

The 1st-place Solution for ECCV 2022 Multiple People Tracking in Group Dance Challenge

Yuang Zhang^{1,2}, Tiancai Wang¹, Weiyao Lin², Xiangyu Zhang¹
¹MEGVII Technology
²Shanghai Jiao Tong University

Abstract

We present our 1st place solution to the Group Dance Multiple People Tracking Challenge. Based on MOTR: End-to-End Multiple-Object Tracking with Transformer [13], we explore: 1) detect queries as anchors, 2) tracking as query denoising, 3) joint training on pseudo video clips generated from CrowdHuman dataset [10], and 4) using the YOLOX [4] detection proposals for the anchor initialization of detect queries. Our method achieves 73.4% HOTA on the DanceTrack test set, surpassing the second-place solution by +6.8% HOTA.

1. Introduction

The DanceTrack [11] dataset with uniform appearance is proposed to re-emphasize motion-based tracking. Predominant approaches [2, 14] for MOT Challenge [5, 9] are mainly benefited from the tracking-by-detection pipeline with strong object detectors, such as YOLOX [4]. It partially results from the crowd scenarios with simple motion in the dataset and MOTA metric [1] bias towards detection performance. As a result, multi-object trackers directly tuned on MOT Challenge fail to well model the motions in complex group dance scenarios.

In contrast, MOTR [13] models tracked instances by “track query” and extends DETR [3, 15] for learnable multi-object tracking by iterative prediction. As indicated in the Table 3 of DanceTrack Paper [11], MOTR achieved the best performance on AssA metric [8], which shows promising potential to associate people in a complex group dance. However, its detection performance is inferior to those YOLOX-based trackers, like OC-SORT [2].

Our goal is to achieve YOLOX-level detection performance and further extend the association capability of MOTR. We use YOLOX [4] for proposal generation and the proposal boxes are used for the detect query initialization for MOTR [13]. With the simple framework illustrated in Figure 1, our method outperforms the second-place solu-




Figure 1. Overview of our solution. It combines state-of-the-art detector YOLOX [4] and online tracker MOTR [13]. YOLOX is employed to generate high-quality object proposals for each frame. The proposals are then used to provide a good initialization for detect queries of MOTR. It means detection process is decoupled from MOTR, which only performs the tracking association.

tion nearly 10% on AssA metric and achieves the second-highest DetA performance (only 0.02% behind the solution with the highest DetA) in the group dance challenge.

2. Method

As shown in Figure 1, the overall framework is based on MOTR [13]. A pretrained YOLOX [4] detector is used to generate high-quality proposals for the temporal association in MOTR. It combines their advantages on detection and association performance.

2.1. Detect Queries as Anchors

Multi-object tracking is a relatively local task since the moving distance of objects is usually small within short time interval. So it will be better to provide more local information for the detect queries employed in MOTR. Following Anchor DETR [12] and DAB-DETR [7], we replace the learnable positional embedding (PE) of detect query in MOTR with the sine-cosine PE of dense anchors, producing an anchor-based MOTR tracker.

2.2. Tracking as Query Denoising

DN-DETR [6] by Li *et al.* introduces query denoising (QD) as a way of stabilizing bipartite matching in DETR [3], simultaneously accelerating convergence and improving object detection performance. Li *et al.* feed the jittered ground-truth bounding boxes into DETR decoder and the decoder is trained to recover the original bounding boxes.

We observe that the query denoising task can be inherently similar to the task of tracking instances across frames: one bounding box prediction of previous frame can be seen as the current frame one with noise. Based on the observation, we adjust the query denoising in [6] so that the scale of random noise added to the ground-truth boxes is comparable to the scale of object movement across frames for multi-object tracking.

Specifically, we choose the noise scale hyperparameters $\lambda_1 = 0.1, \lambda_2 = 0.05$ as opposed to $\lambda_1 = \lambda_2 = 0.4$ in DN-DETR [6]. For each ground-truth bounding box (x, y, w, h) , the center is shifted by a random noise $(\Delta x, \Delta y)$ where $\Delta x \sim \text{Uniform}(-0.05w, 0.05w)$ and $\Delta y \sim \text{Uniform}(-0.05h, 0.05h)$. The jittered width and height are uniformly sampled from $(0.95w, 1.05w) \times (0.95h, 1.05h)$. Note that noise we choose has a much smaller scale compared to DN-DETR, and is more suitable for crowd scenarios and more consistent with the natural bounding box offsets across frames.

2.3. Joint Training on CrowdHuman

To enhance the detection performance, we use the images from CrowdHuman [10] dataset to generate pseudo videos for the joint training of MOTR tracker. Similar to the joint training of MOT17 and CrowdHuman in MOTR [13], we generate pseudo video clips for CrowdHuman and perform the joint training with DanceTrack. The length of pseudo video clip is set to 5 for simplicity.

