

Domain-guided Masked Autoencoders for Unique Player Identification

Bavesh Balaji^{*†}, Jerrin Bright^{*}, Sirisha Rambhatla[†], Yuhao Chen^{*}, Alexander Wong^{*}, John Zelek^{*} and David A Clausi^{*}

^{*} *Vision and Image Processing Lab, University of Waterloo*

[†] *Critical ML Lab, University of Waterloo*

{bbalaji, jerrin.bright, sirisha.rambhatla, yuhao.chen1, alexander.wong, jzelek, dclausi}@uwaterloo.ca

Abstract—Unique player identification is a fundamental module in vision-driven sports analytics. Identifying players from broadcast videos can aid with various downstream tasks such as player assessment, in-game analysis, and broadcast production. However, automatic detection of jersey numbers using deep features is challenging primarily due to: a) motion blur, b) low resolution video feed, and c) occlusions. With their recent success in various vision tasks, masked autoencoders (MAEs) have emerged as a superior alternative to conventional feature extractors. However, most MAEs simply zero-out image patches either randomly or focus on where to mask rather than how to mask. Motivated by human vision, we devise a novel domain-guided masking policy for MAEs termed *d*-MAE to facilitate robust feature extraction in the presence of motion blur for player identification. We further introduce a new spatiotemporal network leveraging our novel *d*-MAE for unique player identification. We conduct experiments on three large-scale sports datasets, including a curated Baseball dataset, the SoccerNet dataset, and an in-house Ice Hockey dataset. We preprocess the datasets using an upgraded Keyframe Identification (KfID) module by focusing on frames containing jersey numbers. Additionally, we propose a keyframe-fusion technique to augment keyframes, preserving spatial and temporal context. Our spatiotemporal network showcases significant improvements, surpassing the current state-of-the-art by 8.58%, 4.29%, and 1.20% in the test set accuracies, respectively. Rigorous ablations highlight the effectiveness of our domain-guided masking approach and the refined KfID module, resulting in performance enhancements of 1.48% and 1.84% respectively, compared to original architectures.

Keywords—Masked Autoencoder; Domain-guided Masking; Motion Blur; Keyframe Identification; Jersey Number Recognition

I. INTRODUCTION

Unique player identification in real-world broadcast videos is a critical challenge that has been extensively researched over the years [1]–[5]. Accurate identification of individual players in sports holds significant importance in various contexts ranging from performance analysis to tactical evaluation of games [6], [7] by coaches and scouts. However, precise player identification is inherently challenging due to the fast-paced nature of most sports. The rapid movements and unpredictable actions of the players often result in motion blur and occlusions, as shown in Fig. 1. Addressing these challenges necessitate the development of adaptive techniques capable of effectively handling motion blur, occlusions, and other common issues encountered in real-world sporting environments.



Figure 1: Example frames of various tracklets from three large-scale sports datasets showcasing the challenges: motion blur, occlusions and low-resolution.

Previous works on jersey number recognition including [3], [4], [8] have primarily focused on extracting spatial features from static images. However, these approaches often rely on conventional convolutional networks [9], which struggle to handle the aforementioned non-idealities. Recent works including [1], [10], leverage the temporal features using Transformers [11] and LSTMs [12] to identify players from video. Despite these advancements, the absence of jersey numbers in numerous frames has been identified to be detrimental to existing models [2]. To tackle this issue, [2] proposed a unique Keyframe Identification (KfID) module to effectively capture keyframes from player tracklets. While KfID effectively isolates frames with clear jersey numbers, its strict thresholding criteria and sole usage of color for frame association leads to the loss of desirable frames, thereby reducing the amount of data. Nevertheless, within the realm of unique player identification, no prior work, to the best of our knowledge, has specifically focused on effectively capturing spatial context from sports video that deals with common problems like motion blur and occlusion.

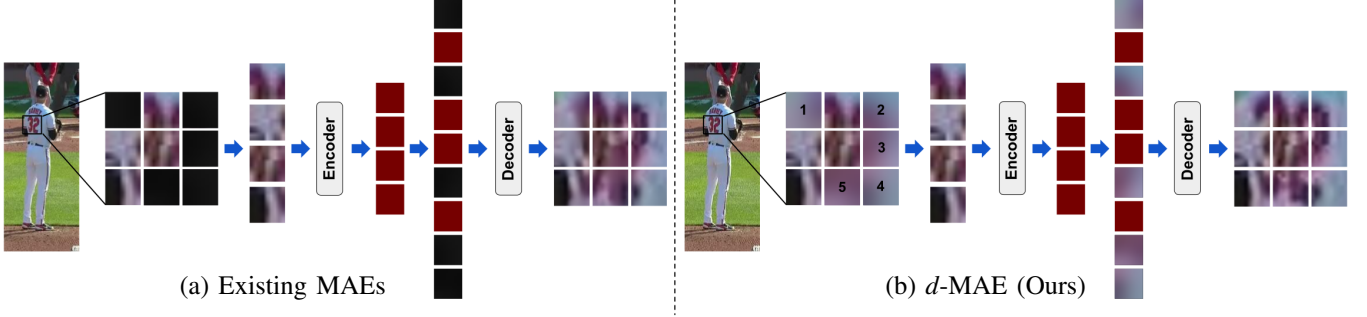


Figure 2: **Comparison of d -MAE with existing MAEs.** (a) Existing MAEs zero-out/ blackout patches randomly while (b) We introduce motion blur artifacts on random patches. The masked patches in (b) are numbered from 1-5.

