

# MrTrack: Register Mamba for Needle Tracking with Rapid Reciprocating Motion during Ultrasound-Guided Aspiration Biopsy

Yuelin Zhang<sup>1</sup>, Qingpeng Ding<sup>1</sup>, Long Lei<sup>2</sup>, Yongxuan Feng<sup>1</sup>,  
Raymond Shing-Yan Tang<sup>3</sup>, and Shing Shin Cheng<sup>1✉</sup>

<sup>1</sup>Department of Mechanical and Automation Engineering,

<sup>2</sup>Department of Computer Science and Engineering,

<sup>3</sup>Department of Medicine and Therapeutics and Institute of Digestive Disease,  
The Chinese University of Hong Kong

**Abstract.** Ultrasound-guided fine needle aspiration (FNA) biopsy is a common minimally invasive diagnostic procedure. However, an aspiration needle tracker addressing rapid reciprocating motion is still missing. MrTrack, an aspiration needle tracker with a mamba-based register mechanism, is proposed. MrTrack leverages a Mamba-based register extractor to sequentially distill global context from each historical search map, storing these temporal cues in a register bank. The Mamba-based register retriever then retrieves temporal prompts from the register bank to provide external cues when current vision features are temporarily unusable due to rapid reciprocating motion and imaging degradation. A self-supervised register diversify loss is proposed to encourage feature diversity and dimension independence within the learned register, mitigating feature collapse. Comprehensive experiments conducted on both robotic and manual aspiration biopsy datasets demonstrate that MrTrack not only outperforms state-of-the-art trackers in accuracy and robustness but also achieves superior inference efficiency. Project page: <https://github.com/PieceZhang/MrTrack>

**Keywords:** Ultrasound Needle Tracking · Mamba · Temporal Context.

## 1 Introduction

Ultrasound (US)-guided fine needle aspiration (FNA) has emerged as a pivotal minimally invasive diagnostic procedure for percutaneous interventions [22], facilitating the retrieval of cellular material by rapid reciprocating movement [23]. Despite its growing clinical importance, a dedicated needle tracker addressing the rapid needle aspiration is still missing. In practice, the imaging environment during FNA presents unique challenges for a US needle tracker, where reciprocating rapid needle motion during the aspiration [23] induces abrupt nonlinear tip displacement and degraded imaging quality. Although recent models enhance needle tracking by deep learning [13], they were mainly designed for stable environments, yet the nonlinear needle motion and artifact-prone imaging of aspiration biopsy can render these methods ineffective.

Recent developments in tracking reveal that capturing long-range temporal dependencies can enhance tracking [7,14,36]. However, conventional CNN-based [7] or transformer-based [36] trackers struggle with computational constraints when integrating temporal context. In contrast, Mamba [9], a recently developed framework based on state space models (SSMs) [12], has shown its superiority and efficiency in modeling long-term temporal features with satisfactory complexity growth. Although there exist Mamba trackers with temporal modeling [14,28,35,33], their performance and efficiency suffer from inefficient temporal integration. The register then offers a promising solution [4], which is a learnable token that aggregates global context into a compact representation for efficient downstream processing. It calls for an efficient register mechanism to model a spatiotemporal profile for the aspiration needle. The global context from historical search maps should be sequentially *extracted* and stored using the register. The stored latent features from registers should also be effectively *retrieved* to provide external cues for the tracking of the current frame, thus compensating for rapid motion and abrupt appearance changes of the target being tracked.

In this work, **MrTrack**, an ultrasound needle tracker for FNA based on Mamba and register, is proposed with a **register extractor** and a **register retriever**. To the best of our knowledge, it is the first US needle tracker designed for the FNA procedure to robustly track the needle tip undergoing fast reciprocating motion. A **register diversify loss (RD loss)** is also proposed as a self-supervised regularization to prevent collapse in the learned register by encouraging register diversity and dimension independence. The experimental results on both robotic and manual aspiration biopsies show state-of-the-art performance and efficiency of MrTrack. The contributions of MrTrack are fourfold: **1)** The proposed Mamba-based register extractor sequentially condenses historical frames into compact tokens using a trainable register, storing spatiotemporal profiles to provide external cues from past frames when the needle undergoes appearance variation. **2)** The proposed Mamba-based register retriever retrieves temporal cues from historical registers and integrates them with the template map, assisting in tracking an abruptly moving needle using temporal context. **3)** The self-supervised register diversify loss (RD Loss) is proposed to prevent feature collapse by promoting diversity and dimension independence, ensuring an informative register that representationally stores the varying temporal context. **4)** MrTrack achieves highly efficient inferencing, which enables real-time tracking of rapid needle motion under high-speed US imaging.