2.4. Using YOLOX Proposals

To further improve the detection performance, a pretrained YOLOX [4] detector is used to generate the bounding box proposals. To combine it with MOTR in a elegant manner and keep the end-to-end feature of MOTR, we introduce a very simple pipeline. The proposals generated by YOLOX are served as the anchor initialization of queries for the modified anchor-based MOTR tracker. The high-quality proposal provides a good initialization for object localization. The transformer decoder of MOTR predicts the relative offsets compared to the initial localization. With the employment of YOLOX, the detection and association processes of MOTR are fully decoupled and MOTR is only used to perform end-to-end association.




Figure 2. The contribution of each part to HOTA metric on DanceTrack validation set.




Figure 3. The contribution of each component to AssA and DetA metrics on DanceTrack validation set.

3. Experiments

We evaluate the aforementioned parts in our method on the DanceTrack validation set in this section. The improvements on HOTA, AssA and DetA in each part are summarized in Figure 2 and Figure 3.

3.1. Implementation Details

For both training and inference, we generate object bounding boxes with YOLOX for each frame individually. For the DanceTrack dataset, we use the pretrained YOLOX

Table 1. The performance comparison of MOTR official implementation and our baseline on the DanceTrack validation set.

MOTR	HOTA	DetA	AssA
Official repo	51.7	69.0	39.0
Our baseline	57.1	66.2	49.5

model available in the DanceTrack dataset GitHub repo¹. The model weight is also used in OC-SORT [2]. For the CrowdHuman dataset, we use the YOLOX model for testing on MOT17 provided by ByteTrack [14].

For MOTR tracker, our baseline implementation is based on the official repo² with several modifications:

- Object bounding boxes are clipped by the frame boundary after data augmentation.
- The HSV augmentation applied in the training of the YOLOX detector is adopted.
- Both detection and track queries are formulated as anchors. The learnable query PE is replaced with the sine-cosine PE of 4D learnable anchors.
- Track queries are filtered by a score threshold of 0.5 during training, which creates false positive (FP) and false negative (FN) track queries to enhance the handling of FPs and FNs during inference. In this way, we do not manually insert negative or drop positive track queries to create hand-crafted FP and FN, *i.e.*, $p_{drop} = 0$ and $p_{insert} = 0$.

These modifications boost HOTA performance by 5.4%, as shown in Table 1.

We train for 5 epochs with a fixed clip size of 5. The training set of the DanceTrack [11] dataset and the training and validation set of the CrowdHuman [10] are used for training. The initial learning rate 2×10^{-4} is dropped by a factor of 10 at the 4th epoch. If YOLOX is not used, we use 300 learnable anchors with corresponding learnable query embedding as the detect query for MOTR. Otherwise, we use 10 learnable anchors and concatenate them with YOLOX proposals to recall objects missed by the YOLOX detector. For inference, the score threshold for detection is set to 0.5 and the maximum age for negative track queries is set to 20. Other hyperparameters and supervision are the same as the original MOTR.

3.2. Ablation study

CrowdHuman joint training and YOLOX proposal As shown in Table 2, utilizing YOLOX prediction as MOTR detect query initialization *consistently improves all three*

¹<https://github.com/DanceTrack/DanceTrack>

²<https://github.com/megvii-research/MOTR>

Table 2. Ablation study of CrowdHuman joint training and YOLOX proposal on the DanceTrack validation set.

CrowdHuman	YOLOX	HOTA	DetA	AssA
		57.1	66.2	49.5
	✓	60.7	74.8	49.6
✓		56.7	73.7	43.9
✓	✓	63.7	76.6	53.2

Table 3. Ablation study on query denoising on the DanceTrack validation set. The scale of noise λ_1 and λ_2 follows the definition in DN-DETR [6].

λ_1	λ_2	HOTA	DetA	AssA
No QD		63.7	76.6	53.2
0.4	0.4	63.1	77.7	51.5
0.1	0.05	64.5	78.7	53.0

metrics (HOTA, DetA, and AssA) regardless of whether the CrowdHuman dataset is used.

It is worth noticing that with the joint training on the CrowdHuman dataset, using YOLOX for proposal generation significantly improves association performance, as indicated by the nearly 10% improvement in AssA.

While using the CrowdHuman dataset alone does not help improve tracking due to the significant drop in association performance, it further improves both the detection and association with the help of YOLOX.

Query denoising (QD) We add query denoising as an auxiliary task to our method. Experimental results in Table 3 show that using query denoising greatly improves detection performance and further improves HOTA by 0.8%. Using query denoising with the default scale of noise (0.4) hurts association performance. We attribute this to the gap between denoising and tracking tasks, as the artificial noise is much larger than the cross-frame motion of the instances. Our choice of noise range achieves a 2.1% improvement in DetA, while almost retaining AssA performance.