Recently, masked autoencoders (MAEs), inspired from masked language modeling [13], [14], have emerged as a promising self-supervised learning method for robust spatial feature extraction utilizing vision transformers [15]. MAEs [16] aim to learn semantic representations by zeroing-out random patches and reconstructing the input image using the visible patches. Although MAEs are robust to occlusions, they perform suboptimally in the presence of motion blur as illustrated in Fig. 2(a). This raises critical questions concerning the design of MAEs in learning representations from different *domains* such as sports. *Where should we mask? How should we mask? Is zeroing-out the only way to do it?* While significant research has tackled the first question [17]–[19], the latter two remain underexplored.

To address these underexplored questions for jersey number recognition, we propose a novel spatiotemporal network that utilizes MAEs with a new masking strategy specific to the domain of sports (domain-guided). To improve our MAE’s robustness to motion blur, we introduce motion blur artifacts on random patches instead of zeroing them out. This masking policy proves to be advantageous in extracting effective visual representations, considering the prevalence of motion blur in sports data. Fig. 2 illustrates the different masking strategies employed by the vanilla-MAE and our proposed approach. Doing so, we outperform the state-of-the-art jersey number recognition networks by 8.58%, 4.29% and 1.2% on three large-scale sports datasets. Furthermore, we quantitatively validate our masking strategy against conventional MAEs, demonstrating superior performance. Moreover, we enhance the KfID module proposed in [2] to capture keyframes containing vital information, resulting in a significant improvement of 1.84% compared to the existing module. In summary our contributions include the following:

- 1) We introduce a novel jersey number recognition network that utilizes MAEs coupled with a transformer decoder to capture robust features from low-resolution blurred tracklets.
- 2) We propose a new domain-guided masking strategy, termed d -MAE, specifically tailored to player identi-

fication, enhancing model robustness to motion blur.

- 3) We refine the KfID module by improving its jersey number localization and its ability to capture fine-grained semantic representations of keyframes.
- 4) Addressing the issue of limited data, we introduce a keyframe fusion technique to augment meaningful data, thereby enriching the training process.
- 5) We validate that our model outperforms existing state-of-the-art methods on three large-scale datasets spanning different sports.

II. RELATED WORK

A. Masked Image Modeling (MIM)

Early classical approaches such as classical inpainting [20], [21] and texture synthesis [22], [23] mainly focused on denoising small portions of an image. Hence, these were not very effective in reconstructing very large regions of mask. They also had challenges with filling in objects that were partially occluded. Therefore, to address these issues, Vincent *et al.* [24] use denoising autoencoders to reconstruct corrupted images. Pathak *et al.* [25] introduced context encoders to fill out large holes (rectangular masks) created in an image. Following the success of Masked language modeling using transformers in NLP, a similar approach to predicting missing pixels was then pursued in computer vision. Chen *et al.* [26] use a sequence transformer to autoregressively predict pixels. Other works such as [15], [16], [27] focus on representing an image as discrete tokens and mask them randomly to reconstruct them. However, all of the above methods focus on one masking strategy:- zeroing-out random image patches. More recent works such as [18] explore the use of attention maps, while [19] uses an adversarial mask-generator to learn where to mask. Bandara *et al.* [17] uses a novel token sampling strategy to sample tokens with high spatiotemporal information.

Research Gap in MIM. The above methods on masked autoencoders are focused on the *where to mask* sampling strategy rather than the *how to mask* strategy. We explore into this unexplored realm by developing a domain-guided

masking strategy specifically designed to address the challenge of highly prevalent motion blur in sports videos.

B. Jersey Number Recognition from Static Images

The advent of deep learning facilitated the use of jersey numbers rather than player appearances for unique player identification. Earlier works [3], [4], [28] use CNNs to predict the jersey number: Gerke *et al.* [28] recognize jersey numbers directly from soccer images while Li *et al.* [4] and Liu *et al.* [3] propose a unified solution to detect and classify jersey numbers using Spatial Transformer Networks (STNs) and pose-guided recurrent CNNs, respectively. Vats *et al.* [8] leverage a multi-task loss function to perform holistic and digit-wise predictions. Bhargavi *et al.* [29] present a multi-stage network that takes advantage of pose to localize jersey numbers before detecting them using a secondary classifier.

C. Jersey Number Recognition from Player Tracklets

More recent works [1], [5], [10] aim at leveraging the temporal aspect and capture temporal cues from tracklets. Vats *et al.* [1] developed a transformer-based architecture to recognize jersey numbers from Ice Hockey videos. Chan *et al.* [10], on the other hand, utilize LSTMs to extract temporal characteristics from player tracklets. Furthermore, they also employ 1D CNNs as a late score-level fusion method for classification. Liu *et al.* [5] adopt a two-stage approach to perform the detection and classification of players from American football videos. Balaji *et al.* [2] propose a novel KfID module to remove redundant frames that contain no essential information about the jersey number.

Research Gap. While previous methods on player identification from tracklets prioritize enhancing the temporal representation of their networks for accurate jersey number recognition, our focus lies on improving spatial feature extraction, which contains richer information and enhances our temporal decoder. To achieve this, we leverage MAEs tailored for sports data. The proposed system depicted in Fig. 3 is designed to handle motion blur and occlusions, thereby aiding in determining player identities in a more reliable manner.