## 2 Related Work

**Ultrasound needle tracking** By free-hand [17] or robotics-driven [19], US is commonly adopted in many clinical practices [2,18], including US-guided needle insertion. To track a US needle, recent learning-based methods have been proposed, including discriminative correlation filters [24], convolutional neural networks [21], and transformer [29,30]. However, these methods were largely developed for relatively stable needle dynamics. In contrast, the dynamic envi-

ronment of aspiration biopsy, where rapid needle motion and significant imaging artifacts frequently obscure fine needle features [23], exposes the limitations of existing trackers. Although some needle trackers incorporate historical motion [29] to compensate for temporary tip loss, the reciprocating and abrupt movements in aspiration can limit the effectiveness of such motion integration, where the motion can be highly nonlinear and unpredictable.

**Mamba** Recently, based on SSMS [12], Mamba [9] outperforms transformers, leveraging its selective scan mechanism to integrate information over extended temporal ranges with a linear complexity [32]. Mamba-based trackers have demonstrated their capability in tracking tasks and prove the importance of modeling long-term temporal dependencies [14,28]. In contrast, transformer-based [36] or CNN-based [7] trackers with temporal modeling tend to struggle under similar computational constraints due to higher complexity. The existing Mamba trackers with temporal-context modeling face challenges due to the lack of an effective and efficient temporal extraction and integration paradigm [14,28]. As a trainable prompt, the register has been applied in transformer [4] and Mamba [25] for its ability to suppress artifacts and, more critically, function as a compact global representation. The learnable register can thus be helpful for aspiration needle tracking by adapting rapid motion with temporal context.

**Self-supervised regularization** Advances in self-supervised learning have demonstrated that enforcing feature invariance can yield powerful representations [10,34]. Recent approaches like BYOL [8], Barlow Twins [31], and VICReg [1] promote feature diversity and prevent collapse through variance preservation and decorrelation of embedding dimensions. However, the self-supervised regularization for training learnable prompts like the register in Mamba is still missing.

### 3 Method

#### 3.1 Preliminaries on Mamba

The advancement of Mamba is based on the structured state-space models (SSMs), which rely on a continuous linear time-invariant (LTI) system. It can be expressed with the following linear ordinary differential equation (ODE)

$$\begin{aligned} h'(t) &= \mathbf{A}h(t) + \mathbf{B}u(t), \\ y(t) &= \mathbf{C}h(t), \end{aligned} \tag{1}$$

where  $\mathbf{A} \in \mathbb{R}^{N \times N}$ ,  $\mathbf{B} \in \mathbb{R}^{N \times 1}$ , and  $\mathbf{C} \in \mathbb{R}^{1 \times N}$ . Following the equations, the input  $u(t) \in \mathbb{R}$  is mapped to the output  $y(t) \in \mathbb{R}$  by a linear transformation with  $h(t) \in \mathbb{R}^{N \times 1}$ . A discretization is then required for this continuous system. With  $\Delta$  denoting the timescale step size, the discretized system is given by

$$\begin{aligned} h_t &= \bar{\mathbf{A}}h_{t-1} + \bar{\mathbf{B}}u_t, \\ y_t &= \mathbf{C}h_t, \end{aligned} \tag{2}$$

where  $\mathbf{A}$  and  $\mathbf{B}$  is discretized to be  $\bar{\mathbf{A}} = \exp(\Delta \cdot \mathbf{A})$  and  $\bar{\mathbf{B}} = (\Delta \cdot \mathbf{A})^{-1}(\exp(\Delta \cdot \mathbf{A}) - I) \cdot \Delta \cdot \mathbf{B}$ . In Mamba [9], the selective scan mechanism is incorporated to

learn  $\mathbf{B}, \mathbf{C}, \Delta$  from input  $x$  by an input-dependent way, which are given by  $\mathbf{B} = \text{Linear}(x)$ ,  $\mathbf{C} = \text{Linear}(x)$ ,  $\Delta = \text{Softplus}(\tilde{\Delta} + \text{Linear}(x))$ , where  $\tilde{\Delta}$  is learnable.