Extra Association in post-processing We find that keeping tracked queries with negative predictions (lost or occluded) for 20 consecutive frames produces the best HOTA result. To associate objects that are occluded for more than 20 frames, we design an extra post-processing step that connects the tracks that exit and reappear within 20 to 100 frames so long as the linkage is unique (*i.e.*, there are no entrance or exit of other objects during that time interval).

For the prediction with validation HOTA 64.5 in the previous subsection, the post-processing step connects a total of 21 tracks in the validation set. This operation links the

Table 4. Ablation study on extra association in post-processing on the DanceTrack validation set.

Extra Association	HOTA	DetA	AssA
	64.5	78.7	53.0
✓	65.1	78.7	54.0

Table 5. The final leaderboard of the group dance challenge. We ranked 1st on the DanceTrack test set. The individual rankings for each metric are listed in parentheses.

Ranking	User	HOTA	DetA	AssA
1	mfv	73.4 (1)	83.7 (2)	64.4 (1)
2	C-BIoU	66.6 (2)	81.3 (13)	54.7 (2)
3	ymzis69	64.6 (3)	82.5 (3)	50.7 (3)

tracks of objects that have been obscured for more than 1 second without affecting the detection performance. The final AssA is improved by 1.0% and HOTA is improved by 0.6% as a result.

4. Final Result

The final result of the group dance challenge is shown in Table 5. Our learning-based solution shows a much stronger association performance (+9.7% AssA and +6.8% HOTA) compared to the matching-based solutions applied by C-BIOU and ymzis69. Note that the model for submission is jointly trained on the training and validation sets of DanceTrack (+1% HOTA), and four models are ensembled together (+ 2% HOTA).

References

- [1] Keni Bernardin and Rainer Stiefelhagen. Evaluating multiple object tracking performance: the clear mot metrics. *EURASIP Journal on Image and Video Processing*, 2008:1–10, 2008. 1
- [2] Jinkun Cao, Xinshuo Weng, Rawal Khirodkar, Jiangmiao Pang, and Kris Kitani. Observation-centric sort: Rethinking sort for robust multi-object tracking. *arXiv preprint arXiv:2203.14360*, 2022. 1, 3
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 1, 2
- [4] Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*, 2021. 1, 2
- [5] Laura Leal-Taixé, Anton Milan, Ian Reid, Stefan Roth, and Konrad Schindler. Motchallenge 2015: Towards a benchmark for multi-target tracking. *arXiv preprint arXiv:1504.01942*, 2015. 1
- [6] Feng Li, Hao Zhang, Shilong Liu, Jian Guo, Lionel M Ni, and Lei Zhang. Dn-detr: Accelerate detr training by introducing query denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13619–13627, 2022. 2, 3
- [7] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. Dab-detr: Dynamic anchor boxes are better queries for detr. *arXiv preprint arXiv:2201.12329*, 2022. 1
- [8] Jonathon Luiten, Aljosa Osep, Patrick Dendorfer, Philip Torr, Andreas Geiger, Laura Leal-Taixé, and Bastian Leibe. Hota: A higher order metric for evaluating multi-object tracking. *IJCV*, 129(2):548–578, 2021. 1
- [9] Anton Milan, Laura Leal-Taixé, Ian Reid, Stefan Roth, and Konrad Schindler. Mot16: A benchmark for multi-object tracking. *arXiv preprint arXiv:1603.00831*, 2016. 1
- [10] Shuai Shao, Zijian Zhao, Boxun Li, Tete Xiao, Gang Yu, Xiangyu Zhang, and Jian Sun. Crowdhuman: A benchmark for detecting human in a crowd. *arXiv preprint arXiv:1805.00123*, 2018. 1, 2, 3
- [11] Peize Sun, Jinkun Cao, Yi Jiang, Zehuan Yuan, Song Bai, Kris Kitani, and Ping Luo. Dancetrack: Multi-object tracking in uniform appearance and diverse motion. *arXiv preprint arXiv:2111.14690*, 2021. 1, 3
- [12] Yingming Wang, Xiangyu Zhang, Tong Yang, and Jian Sun. Anchor detr: Query design for transformer-based detector. *arXiv preprint arXiv:2109.07107*, 2021. 1
- [13] Fangao Zeng, Bin Dong, Yuang Zhang, Tiancai Wang, Xiangyu Zhang, and Yichen Wei. Motr: End-to-end multiple-object tracking with transformer. In *European Conference on Computer Vision (ECCV)*, 2022. 1, 2
- [14] Yifu Zhang, Peize Sun, Yi Jiang, Dongdong Yu, Zehuan Yuan, Ping Luo, Wenyu Liu, and Xinggang Wang. Byte-track: Multi-object tracking by associating every detection box. *arXiv preprint arXiv:2110.06864*, 2021. 1, 3
- [15] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *ICLR*, 2020. 1