III. METHODOLOGY

The overview of the proposed architecture is shown in Fig. 3. Initially, the frames of a player tracklet are processed in the KfID module, where keyframes crucial for identifying jersey numbers are extracted. Subsequently, these keyframes are passed to our d -MAE encoder, which captures features with rich semantic representations of each keyframe. These spatial features are then passed to the temporal transformer decoder, which extracts temporal cues and predicts the jersey number associated with the tracklet. More details on the individual modules are explained in this section.

A. Domain-guided Masked Autoencoders

The application of MAE becomes particularly significant for jersey number detection due to the dynamic nature of the game. Players are frequently in motion leading to challenges such as occlusion, low-visibility and motion blur. However, we recognize that the conventional masking policy (zeroing-out) used in MAEs do not encompass the diverse conditions encountered in real-world scenarios. For instance, viewing an image through a window filled with water droplets provides a form of masking that is not similar to zeroing-out image patches. Similarly, broadcast feeds of fast-paced sports like Soccer, Basketball, or Ice Hockey highlight the occurrence of such visual distortions and blurring effects that affect visual clarity.

Motivated by the need to recover missing or occluded spatial information to acknowledge these diverse scenarios, we build on the proficiency of MAE to reconstruct missing patches within the pixel space by introducing a domain-guided masking strategy. Particularly, during the pre-training stage, we incorporate motion blur to the patches instead of simply zeroing-out them, thereby infusing domain knowledge in the process. This approach facilitates reliable and accurate prediction of jersey numbers in dynamic sports scenarios. We incorporate an additional supervision to the pre-training objective, improving the feature extraction process.

Pre-training. The input image $\mathbf{I} \in \mathbb{R}^{H \times W \times C}$ is split into K patches $\mathbb{I} = \{\mathbf{I}_k \in \mathbb{R}^{P^2 \times D}\}_{k=1}^K$ where P is the patch size, D is the embedding dimension and $K = HW/P^2$. A random subset of the patches $\mathbb{S} \subseteq \mathbb{I}$ are then masked by introducing motion blur artifacts to the pixels, resulting in the set

$$\mathbb{I}_{\text{masked}} = \{m(\mathbf{x}) : \mathbf{x} \in \mathbb{S}\},$$

where $m : \mathbb{R}^{P^2 \times C} \rightarrow \mathbb{R}^{P^2 \times C}$ is the mask applied to random patches. The unmasked patches $\mathbb{I}_{\text{unmasked}} = \mathbb{I} \setminus \mathbb{S}$ are then converted to unmasked tokens

$$[\text{unmask}] = \text{concat}(\mathbf{x} : \mathbf{x} \in \mathbb{I}_{\text{unmasked}})$$

and passed to the MAE encoder to extract latent spatial features

$$\mathcal{F}_s = f([\text{unmask}]),$$

where $f : \mathbb{R}^{K \times P^2 \times D} \rightarrow \mathbb{R}^{K \times P^2 \times D}$ denotes the MAE encoder. The masked tokens

$$[\text{mask}] = \text{concat}(\mathbf{x} : \mathbf{x} \in \mathbb{I}_{\text{masked}})$$

are then used along with the latent features to generate the reconstructed image

$$\hat{\mathbf{I}} = g(\text{concat}([\text{mask}], \mathcal{F}_s)),$$

where $g : \mathbb{R}^{K \times P^2 \times D} \rightarrow \mathbb{R}^{H \times W \times C}$ denotes the MAE decoder.

To induce motion blur in the selected patches, we employ an oriented motion blur filter \mathbf{K}' characterized by two