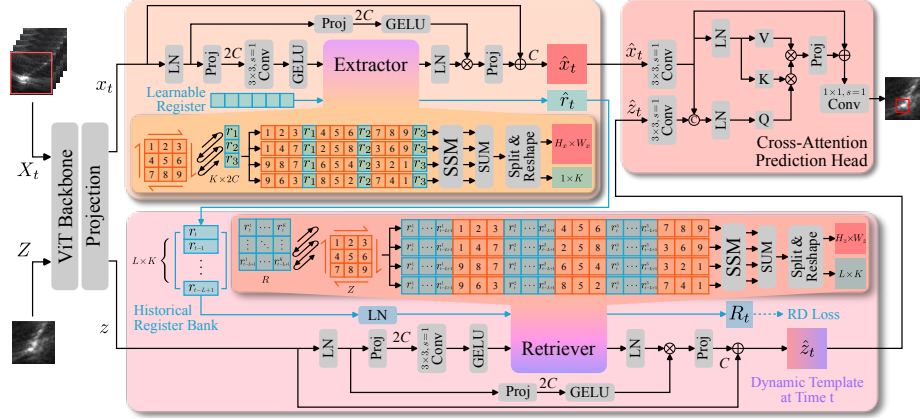


Fig. 1: Overview of MrTrack. The notes (e.g.  $3 \times 3, s = 1$ ) on convolution layers represent the kernel size and stride.

### 3.2 Overview of MrTrack

As shown in Fig. 1, a ViT-Base [5] is adopted as the backbone, where the embedded features of the template map  $z \in \mathbb{R}^{H_z \times W_z \times C}$  and search map  $x_t \in \mathbb{R}^{H_x \times W_x \times C}$  ( $C = 256$ ) are obtained, where  $t$  denotes the current time step. The  $x_t$  and  $z$  are then passed to the Mamba-based register extractor and retriever, respectively, to extract registers from the search map and retrieve information into the template map. The register is trained in a self-supervised manner by the register diversify loss  $\mathcal{L}_{RD}$ . With the obtained  $\hat{x}_t$  and  $\hat{z}_t$ , the tracking result is predicted by the cross-attention prediction head.

### 3.3 Mamba-based Register Extractor and Register Retriever

The proposed Mamba-based register extraction-retrieval framework benefits aspiration needle tracking by aggregating long-term visual features at minimal overhead. The extractor learns compact global representation from a search map  $x_t$  through a trainable register  $r \in \mathbb{R}^{k \times 2C}$  ( $k = 8$ ), where  $r$  is shareable for all frames. Specifically, after a series of layer normalization, projection, and convolution,  $x_t$  is passed to the extractor together with  $r$ , in which a cross-map scanning is performed to evenly insert the register *behind* each image segment throughout the  $x_t$ , as shown in Fig. 1. This insertion ensures a thorough feature translation. The fused vectors are then learned by the SSM, where not only the feature from  $x_t$  is conveyed into the output  $r_t$  but also the  $x_t$  itself is modeled

to be  $\hat{x}_t$  with richer representation and suppressed artifacts [25]. By the extractor, the dedicated compact descriptor for each frame is efficiently obtained, then stored in the register bank  $R_t = [r_t, \dots, r_{t-L+1}] \in \mathbb{R}^{L \times k \times 2C}$ ,  $L = 300$ .

The retriever retrieves the information from the register bank  $R_t$  to the template  $z$ . Through cross-map scanning, the descriptors  $r_i$  in  $R_t$  are evenly inserted back into the sequence of image tokens of  $z$ , yet *before* each image segment. After SSM modeling, the temporal context in  $R_t$  is retrieved into  $\hat{z}_t$ , which is a dynamic template containing global temporal context from time  $t$  to  $t-L$ . Through the retriever, the feature of the initial template is aggregated with the latest transient cues, which can be helpful during aspiration when abrupt and rapid motion happens.

