.

The long video input presents another challenge for this task. Existing models [8] are primarily designed to handle shorter videos. To address long videos, related works attempted to adapt literature models by sparse sampling [1, 4] or window-slicing [6]. However, these approaches lead to information loss or insufficient global contextual encoding. To input the complete video at once, the model of the Badger team [5] re-purposes ActionFormer [9] for video grounding, which integrates local self-attention to extract a feature pyramid from an input video. However, they adopt a late-stage multi-modal fusion network, which results in ineffective multi-modal interaction. Motivated by this, we introduce GroundNLQ, which incorporates well-designed multi-modal multi-scale modules. Our module integrates the textual query and long video deeply in the early stage and constructs the text-aware video feature pyramid to capture temporal intervals of various lengths.

Through the two-stage pre-training pipeline for both the egocentric video feature and the grounding model, our single model GroundNLQ and final ensemble submission surpasses all other teams by a noticeable margin regarding every evaluation metric (in Table 2). Notably, we also achieve a sizeable performance boost compared with the winner approach of the ECCV22 workshop (i.e., $10.06\% \xrightarrow{+81\%} 18.18\%$ on $R1@0.5$ in Table 1.)

2. Methodology

As depicted in Figure 1, we implement a two-stage pre-training process using Ego4D narrations, followed by fine-tuning the model on the annotated dataset. Additionally, we

* This work was done during the first author’s internship in MSR Asia.

Table 1: Comparison among our approach and previous challenge approaches. Results are reported on the blind test split on the Recall@1 metric.

Team name	Video Feature	Text Feature	Grounding model	Long Video Processing	Model Pre-train	IoU=0.3	IoU=0.5
EgoVLP [2]	EgoVLP	EgoVLP	VSLNet	Sparse Sampling		10.46	6.24
ReLER [4]	SlowFast+Omnivore+CLIP	CLIP	ReLER	Sparse Sampling		12.89	8.14
CONE [6]	EgoVLP	EgoVLP/CLIP	CONE	Window Slicing		15.26	9.24
Badgers@UW. [5]	SlowFast+Omnivore+EgoVLP	CLIP	ActionFormer+AdaAttN	full		15.71	9.57
InternVideo [3]	EgoVLP+InternVideo	EgoVLP	VSLNet	Sparse Sampling		16.46	10.06
NaQ++ [7]	EgoVLP+InternVideo	CLIP	ReLER	Sparse Sampling	✓	21.70	13.46
GroundNLQ	EgoVLP+InternVideo	CLIP	GroundNLQ	full	✓	25.67	18.18

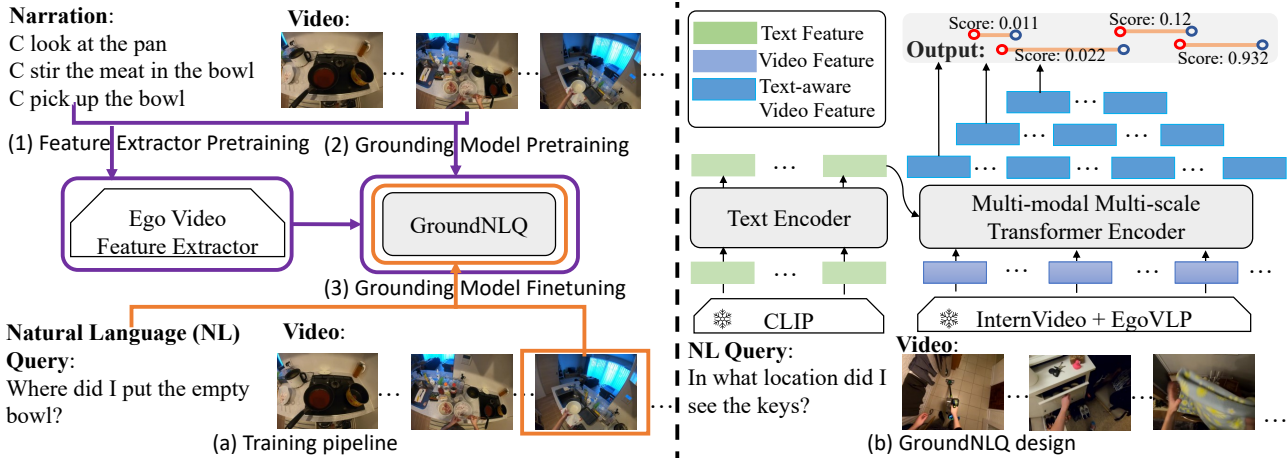


Figure 1: Left (a) depicts the overall training pipeline (§ 2.3). Right (b) shows our grounding model GroundNLQ (§ 2.2).

design a novel grounding model, GroundNLQ. This section elaborates on the specifics of both the features (§ 2.1) and the grounding model (§ 2.2), as well as the implementation details of the training pipeline (§ 2.3).