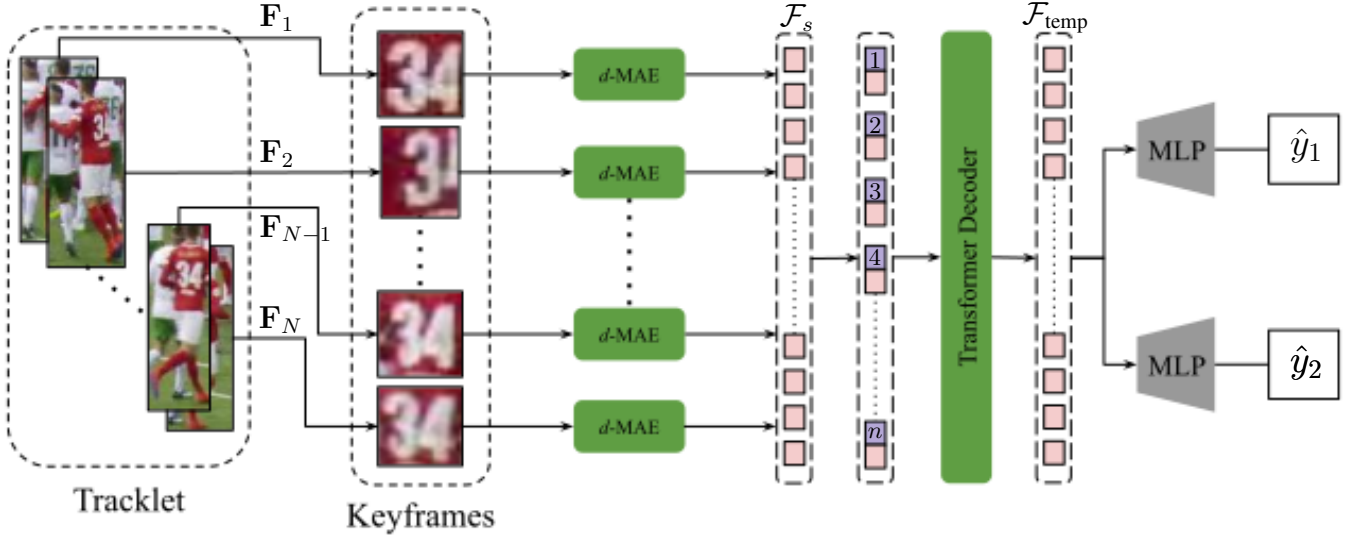


Figure 3: **Overall architecture.** Given a tracklet \mathbb{T} consisting of N frames, we pass \mathbb{T} through the KfID module to extract $n \leq N$ keyframes that contain the jersey number. Each keyframe is passed as an input to our d -MAE encoder to extract spatial features \mathcal{F}_s . These features are then fed to the temporal transformer decoder to extract temporal features $\mathcal{F}_{\text{temp}}$. Two classification heads are utilized to compute the predicted digits of the jersey number \hat{y}_1 and \hat{y}_2 respectively.

parameters: the angle of rotation (ω) and the scale factor (s_f). The filter is centered at $(k_s/2, k_s/2)$ where k_s is the kernel size. Eq. (1) denotes the motion blur filter used to apply motion blur on image patches.

$$m(\mathbf{I}_k)(x, y) = \frac{\mathbf{K}'}{\sum_{i,j} \mathbf{K}'_{ij}} * \mathbf{I}_k(x, y) \quad (1)$$

$$\mathbf{R} = \begin{bmatrix} s_f \cos(\omega) & -s_f \sin(\omega) & \frac{k_s}{2} (1 - s_f \cos(\omega)) + \frac{k_s}{2} s_f \sin(\omega) \\ s_f \sin(\omega) & s_f \cos(\omega) & \frac{k_s}{2} (1 - s_f \cos(\omega)) - \frac{k_s}{2} s_f \sin(\omega) \end{bmatrix} \quad (2)$$

where $*$ denotes the convolution operation and $m(\cdot)$ is the masking strategy we employ at every pixel position (x, y) of an image patch \mathbf{I}_k . The rotation matrix \mathbf{R} used to generate the oriented filter \mathbf{K}' is shown in Eq. (2).

This tailored approach facilitates our d -MAE in capturing crucial cues necessary for the accurate reconstruction of the keyframes in the presence of motion blur.

B. Transformer Decoder

To capture the temporal cues within the tracklet, we extend our MAE module, by introducing a transformer decoder. Specifically, after the pre-training stage, we discard the decoder of d -MAE during finetuning, and pass the original unmasked keyframes directly to the d -MAE encoder. The extracted latent spatial features \mathcal{F}_s are fed to the temporal transformer decoder to perform jersey number recognition. Leveraging the standard vision transformer (ViT) architecture for our decoder, we utilize self-attention to capture long-range dependencies between the spatial features of different

frames within a player tracklet. By employing the self-attention mechanism on the spatial tokens, we facilitate the model's ability to understand the global context and intricate connections between keyframes of a tracklet. The resulting representation $\mathcal{F}_{\text{temp}}$, encapsulates rich cues on the jersey number, which are crucial for player identification.

C. Keyframe Identification

The Keyframe Identification (KfID) module was proposed in [2] with the objective to capture critical frames of a player tracklet where jersey numbers are visible. It consists of three critical components – 1) *Jersey Number Localization (JNL)* localizes all digits in a particular frame \mathbf{F}_i ; 2) *RoI-based filtering* captures digits of our player of interest by filtering all the digits that are within a preset region of interest (RoI); 3) *Local Histogram Correlation (LHC)* creates the holistic representation of the jersey number of our player of interest by merging digits detected in a frame \mathbf{F}_i ; and 4) *Global Histogram Correlation (GHC)* clusters different frames using their spatial color (HSV) layout to find keyframes that contain the jersey number of our player of interest.

In mathematical terms, given a player tracklet $\mathbb{T} = \{\mathbf{F}_i : \mathbf{F}_i \in \mathbb{R}^{H \times W \times 3}\}_{i=1}^N$ consisting of N frames,

$$\text{KfID}(\mathbb{T}) = \mathbb{T} \setminus \{\mathbf{F}_{n_1}, \mathbf{F}_{n_2}, \dots, \mathbf{F}_{n_k}\} \quad (3)$$

where $\{\mathbf{F}_{n_1}, \mathbf{F}_{n_2}, \dots, \mathbf{F}_{n_k}\}$ denotes the set of noisy frames that need to be removed as they contain no relevant information regarding the jersey number and affect the performance of the models.

While the KfID module contributes to a significant improvement over the existing frameworks for jersey number identification by filtering out undesirable frames from the tracklet, we observe that its JNL module fails to recognize smaller digits and provides false positives in noisy scenarios. Furthermore, the GHC module, relying solely on color for frame association, tends to produce spurious clusters. To address these issues and further enhance the KfID module, we introduce architectural modifications to the JNL and GHC modules. Additionally, we incorporate a keyframe fusion-based data augmentation to address challenges with limited labelled data.

JNL Module. The vanilla KfID module utilizes a finetuned YOLOv5 [30] model to detect digits from the input frames of player tracklets. However, it tends to perform poorly in noisy and crowded scenarios, particularly when dealing with the detection of smaller digits. To address these issues, we use a pretrained deformable-DETR [31] model and finetune it on our digit detection dataset captured from SoccerNet images for reliable digit detection. We chose the deformable-DETR network because of its ability to accommodate high-resolution feature maps in a transformer network, enabling capturing of fine-grained representations of smaller objects.