### 3.4 Register Diversify Loss

To ensure that the learned register does not collapse and to encourage distinctiveness in the register bank, the register diversify loss (RD loss) is proposed. RD loss comprises two components: a variance promoting term  $\mathcal{L}_{var}$  and a cross-register diversify term  $\mathcal{L}_{div}$ . The variance promoting term  $\mathcal{L}_{var}$  penalizes feature dimensions whose standard deviation falls below a threshold  $\tau$  ( $\tau = 1$  in this paper) to encourage diversity. The *softplus* is applied here as a smooth penalty with differentiability, given by  $\text{softplus}(v) = \ln(1 + e^v)$ .  $\mathcal{L}_{var}$  is then defined by

$$\mathcal{L}_{var} = \frac{1}{d} \sum_{i=1}^d \text{softplus}(\tau - \sqrt{V_i + \epsilon}), \quad (3)$$

where  $d$  is the dimensionality,  $\epsilon$  is a small constant,  $V_i$  is variance of the  $i_{\text{th}}$  dimension. The variance of the learnable register is then promoted with  $\mathcal{L}_{var}$ .

Except for promoting variance of the register, the cross-register diversify term  $\mathcal{L}_{div}$  is proposed to encourage the diversity within the register bank, thus preventing the register extractor from encoding similar features for different frames. Assume each descriptor  $r_i$  in  $R_t = [r_t, \dots, r_{t-L}]$  is defined by  $r_i = [r_{i,1}, \dots, r_{i,d}] \in \mathbb{R}^{k \times d}$ , where  $d = 2C$ ,  $\mathcal{L}_{div}$  is presented with a covariance format given by

$$\mathcal{L}_{div} = \frac{1}{dL(L-1)} \sum_{p=1}^d \sum_{\substack{i,j=1 \\ i \neq j}}^L [(r_{i,p} - \bar{r})(r_{j,p} - \bar{r})]^2, \text{ where } \bar{r} = \frac{1}{L} \sum_{n=1}^L r_{n,p}. \quad (4)$$

$\mathcal{L}_{div}$  drives the off-diagonal elements of the covariance close to zero thus promoting decorrelation. Different from the covariance regularization in [1] that decorrelates the dimension within the same sample,  $\mathcal{L}_{div}$  decorrelates dimensions between different samples  $r_i$  from the register bank, thus encouraging the model to learn independently varying features for different frames. By combining Eq. 3 and Eq. 4 with scalar  $\alpha = \beta = 0.01$ , the RD loss  $\mathcal{L}_{RD}$  is given by

$$\mathcal{L}_{RD} = \alpha \mathcal{L}_{var} + \beta \mathcal{L}_{div}. \quad (5)$$

This self-supervised regularization prevents collapse by encouraging dimension independence in the learnable register and promoting diversity in register bank.

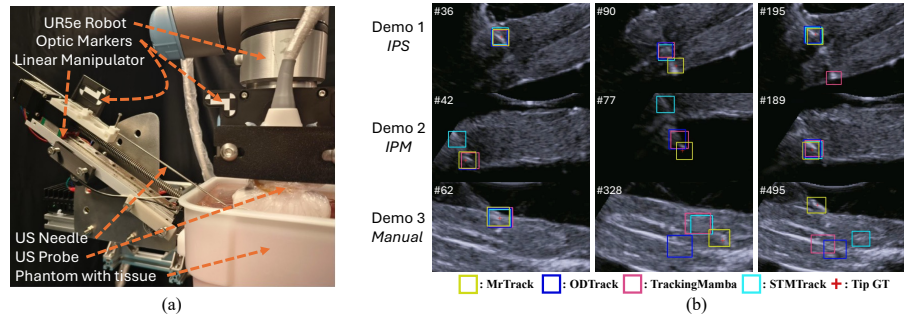


Fig. 2: Experiment setup (a) and tracking demonstration (b). More demonstrations are shown in the supplementary video.