2.1. Text and Video Representations

For video features, we employ video feature extractors that have been pre-trained on Ego4D narrations, namely, InternVideo [3] and EgoVLP [2]. We concatenate these features along the channel dimension to formulate the final video feature, mirroring the approach used in [3]. Specifically, we extract each video feature for a short snippet (approximately 0.53s) and feed the long video input sequence into our grounding model. In contrast, models such as VSLNet [1, 8] and ReLER [4] employ sparse sampling for video features to generate a fixed-length frame sequence (e.g., 128 features in VSLNet). For textual features, we use the CLIP [10] text encoder to extract the textual token feature sequence of the natural language (NL) query.

Additionally, we conduct feature projection to map the video and text features into the same embedding space. Specially, we use 2 layers of the 1D Convolution network to project video features, and 2 layers of the linear network to project text features, respectively. Each layer is also stacked with layer normalization and ReLU activation. After the

projection, we add the sin-cos position embedding as in [11] into the projected video features to embed temporal-aware position information.

2.2. GroundNLQ Model

Multi-modal Transformer Encoder. For the text encoder, we stack 4 vanilla transformer encoder blocks to learn the contextual token features. Each text transformer encoder block incorporates the multi-head self-attention (MHA) layer and feed-forward network (FFN).

Regarding the video encoder, we use another stack of four transformer encoder blocks to learn text-aware video features. Each video transformer encoder block consists of a local MHA layer, a cross-modal attention layer, and an FFN. We leverage window-based local self-attention to decrease the high computational burden of long video modeling, as demonstrated in previous studies [9, 12]. This approach efficiently handles long sequence inputs while significantly reducing computation costs and maintaining comparable performance. Additionally, we integrate cross-modal attention to infuse text information into the video features, enabling the early-stage fusion of text and video features. This contrasts with the models by the Badgers team [5], which employ late-stage fusion following unimodal feature encoding. This early-stage fusion is critical,

Table 2: Performance comparison of the top three teams on the public leaderboard. For our submission, we report the performances of both validation and blind test splits.

Team-name	Validation				Test Private			
	R1 @0.3	R1 @0.5	R5 @0.3	R5 @0.5	R1 @0.3	R1 @0.5	R5 @0.3	R5 @0.5
ego-env	-	-	-	-	23.28	14.36	27.25	17.58
asL_nlq	-	-	-	-	24.13	15.46	34.37	23.18
GroundNLQ*	26.63	18.04	52.46	40.00	24.02	16.91	40.23	28.67
GroundNLQ	26.98	18.83	53.56	40.00	24.50	17.31	40.46	29.17
Ensemble	27.20	18.91	54.42	39.98	25.67	18.18	42.05	29.80

given the inherent challenge of aligning video content with textual queries; it ensures comprehensive learning and fusion of multi-modal information.

Multi-scale Transformer Encoder. We apply a stack of 6 multi-scale transformer encoder blocks to learn the text-aware video feature pyramid from the text-aware video features. Each block within the multi-scale transformer encoder consists of an MHA layer, a max-pooling layer, and an FFN. We use a max-pooling operator with a stride of 2 to downsample the features, facilitating the capture of longer intervals. The feature pyramid output is a combined result of all 6 layers’ outputs with the inputs. This produces 7 text-aware video feature sequence levels of varying lengths for moment prediction. Contrary to the approach in ReLER [4], which slices the video input sequence into different numbers of clip segments using varying window lengths, our multi-scale mechanism does not employ splitting and is more akin to the feature pyramid network [13].

Prediction Head and Loss Function. Lastly, we implement two layers of a 1D Convolution network for the classification and regression heads, respectively, following a similar approach as [5, 9]. The classification head outputs a probability score for each interval feature of the pyramid, while the regression head outputs the boundary distances from the current interval. The model has dual learning objectives: background/foreground classification and boundary regression. The loss function is a sum of the binary classification loss and IoU regression losses.