GHC Module. Balaji *et al.* [2] utilize the spatial color layout of the holistic detections from different frames of a tracklet and employ clustering techniques to associate frames containing the jersey number of the player of interest. However, when multiple players from the same team appear in a tracklet, this approach fails since it solely leverages color to cluster frames irrespective of their jersey numbers. To overcome this challenge, we propose extracting deep semantic features such as shape and texture of each frame using a ResNet backbone. These features are then clustered to accurately identify the keyframes within a player tracklet.

Keyframe Fusion. To tackle the challenge of limited labeled data from the incorporation of the KfID module, we introduce a strategic *fusion augmentation* technique to counter the shortcomings. Specifically, we implement this augmentation by randomly selecting n consecutive frames from a specific sequence within a tracklet and merging them together. This fusion captures rich visual representation of the digits especially in scenarios with noise or fast-paced movements in the keyframes. By fusing all the frames within a sequence of a tracklet, we preserve the temporal flow of information. This ensures that the fusion augmentation doesn't introduce additional noise by fusing frames that are too distant to each other.

D. Loss Functions

Siamese Loss. Inspired by the existing literature from 3D vision [32], [33], we employ an additional Siamese objective apart from the MSE loss to supervise the reconstructed image from the MAE decoder while pre-training it. Given a

prediction ($\hat{\mathbf{I}}$) and groundtruth (\mathbf{I}) image, we extract features from both the images using a pretrained feature extractor $h(\cdot)$ and minimize the discrepancy using ℓ_1 -loss.

$$\mathcal{L}_{\text{siamese}} = \|h(\hat{\mathbf{I}}) - h(\mathbf{I})\|_1 \quad (4)$$

We use ResNet as $h(\cdot)$ to extract features from the images and incorporate ℓ_1 -loss instead of the MSE metric used in [33]. Ablations on the different setups are detailed in Table V.

The total loss of our MAE network in pre-training stage is as follows:

$$\mathcal{L}_{\text{mae}} = \sigma_1 \|\hat{\mathbf{I}} - \mathbf{I}\|_2 + \sigma_2 \mathcal{L}_{\text{siamese}} \quad (5)$$

where σ_1 and σ_2 are learnable weights.

Multi-head Classification Loss. We employ a multi-task loss function using cross-entropy to effectively classify jersey numbers [2] from the transformer decoder, as shown in Eq. (6), where $y_1 \in \mathbb{R}^{11}$ and $y_2 \in \mathbb{R}^{11}$ are ground-truth digits of the jersey number (10 digits + 1 null class), and $\hat{y}_1 \in \mathbb{R}^{11}$ and $\hat{y}_2 \in \mathbb{R}^{11}$ are the predictions made by the spatiotemporal network.

$$\mathcal{L}_{\text{class}} = - \sum_{i=0}^{10} y_1^i \log \hat{y}_1^i - \sum_{j=0}^{10} y_2^j \log \hat{y}_2^j \quad (6)$$

IV. EXPERIMENTS

A. Datasets

We evaluate the performance of our model, along with the state-of-the-art methods, on three large-scale player tracklet datasets comprising videos from sports including Ice Hockey, Baseball, and Soccer. The dataset splitup is outlined in Table I. Example tracklets for these datasets is illustrated in Fig. 1.

SoccerNet Dataset. The SoccerNet player recognition dataset is the largest open source jersey number dataset in the world, comprising a total of 4,064 player tracklets. Each tracklet is predominantly focused on one player, and the label for the entire tracklet is the jersey number of that particular player. To facilitate model evaluation and training, the dataset has been partitioned into four distinct subsets: training, validation, testing, and challenge sets. The test and challenge sets are two different test sets sampled from different distributions of games. This helps us in evaluating the generalizability of models and their robustness to slight distribution shifts.

Baseball Dataset. We have curated a player identification dataset from Baseball videos, which is built based on the Baseball 3D pose dataset introduced in [34]. The dataset comprises of 150 player tracklets sampled from over 1000 videos. Here, we utilize the videos from the aforementioned dataset and assign jersey number labels to each tracklet.

Table I: Dataset split-up for training, validation and testing.

Dataset	SoccerNet Dataset			Ice Hockey Dataset			Baseball Dataset		
	Tracklets	Images	Keyframes	Tracklets	Images	Keyframes	Tracklets	Images	Keyframes
Train	1,141	587,543	71,021	2,829	540,339	162,101	105	21,050	18,344
Validation	286	146,886	19,609	176	33,616	11,084	15	2,962	2,571
Test	1,211	565,758	70,445	505	96,455	28,937	30	5,988	4,640
Challenge	1,426	750,092	101,307	-	-	-	-	-	-
Total	4,064	2,052,306	262,382	3,510	670,410	202,122	150	30,000	25,555

Table II: Quantitative comparison of our model with the state-of-the-art on the three datasets.