## 4 Experiments and Results

### 4.1 Experiment Setup and Data Collection

As shown in Fig. 2(a), a robotic US needle insertion platform and a US imaging system with a Wisonic Clover 60 US machine and a Wisonic C5-1 convex transducer (sampling frequency 4.5 MHz) was adopted for the experiments. A 20 gauge fine aspiration needle [6] was used to insert into a phantom made with fresh pork supported by solidified agar. The ground truth position of the needle tip was collected by an optic tracker (ClaroNav MicronTracker 3), which has an RMSE of 0.189 mm in our setup. Datasets for both robotic and manual needle manipulation experiments simulating the aspiration biopsy procedure were collected. During the robotic insertion, the needle was installed on a linear manipulator, while the US probe was held by a UR5e robot arm. The needle was first inserted at a constant velocity of 20 mm/s to reach the target with a depth of more than 50 mm. After reaching the target, the needle would move back and forth five times at a velocity of 30 mm/s and a depth of 15 mm to perform aspiration. Both in-plane-static (IPS) and in-plane-moving (IPM) techniques [13,29] were performed at three needle insertion angles ( $30^\circ$ ,  $45^\circ$ ,  $60^\circ$ ). The robotic insertion dataset includes 54,428 frames in videos ( $1920 \times 1080$ , 30 FPS) from 239 trials (105 IPS, 134 IPM). After sampling it by 5 to remove redundancy, the dataset was divided into training (7,619 samples), validation (1,088 samples), and testing (2,177 samples) sets in the ratio 7:1:2. During the manual insertion, both the needle and probe were held by the user. Four users were asked to perform four aspirations each, with the needle being moved back and forth two to three times per second. It resulted in a dataset that includes 16 videos (10,948 frames), which is only for testing to validate generalization.

During the evaluation, all models were implemented using PyTorch and trained on the same dataset with four NVIDIA RTX 4090 GPUs. All models adopt the same training strategy (350 epochs, batch size 32, AdamW optimizer) and the same input size of  $384 \times 384$  and  $192 \times 192$  for search and template map.

Scaling, blur, and position shifting were adopted for image augmentation. The learning rate was set to 5e-4 and dropped by 10 after 200 epochs.

Table 1: Evaluation and ablation study results. The methods with the best and the second best performance are in **red** and **cyan**. AUC and  $P$  are reported in percentage (%), and Err and SD are reported in mm. In ablation studies, the performance difference relative to the baseline has been highlighted.

Method	Robotic									Manual			FPS
	In-plane-static (IPS)			In-plane-moving (IPM)			Mean						
	AUC $\uparrow$	$P_{\uparrow}$	Err $\pm$ SD $\downarrow$	AUC $\uparrow$	$P_{\uparrow}$	Err $\pm$ SD $\downarrow$	AUC $\uparrow$	$P_{\uparrow}$	Err $\pm$ SD $\downarrow$	AUC $\uparrow$	$P_{\uparrow}$	Err $\pm$ SD $\downarrow$	
SiamRPN++ [15]	44.1	60.0	4.71 $\pm$ 3.10	58.2	75.0	3.28 $\pm$ 3.00	50.3	66.7	4.14 $\pm$ 3.06	42.0	59.3	5.50 $\pm$ 3.73	31.3
SiamCAR [11]	45.3	68.8	4.20 $\pm$ 2.38	59.7	80.5	2.84 $\pm$ 2.25	52.0	74.6	3.66 $\pm$ 2.33	40.3	57.1	5.72 $\pm$ 3.49	48.2
SiamBAN [3]	52.9	73.1	3.90 $\pm$ 2.51	61.5	89.8	2.39 $\pm$ 2.09	56.2	79.6	3.30 $\pm$ 2.34	49.0	69.5	3.60 $\pm$ 3.18	33.2
SwinTrack [16]	51.0	72.6	3.93 $\pm$ 1.83	64.4	93.1	1.91 $\pm$ 1.72	57.1	80.6	3.13 $\pm$ 1.79	48.2	70.6	3.62 $\pm$ 2.85	60.1
STMTrack [7]	54.7	76.9	3.80 $\pm$ 2.00	64.1	92.5	2.10 $\pm$ 1.98	59.3	84.6	3.14 $\pm$ 2.00	51.5	76.8	3.05 $\pm$ 2.42	24.4
TrackingMamba [26]	55.2	79.7	3.62 $\pm$ 1.71	63.9	92.9	1.90 $\pm$ 1.62	59.5	86.3	2.93 $\pm$ 1.72	54.1	77.0	3.39 $\pm$ 2.98	63.1
ODTrack [36]	57.6	78.6	3.76 $\pm$ 1.73	63.5	94.5	1.81 $\pm$ 1.69	60.8	85.0	2.98 $\pm$ 1.71	55.4	76.2	3.20 $\pm$ 2.11	53.2
MrTrack	63.6	92.0	2.51 $\pm$ 1.31	67.2	93.2	1.88 $\pm$ 1.10	65.4	92.3	2.29 $\pm$ 1.22	59.0	83.9	2.80 $\pm$ 1.59	73.9