2.3. Implementation Details

Egocentric Video Feature Extractor Pre-training. EgoVLP [2] undergoes pre-training with 3.8M paired egocentric video clips and corresponding narrations (i.e., EgoClip [2]) via a CLIP-like contrastive loss. Despite a significantly smaller amount of dataset than the 400M image-text pairs in CLIP or 136M video clip-text pairs in Howto100M, EgoVLP demonstrates the necessity of egocentric data pre-training for egocentric video tasks.

Additionally, InternVideo [3] translates textual narrations into corresponding verb and noun class labels for each EgoClip video clip, utilizing the VideoMAE [14] backbone to train the 1-of-K classification for separate verb and action labels. The result is three video features pre-trained on Ego4D narrations: EgoVLP, InternVideo-Verb, and InternVideo-Noun. We concatenate these features along the channel dimension for the grounding model’s video input.

Grounding Model Pre-training. NaQ [7] presents a data augmentation strategy that converts standard Ego4D video-text narrations into training data for the grounding model. We adhere to NaQ’s data collection steps, assembling all training videos for episode memory benchmarks and collating corresponding <video, narration, moment> tuples. Ground-truth temporal boundaries are initialized using the EgoClip boundary and refined via the temporal response jittering strategy [7]. We then use this pre-training data to train our model with our training loss function. The training utilizes 4 V100 GPUs with a batch size of 4 per GPU, lasting approximately 4 days. The total and warmup epochs are 10 and 4, respectively, with a maximum learning rate of $2e-4$. The best-performing model epoch is chosen based on inference on the validation split in zero-shot mode.

Grounding Model Fine-tuning. The fine-tuning stage involves initializing the prediction head from scratch due to moment boundary discrepancies between pre-training and training data. We initially train our model on the train split, determine the optimal epoch number, and report results on the validation split. For leaderboard submission, we train on the combined train+val splits and apply the model with the optimal epoch number to the private test split. This stage uses 2 V100 GPUs with a batch size of 2 per GPU, taking about 8 hours. The total and warmup epochs are 10 and 4, respectively, with a maximum learning rate of $1e-4$.

Model Ensemble We also explore a model variant called GroundNLQ*, wherein the primary modification is the integration of the cross-modal layer with the multi-scale trans-

former encoder. This new multi-scale block comprises the MHA layer, the max-pooling layer, the cross-modal layer, and the FFN. For leaderboard submission, we ultimately ensemble the predictions from two models (i.e., GroundNLQ and GroundNLQ*).

3. Experiment

3.1. Data Analysis

Table 3: Statistics showing the number of videos and textual narrations/queries in different splits. PT is short for Pre-Train.

Type	Feature PT	Model PT	Train	Validation	Blind Test
Video	3.8M	5,130	1,271	415	333
Text	3.8M	899,219	13,849	4,552	4,005

Table 3 presents the dataset statistics for training the model. For feature pre-training, Ego4D videos are segmented into short clips to learn video representation. For model training, we employ original long videos to train the grounding model.

3.2. Result Analysis

Table 2 displays our primary leaderboard results. Our best leaderboard submission originates from the ensembled model. Significantly, our results outpace all other teams by a substantial margin across all metrics, particularly the R5 metrics. Furthermore, our single model also exhibits robust performance, thereby demonstrating the superior efficacy of GroundNLQ and the staged training pipeline.

3.3. Ablation Analysis

Table 4: Ablations study for InternVideo video feature and model pre-training stage. Results are reported on the validation split.

InternVideo	R1@0.3	R1@0.5	R5@0.3	R5@0.5
✗	21.81	14.28	45.56	32.54
✓	26.98	18.83	53.56	40.00

(a) Effects of the InternVideo video feature.

Pre-train	R1@0.3	R1@0.5	R5@0.3	R5@0.5
✗	16.74	11.47	39.02	27.39
✓	26.98	18.83	53.56	40.00

(b) Effects of the model pre-training stage.