Method	SoccerNet		Ice Hockey	Baseball
	Test Acc	Challenge Acc	Test Acc	Test Acc
Gerke et al [28]	32.57	35.79	61.20	64.47
Vats et al [8]	46.73	49.88	83.17	87.61
Li et al [4]	47.85	50.60	81.15	88.29
Vats et al [1]	52.91	58.45	85.14	89.46
Balaji et al [2]	68.53	73.77	92.50	93.68
Ours	77.31 $\uparrow 8.58$	81.92 $\uparrow 8.15$	96.79 $\uparrow 4.29$	94.70 $\uparrow 1.02$

Ice Hockey Dataset. We utilize the dataset introduced in [1] to evaluate our model’s performance on a fast-paced game with high motion blur and heavy equipment. The dataset consists of 3,510 player tracklets generated from 84 NHL videos. The average length of a player tracklet is 191 frames, sampled at 30 fps.

B. Implementation Details

Data Augmentation. We follow the data augmentation pipelines outlined in [2] for the SoccerNet dataset and [1] for the Ice Hockey dataset. For the Baseball dataset, our main augmentation strategies include color jitter and random rotation within ± 10 degrees. Subsequently, all images are resized to 224×224 and normalized using the ImageNet mean and standard deviation.

Model Settings. We use the ViT-B variant of MAE for spatial feature extraction. Due to computational constraints, instead of pre-training it from scratch, we utilize a pretrained version of MAE that leverages the zeroing-out policy and further pre-train it on an in-house static jersey number dataset using our masking strategy. We follow the training pipeline outlined in the paper [16] with modifications to incorporate the masking strategy to the MAE. The parameters ω and s_f of the motion blur filter are empirically selected upon experimentation from a range of $0^\circ - 90^\circ$ and $0 - 2.0$ respectively. For the temporal transformer decoder, we use the standard ViT [15] with 8 attention heads and 4 transformer layers for ideal performance.

Training Details. We trained our model for 20,000 iterations on the SoccerNet and Ice Hockey dataset, and 10,000 iterations on the Baseball dataset with a batch size of 16. The AdamW optimizer was employed with an initial learning rate of $3e-4$, with a learning rate scheduler that decreased the learning rate after every 2000 iterations for the initial

6000 iterations. All experiments were carried out using a single NVIDIA 2080Ti GPU with 12GB RAM.

C. Results

Comparison with SOTA. We conduct extensive experiments on the aforementioned datasets to assess the efficacy of our proposed architecture compared to existing state-of-the-art models on jersey number recognition. Table II outlines our model’s performance in comparison with existing works. The data processed by the KfID module serves as input to all the existing jersey number recognition networks for fair evaluation. Table II demonstrates that our model consistently outperforms the existing techniques on all the three datasets. This demonstrates our model’s ability to capture clearer spatial features, in dynamic sports scenarios characterized by challenges such as low-resolution, motion blur and occlusion.

Qualitative Results. We present the qualitative results of our model on different tracklets from the datasets in Fig. 4. To obtain predictions for each image individually, we pass the same image feature S times to our transformer decoder, where S represents the number of tokens. For the overall tracklet prediction (TP), we pass the features of all the images within a tracklet to the transformer decoder. The results demonstrate that our model consistently predicts jersey numbers reliably, even in extremely blurry scenarios. However, it is clear that our model yields suboptimal results when the jersey number is not entirely visible in a majority of the frames, like example 4 in Fig. 4.

D. Ablation Study

Spatial Backbone. To understand the effectiveness of MAEs and the domain-guided masking strategy, we conduct an experiment with different backbones (conventional



Figure 4: **Qualitative results.** Performance of our model on five different player tracklets from all the three datasets. We find our model’s prediction for each image separately and for the entire tracklet (**Pred**). **GT** represents the ground-truth value for the entire tracklet.

backbones and different MAEs) and masking strategies (for the MAEs) on the SoccerNet dataset as shown in Table VI. The results illustrate that leveraging MAEs as spatial feature extractors results in a 12.2% increase compared to conventional backbones. This performance improvement is mainly credited to the self-supervised objective used in its pre-training stage. Furthermore, the observed accuracy boost of 1.84% with the use of motion blur as the masking strategy validates our hypothesis that a domain-guided masking strategy enhances robustness to real-world broadcast data.

Keyframe Identification Module. We evaluate the performance of our model on all 3 datasets with and without the

processing from the KfID module in Table III. Additionally, we quantitatively demonstrate the impact of our modifications in the JNL and GHC components on the overall improvement of the model performance in Table IV.

The results in Table III underscore the importance of the KfID module, as its incorporation leads to an increase of 35.08%, 41.66% and 5.73% in the test accuracies of Ice Hockey, SoccerNet and Baseball datasets respectively. This shows the amount of spurious frames in real-world data and necessitates the need to identify useful detections.

Table IV illustrates the efficacy of our modified KfID as it improves our overall model accuracy by 1.84%. Further

Table III: Results with and without KfID Module. (†) - with the KfID module.

Dataset	Test Acc	Challenge Acc
Ice Hockey	61.71	-
Baseball	88.43	-
SoccerNet	35.65	35.98
Ice Hockey (†)	96.79 ↑35.08	-
Baseball (†)	94.70 ↑5.73	-
SoccerNet (†)	77.31 ↑41.66	81.92 ↑45.94

Table IV: Effect of different proposed components of the KfID module on the overall accuracy in the SoccerNet dataset.

Module	GHC old	GHC new(ours)
JNL old	75.47	76.54
JNL new(ours)	75.92	77.31

examining the findings presented in Table IV provides valuable insights into the significance of each implemented design change within the KfID module. Notably, the integration of our GHC module yields a 1.39% enhancement in overall accuracy, underscoring the limitations associated with relying solely on color for clustering frames and emphasizing the need for deep features. Additionally, the incorporation of our JNL module demonstrates a 0.77% improvement over the original JNL module in accuracy on the SoccerNet dataset. This underscores the critical role of robustly detecting smaller digits, thereby contributing to the refinement of our model’s performance.