Ablations	Robotic (Mean)			Manual (Mean)			FPS
	AUC $\uparrow$	$P_{\uparrow}$	Err $\pm$ SD $\downarrow$	AUC $\uparrow$	$P_{\uparrow}$	Err $\pm$ SD $\downarrow$	
Baseline: MrTrack	65.4	92.3	2.29 $\pm$ 1.22	59.0	83.9	2.80 $\pm$ 1.59	73.9
$v_{r1}$ : $L = 100, k = 8$	63.5(-1.9)	92.2(-0.1)	2.27(-0.02) $\pm$ 1.73(+0.51)	59.7(+0.7)	82.1(-1.8)	2.94(+0.14) $\pm$ 1.72(+0.13)	76.2(+2.3)
$v_{r2}$ : $L = \infty, k = 8$	65.9(+0.5)	94.0(+1.7)	2.19(-0.10) $\pm$ 1.34(+0.12)	57.8(-1.2)	80.6(-3.3)	2.94(+0.14) $\pm$ 1.80(+0.21)	58.0(-15.9)
$v_{r3}$ : $L = 300, k = 2$	65.1(-0.3)	92.2(-0.1)	2.24(-0.05) $\pm$ 1.25(+0.03)	58.1(-0.9)	82.9(-1.0)	2.92(+0.12) $\pm$ 1.55(-0.04)	73.7(-0.2)
$v_{r4}$ : $L = 300, k = 32$	62.4(-3.0)	91.2(-1.1)	2.31(+0.02) $\pm$ 1.39(+0.17)	53.8(-5.2)	81.9(-2.0)	3.19(+0.39) $\pm$ 1.99(+0.40)	68.0(-5.9)
$v_{M1}$ : <i>Trans E &amp; R</i>	59.0(-6.4)	85.8(-6.5)	3.10(+0.81) $\pm$ 1.40(+0.18)	55.3(-3.7)	78.8(-5.1)	2.99(+0.19) $\pm$ 1.65(+0.06)	61.3(-12.6)
$v_{M2}$ : w/o Register	61.2(-4.2)	88.5(-3.8)	2.49(+0.20) $\pm$ 1.79(+0.57)	54.2(-4.8)	79.4(-4.5)	3.38(+0.58) $\pm$ 2.03(+0.44)	76.1(+2.2)
$v_{RDL}$ : w/o $\mathcal{L}_{RD}$	64.1(-1.3)	90.5(-1.8)	2.40(+0.11) $\pm$ 1.65(+0.43)	56.4(-2.6)	83.5(-0.4)	2.88(+0.08) $\pm$ 1.77(+0.18)	73.9(-0.0)

## 4.2 Results and Discussion

The evaluation was performed between MrTrack and seven trackers, including classic Siamese trackers [3,11,15], a tracker with motion integration [16], a Mamba-based tracker [26], and trackers with temporal modeling [7,36]. The results are reported by four metrics, including average error (Err) and standard deviation (SD) in mm, and commonly used tracking evaluation metrics area under curve (AUC) [27] and precision ( $P$ ) [20] in percentage (%). The inference speed (FPS) is also reported. As shown in Tab. 1, MrTrack achieves the best performance for most metrics while maintaining the best efficiency with an inference speed of 73.9 FPS. In Err for robotic and manual insertion, it outperforms the second-best method by 21.8% and 8.2%, respectively. For SD, the improvements are even more substantial, reaching 28.7% and 24.6%. For manual insertion, which involves more abrupt motion and imaging disturbances, MrTrack still achieves consistent SOTA performance, which shows its robustness and generalizability to unseen images and motion, as the manual insertion data was not used for training. The demonstration in Fig. 2(b) also shows that MrTrack achieves

the most accurate tracking of the aspiration needle with reciprocating motion. More examples are provided in the supplementary video.