Table 4 provides an ablation study on both features and models. Table 4a underlines the importance of all features. Using the EgoVLP feature alone, the R1@0.3 performance declines from 26.98% to 21.81%. This noticeable gap leads

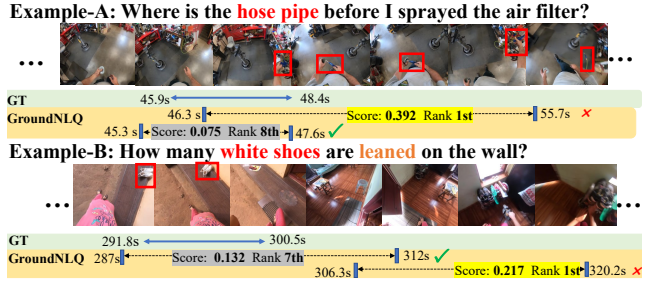


Figure 2: Two failure cases. We show the groundtruth segments and our prediction segments.

to diminished performance in further ensembles. Additionally, we explore the fine-grained ranking in CONE [6]. We employ the top5 prediction of the GroundNLQ model and re-rank these predictions based on the matching score of the EgoVLP model. This also results in a drop in R1@0.3 performance from 26.98% to 20.52%, and the further ensemble does not enhance performance. Table 4b emphasizes the importance of the pre-training stage. As found in NaQ [7], pre-training the grounding model substantially enhances performance.

3.4. Limitation Analysis

Figure 2 shows two failure examples contrasting the groundtruth and predicted segments. In Example-A, the error arises from the imprecise boundary. The first prediction of GroundNLQ captures not only the ground truth event (“I pick up the hose pipe”) but also subsequent events (“I spray the air filter with the hose pipe and then put back the hose pipe”). Consequently, the matched IoU with the ground truth is below the 0.3 threshold, classifying the prediction as false. In Example-B, the error results from an inadequate understanding of the video content and textual query’s nuances. The first prediction of GroundNLQ identifies the event where “I enter the room, take off the shoes and put the shoes in the shoebox near the wall”. While the key objects (shoes and wall) are matched, the alignment neglects detailed attributes of the shoes (their color is white, and they lean against the wall).

4. Conclusion

We present our solution to the Ego4D natural language queries challenge in CVPR 2023. Through our experiments, we highlight the importance of the two-stage pre-training of both the video feature extractor and the grounding model. Regarding the model, we find the importance of early-stage fusion between long video and textual query in the multi-modal multi-scale module. Moreover, we also identify further challenges of the imprecise boundary and fine-grained attribute understanding issues for future work.

References

- [1] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012, 2022.
- [2] Kevin Qinghong Lin, Alex Jinpeng Wang, Mattia Soldan, Michael Wray, Rui Yan, Eric Zhongcong Xu, Difei Gao, Rongcheng Tu, Wenzhe Zhao, Weijie Kong, et al. Egocentric video-language pretraining@ ego4d challenge 2022. *arXiv preprint arXiv:2207.01622*, 2022.
- [3] Guo Chen, Sen Xing, Zhe Chen, Yi Wang, Kunchang Li, Yizhuo Li, Yi Liu, Jiahao Wang, Yin-Dong Zheng, Bingkun Huang, et al. Internvideo-ego4d: A pack of champion solutions to ego4d challenges. *arXiv preprint arXiv:2211.09529*, 2022.
- [4] Naiyuan Liu, Xiaohan Wang, Xiaobo Li, Yi Yang, and Yueting Zhuang. Reler@ zju-alibaba submission to the ego4d natural language queries challenge 2022. *arXiv preprint arXiv:2207.00383*, 2022.
- [5] Sicheng Mo, Fangzhou Mu, and Yin Li. A simple transformer-based model for ego4d natural language queries challenge. *arXiv preprint arXiv:2211.08704*, 2022.
- [6] Zhijian Hou, Wanjun Zhong, Lei Ji, Difei Gao, Kun Yan, Wing-Kwong Chan, Chong-Wah Ngo, Zheng Shou, and Nan Duan. An efficient coarse-to-fine alignment framework@ ego4d natural language queries challenge 2022. *arXiv preprint arXiv:2211.08776*, 2022.
- [7] Santhosh Kumar Ramakrishnan, Ziad Al-Halah, and Kristen Grauman. Naq: Leveraging narrations as queries to supervise episodic memory. *arXiv preprint arXiv:2301.00746*, 2023.
- [8] Hao Zhang, Aixin Sun, Wei Jing, and Joey Tianyi Zhou. Span-based localizing network for natural language video localization. *arXiv preprint arXiv:2004.13931*, 2020.
- [9] Chen-Lin Zhang, Jianxin Wu, and Yin Li. Actionformer: Localizing moments of actions with transformers. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IV*, pages 492–510. Springer, 2022.
- [10] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186, 2019.
- [12] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- [13] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [14] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *arXiv preprint arXiv:2203.12602*, 2022.