Ablations on Siamese Loss. We explore and evaluate different feature extractors for the Siamese loss $\mathcal{L}_{\text{siamese}}$ to assess their influence on the loss function and determine the optimal feature extractor. Furthermore, we compare different similarity and distance metrics to determine the most effective distance metric, which leads to quicker convergence and superior results. These experiments were conducted on the SoccerNet dataset, and the results are depicted in Table V.

Table V illustrates the efficacy of ℓ_1 -loss over other similarity metrics such as the cosine similarity and ℓ_2 norm. The lower performance of cosine similarity can be attributed to its focus on the angle between two vectors rather than their magnitudes. This characteristic is less advantageous for image reconstruction tasks, where the magnitude of each feature holds more importance than the angle between them. Among the other two metrics, we hypothesise that the superior performance of ℓ_1 -loss could be attributed to the sparsity induced by this norm. This ensures that prominent features such as the edges in images are captured effectively. Additionally, the robustness of ℓ_1 -loss to outliers could contribute to its success. The table also demonstrate using ResNet for the Siamese objective yields superior results.

Table V: Impact of feature extractors and metrics for $\mathcal{L}_{\text{siamese}}$ on our overall model performance.

Feature Extractor	ℓ_2 -loss	ℓ_1 -loss	Cosine Similarity
VGG	76.30	76.21	74.52
ResNet	76.45	77.31	74.90
InceptionNet	75.84	75.93	74.66
AlexNet	74.38	74.41	73.93

Table VI: Comparison of different Backbones and masking strategies on the SoccerNet dataset.

Backbone	MAE pretraining	Masking Strategy	Test Acc
ResNet-18	X	-	58.62
ResNet-34	X	-	61.29
ResNet-152	X	-	65.10
ViT-B	✓	Zeroing-Out	75.83
ViT-B	✓	Gaussian Blur	76.47
ViT-B	✓	Motion Blur	77.31

V. CONCLUSION

In this work, we address critical questions concerning the masking policy of MAEs for robust feature extraction in the realm of jersey number recognition. By introducing a novel domain-guided masking strategy, we devise a spatiotemporal network utilizing MAEs and temporal transformers to counter the issues of motion blur and occlusion for the player identification task. Additionally, we refine the existing KfID module to extract more reliable keyframes from the tracklet data. To tackle the issue of limited labelled data resulting due to the KfID module, we leverage a unique keyframe fusion approach to further augment the data. Through quantitative evaluation on different sports including Soccer, Ice Hockey and Baseball, we demonstrate the superior performance of our proposed approach against existing methods. Meticulous ablation studies on the impact of the KfID module and different masking strategies validate the state-of-the-art performance of our network on jersey number recognition.

Future research works involve experimenting with various masking strategies tailored to specific domains, thereby validating the significance of domain-guided masking approaches for robust feature extraction in MAEs.

VI. ACKNOWLEDGEMENT

We acknowledge the support of Stathletes and the Baltimore Orioles through the Mitacs Accelerate Program, as well as the Natural Sciences and Engineering Research Council of Canada (NSERC). Additionally, we express our appreciation to the Digital Research Alliance of Canada for their hardware support.

REFERENCES

- [1] K. Vats, W. J. McNally, P. Walters, D. A. Clausi, and J. S. Zelek, “Ice hockey player identification via transformers and weakly supervised learning,” *2022 IEEE/CVF Conference*

on *Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 3450–3459, 2021.