This result demonstrates the effectiveness of MrTrack with the Mamba-based register, which captures the global spatiotemporal profile of a rapidly moving needle during aspiration. Among the comparison methods, ODTrack [36] also adopts temporal learning yet with worse performance and lower FPS. It may be due to its *token propagation* strategy, which with a fixed-length token to record the temporal context and pass through the whole network, leading to error accumulation. In contrast, MrTrack extracts and stores each frame individually, ensuring that every historical token possesses the same importance, thereby preventing error accumulation. Although TrackingMamba [26] also adopts a Mamba framework, it performs poorly because it only utilizes the vanilla Mamba without temporal context modeling. STMTrack [7] also learns temporal features by simply concatenating historical frames but has lower FPS and worse performance due to the lack of an effective temporal integration strategy. SwinTrack [16] adopts a historical motion fusion strategy to model temporal context, but since the motion in aspiration is nonlinear and abrupt, this strategy may lead to worse performance for this task. Currently, MrTrack still lacks a mechanism to update the template when needle visibility deteriorates. Future work will introduce an explicit template updating strategy to address this issue.

### 4.3 Ablation Study

Ablation studies were carried out on 7 variations.  $v_{r1}$  to  $v_{r4}$  explore the temporal integration capacity of the register, where the results indicate that the baseline structure ( $L = 300, k = 8$ ) with a moderate register bank length  $L$  and a reasonable register scale  $k$  is the most suitable for the aspiration needle tracking task. Although  $v_{r2}$  somewhat improves performance by integrating infinite registers, it also leads to a significant decrease in inference speed, as well as a loss of robustness caused by error accumulation.  $v_{M1}$  replaces the proposed Mamba-based register extractor and retriever with a fundamental self-attention transformer [5] similar to [4]. This results in degradation across all metrics, indicating that the Mamba-based register framework is more effective than transformer.  $v_{M2}$  further shows the effectiveness of the register-based temporal integration, where the register and register bank are removed while keeping Mamba only.  $v_{RDL}$  only applies supervised training without  $\mathcal{L}_{RD}$ , where the result shows the effectiveness, i.e., register diversity and dimension independence are promoted, thus enhancing the whole register framework and aspiration needle tracking performance. The robustness of  $\mathcal{L}_{RD}$  is also showed by avoiding tuning the value of hyper-parameters  $\alpha$  and  $\beta$ .

The results demonstrate the effectiveness of the proposed modules under rapid motion and degraded image quality, where both the learnable register mechanism and the register diversify loss are crucial for robust temporal integration and feature discrimination. It underscores the importance of capturing long-term temporal context, which not only improves tracking accuracy but also enhances stability across various insertion techniques.



## 5 Conclusion

In this paper, a US needle tracker, MrTrack, has been proposed for FNA which features fast reciprocating needle motion. It leverages a Mamba-based register mechanism and a self-supervised register diversify loss to ensure robust and accurate needle tip tracking. In the future, more experiments will be conducted on large-scale manual insertion data and real clinical video, as well as on handling needle out-of-plane caused by deflection.

**Acknowledgments.** Research reported in this work was supported in part by Research Grants Council (RGC) of Hong Kong (T45-401/22-N, CUHK 14217822, CUHK 14207823, and AoE/E-407/24-N) and in part by Innovation and Technology Commission of Hong Kong (ITS/234/21, ITS/235/22, ITS/224/23, ITS/225/23, and Multi-scale Medical Robotics Center, InnoHK). The content is solely the responsibility of the authors and does not necessarily represent the official views of the sponsors.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article

## References

1. Bardes, A., et al.: Vicreg: Variance-invariance-covariance regularization for self-supervised learning. arXiv preprint arXiv:2105.04906 (2021)
2. Cai, Y., et al.: Autonomous robotic ultrasound approach for fetoscope tracking by fusing optical and 2d ultrasound data. IEEE Robotics and Automation Letters (2024)
3. Chen, Z., et al.: Siamban: Target-aware tracking with siamese box adaptive network. IEEE Transactions on Pattern Analysis and Machine Intelligence **45**(4), 5158–5173 (2022)
4. Darcet, T., et al.: Vision transformers need registers. arXiv preprint arXiv:2309.16588 (2023)
5. Dosovitskiy, A.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
6. D’Onofrio, M., et al.: Ultrasound-guided percutaneous fine-needle aspiration of solid pancreatic neoplasms: 10-year experience with more than 2,000 cases and a review of the literature. European radiology **26**, 1801–1807 (2016)
7. Fu, Z., et al.: Stmtrack: Template-free visual tracking with space-time memory networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13774–13783 (2021)
8. Grill, J.B., et al.: Bootstrap your own latent-a new approach to self-supervised learning. Advances in neural information processing systems **33**, 21271–21284 (2020)
9. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces. arXiv preprint arXiv:2312.00752 (2023)
10. Gui, J., et al.: A survey on self-supervised learning: Algorithms, applications, and future trends. IEEE Transactions on Pattern Analysis and Machine Intelligence (2024)

11. Guo, D., et al.: Siamcar: Siamese fully convolutional classification and regression for visual tracking. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6269–6277 (2020)
12. Kalman, R.E.: A new approach to linear filtering and prediction problems (1960)
13. Kimbowa, A., et al.: Advancements in needle visualization enhancement and localization methods in ultrasound: a literature review. *Artificial Intelligence Surgery* **4**(3), 149–169 (2024)
14. Lai, S., et al.: Mambavt: Spatio-temporal contextual modeling for robust rgb-t tracking. *arXiv preprint arXiv:2408.07889* (2024)
15. Li, B., et al.: Siamrpn++: Evolution of siamese visual tracking with very deep networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4282–4291 (2019)
16. Lin, L., et al.: Swintrack: A simple and strong baseline for transformer tracking. *Advances in Neural Information Processing Systems* **35**, 16743–16754 (2022)
17. Luo, M., et al.: Recon: Online learning for sensorless freehand 3d ultrasound reconstruction. *Medical Image Analysis* **87**, 102810 (2023)
18. Ma, X., et al.: Feasibility of pointcloud-based ultrasound-ct registration towards automated, robot-assisted image-guidance in spine surgery. In: 2024 International Symposium on Medical Robotics (ISMR). pp. 1–7. IEEE (2024)
19. Ma, X., et al.: Cross-modality registration using bone surface pointcloud for robotic ultrasound-guided spine surgery. *Journal of Medical Robotics Research* p. 2540004 (2025)
20. Muller, M., et al.: Trackingnet: A large-scale dataset and benchmark for object tracking in the wild. In: Proceedings of the European conference on computer vision (ECCV). pp. 300–317 (2018)
21. Mwikirize, C., et al.: Single shot needle tip localization in 2d ultrasound. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 637–645. Springer (2019)
22. Richardson, MD, R.R., Richardson, R.R.: Imaging modalities: Advantages and disadvantages. *Atlas of Acquired Cardiovascular Disease Imaging in Children* pp. 1–4 (2017)
23. Rumack, C.M., Levine, D.: *Diagnostic ultrasound*. Elsevier Health Sciences (2017)
24. Shen, C., et al.: Discriminative correlation filter network for robust landmark tracking in ultrasound guided intervention. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part V 22. pp. 646–654. Springer (2019)
25. Wang, F., et al.: Mamba-r: Vision mamba also needs registers. *arXiv preprint arXiv:2405.14858* (2024)
26. Wang, Q., et al.: Trackingmamba: Visual state space model for object tracking. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* (2024)
27. Wu, Y., et al.: Online object tracking: A benchmark. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2411–2418 (2013)
28. Xie, J., et al.: Robust tracking via mamba-based context-aware token learning. *arXiv preprint arXiv:2412.13611* (2024)
29. Yan, W., et al.: Learning-based needle tip tracking in 2d ultrasound by fusing visual tracking and motion prediction. *Medical Image Analysis* **88**, 102847 (2023)
30. Yan, W., et al.: Task-oriented network design for visual tracking and motion filtering of needle tip under 2d ultrasound. *IEEE Transactions on Medical Imaging* (2024)

31. Zbontar, J., et al.: Barlow twins: Self-supervised learning via redundancy reduction. In: International conference on machine learning. pp. 12310–12320. PMLR (2021)
32. Zhang, H., et al.: A survey on visual mamba. *Applied Sciences* **14**(13), 5683 (2024)
33. Zhang, Y., et al.: Motion-guided dual-camera tracker for endoscope tracking and motion analysis in a mechanical gastric simulator. *arXiv preprint arXiv:2403.05146* (2024)
34. Zhang, Y., et al.: A unified framework for microscopy defocus deblur with multi-pyramid transformer and contrastive learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11125–11136 (2024)
35. Zhang, Y., et al.: Mambaxctrack: Mamba-based tracker with ssm cross-correlation and motion prompt for ultrasound needle tracking. *IEEE Robotics and Automation Letters* (2025)
36. Zheng, Y., et al.: Odtrack: Online dense temporal token learning for visual tracking. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 38, pp. 7588–7596 (2024)