- [2] B. Balaji, J. Bright, H. Prakash, Y. Chen, D. A. Clausi, and J. Zelek, “Jersey number recognition using keyframe identification from low-resolution broadcast videos,” in *Proceedings of the 6th International Workshop on Multimedia Content Analysis in Sports*, ser. MMSports ’23. New York, NY, USA: Association for Computing Machinery, 2023, p. 123130. [Online]. Available: <https://doi.org/10.1145/3606038.3616162>
- [3] H. Liu and B. Bhanu, “Pose-guided r-cnn for jersey number recognition in sports,” *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 2457–2466, 2019.
- [4] G. Li, S. Xu, X. Liu, L. Li, and C. Wang, “Jersey number recognition with semi-supervised spatial transformer network,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2018, pp. 1864–18647.
- [5] H. Liu, C. Aderon, N. Wagon, H. Liu, S. MacCall, and Y. Gan, “Deep learning-based automatic player identification and logging in american football videos,” *arXiv preprint arXiv:2204.13809*, 2022.
- [6] K. Vats, P. Walters, M. Fani, D. A. Clausi, and J. S. Zelek, “Player tracking and identification in ice hockey,” *ArXiv*, vol. abs/2110.03090, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:238419671>
- [7] W.-L. Lu, J.-A. Ting, J. Little, and K. P. Murphy, “Ieee transactions on pattern analysis and machine intelligence learning to track and identify players from broadcast sports videos.” [Online]. Available: <https://api.semanticscholar.org/CorpusID:7428154>
- [8] K. Vats, M. Fani, D. A. Clausi, and J. Zelek, “Multi-task learning for jersey number recognition in ice hockey,” in *Proceedings of the 4th International Workshop on Multimedia Content Analysis in Sports*, 2021, pp. 11–15.
- [9] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” 2015.
- [10] A. Chan, M. D. Levine, and M. Javan, “Player identification in hockey broadcast videos,” *Expert Syst. Appl.*, vol. 165, p. 113891, 2020.
- [11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” *CoRR*, vol. abs/1706.03762, 2017. [Online]. Available: <http://arxiv.org/abs/1706.03762>
- [12] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, pp. 1735–1780, 1997.
- [13] A. Radford and K. Narasimhan, “Improving language understanding by generative pre-training,” 2018. [Online]. Available: <https://api.semanticscholar.org/CorpusID:49313245>
- [14] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” in *North American Chapter of the Association for Computational Linguistics*, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:52967399>
- [15] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, “An image is worth 16x16 words: Transformers for image recognition at scale,” *ArXiv*, vol. abs/2010.11929, 2020.
- [16] K. He, X. Chen, S. Xie, Y. Li, P. Doll’ar, and R. B. Girshick, “Masked autoencoders are scalable vision learners,” *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15 979–15 988, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:243985980>
- [17] W. G. C. Bandara, N. Patel, A. Gholami, M. Nikkhah, M. Agrawal, and V. M. Patel, “Adamae: Adaptive masking for efficient spatiotemporal learning with masked autoencoders,” *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14 507–14 517, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:253553494>
- [18] I. Kakogeorgiou, S. Gidaris, B. Psomas, Y. Avrithis, A. Bursuc, K. Karantzas, and N. Komodakis, “What to hide from your students: Attention-guided masked image modeling,” in *European Conference on Computer Vision*, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:247627906>
- [19] H. Chen, W. Zhang, Y. Wang, and X. Yang, “Improving masked autoencoders by learning where to mask,” *ArXiv*, vol. abs/2303.06583, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:257496764>
- [20] M. Bertalmío, G. Sapiro, V. Caselles, and C. Ballester, “Image inpainting,” *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 2000. [Online]. Available: <https://api.semanticscholar.org/CorpusID:308278>
- [21] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin, “An iterative regularization method for total variation-based image restoration,” *Multiscale Model. Simul.*, vol. 4, pp. 460–489, 2005. [Online]. Available: <https://api.semanticscholar.org/CorpusID:618185>
- [22] C. Barnes, E. Shechtman, A. Finkelstein, and D. B. Goldman, “Patchmatch: a randomized correspondence algorithm for structural image editing,” *ACM Trans. Graph.*, vol. 28, p. 24, 2009. [Online]. Available: <https://api.semanticscholar.org/CorpusID:26169625>
- [23] A. A. Efros and T. K. Leung, “Texture synthesis by non-parametric sampling,” *Proceedings of the Seventh IEEE International Conference on Computer Vision*, vol. 2, pp. 1033–1038 vol.2, 1999. [Online]. Available: <https://api.semanticscholar.org/CorpusID:221583955>

- [24] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol, "Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion," *J. Mach. Learn. Res.*, vol. 11, pp. 3371–3408, 2010. [Online]. Available: <https://api.semanticscholar.org/CorpusID:17804904>
- [25] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros, "Context encoders: Feature learning by inpainting," 2016.
- [26] M. Chen, A. Radford, J. Wu, H. Jun, P. Dhariwal, D. Luan, and I. Sutskever, "Generative pretraining from pixels," in *International Conference on Machine Learning*, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:219781060>
- [27] H. Bao, L. Dong, and F. Wei, "Beit: Bert pre-training of image transformers," *ArXiv*, vol. abs/2106.08254, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:235436185>
- [28] S. Gerke, K. Müller, and R. Schäfer, "Soccer jersey number recognition using convolutional neural networks," *2015 IEEE International Conference on Computer Vision Workshop (ICCVW)*, pp. 734–741, 2015.
- [29] D. Bhargavi, E. P. Coyotl, and S. Gholami, "Knock, knock. who's there? – identifying football player jersey numbers with synthetic data," 2022.
- [30] G. Jocher, A. Chaurasia, A. Stoken, J. Borovec, NanoCode012, Y. Kwon, K. Michael, TaoXie, J. Fang, imyhxy, Lorna, . Yifu), C. Wong, A. V. D. Montes, Z. Wang, C. Fati, J. Nadar, Laughing, UnglvKitDe, V. Sonck, tkianai, yxNONG, P. Skalski, A. Hogan, D. Nair, M. Strobel, and M. Jain, "ultralytics/yolov5: v7.0 - YOLOv5 SOTA Realtime Instance Segmentation," Nov. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.7347926>
- [31] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, "Deformable detr: Deformable transformers for end-to-end object detection," *ArXiv*, vol. abs/2010.04159, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:222208633>
- [32] C. Weng, B. Curless, P. P. Srinivasan, J. T. Barron, and I. Kemelmacher-Shlizerman, "Humannerf: Free-viewpoint rendering of moving people from monocular video," *CoRR*, vol. abs/2201.04127, 2022. [Online]. Available: <https://arxiv.org/abs/2201.04127>
- [33] S. Hu and Z. Liu, "Gauhuman: Articulated gaussian splatting from monocular human videos," 2023.
- [34] J. Bright, Y. Chen, and J. Zelek, "Mitigating motion blur for robust 3d baseball player pose modeling for pitch analysis," in *Proceedings of the 6th International Workshop on Multimedia Content Analysis in Sports*, ser. MMSports '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 6371. [Online]. Available: <https://doi.org/10.1145/3606038.3616163>