

MEMTrack: A DEEP LEARNING-BASED APPROACH TO MICROROBOT TRACKING IN DENSE AND LOW-CONTRAST ENVIRONMENTS

Medha Sawhney*

Dept. of Computer Science
Virginia Tech
medha@vt.edu

Bhas Karmarkar*

Dept. of Mechanical Engg.
Virginia Tech
bhasnk@vt.edu

Eric J. Leaman

Dept. of Mechanical Engg.
Virginia Tech
leamanej@vt.edu

Arka Daw

Dept. of Computer Science
Virginia Tech
darka@vt.edu

Anuj Karpatne

Dept. of Computer Science
Virginia Tech
karpatne@vt.edu

Bahareh Behkam

Dept. of Mechanical Engg.
Virginia Tech
behkam@vt.edu

ABSTRACT

Tracking microrobots is challenging, considering their minute size and high speed. As the field progresses towards developing microrobots for biomedical applications and conducting mechanistic studies in physiologically relevant media (e.g., collagen), this challenge is exacerbated by the dense surrounding environments with feature size and shape comparable to microrobots. Herein, we report Motion Enhanced Multi-level Tracker (MEMTrack), a robust pipeline for detecting and tracking microrobots using synthetic motion features, deep learning-based object detection, and a modified Simple Online and Real-time Tracking (SORT) algorithm with interpolation for tracking. Our object detection approach combines different models based on the object’s motion pattern. We trained and validated our model using bacterial micro-motors in collagen (tissue phantom) and tested it in collagen and aqueous media. We demonstrate that MEMTrack accurately tracks even the most challenging bacteria missed by skilled human annotators, achieving precision and recall of 77% and 48% in collagen and 94% and 35% in liquid media, respectively. Moreover, we show that MEMTrack can quantify average bacteria speed with no statistically significant difference from the laboriously-produced manual tracking data. MEMTrack represents a significant contribution to microrobot localization and tracking, and opens the potential for vision-based deep learning approaches to microrobot control in dense and low-contrast settings. All source code for training and testing MEMTrack and reproducing the results of the paper have been made publicly available <https://github.com/sawhney-medha/MEMTrack>.

1 INTRODUCTION

In recent years, microrobotic systems have burgeoned due to their potential in various fields, including targeted drug delivery, minimally invasive surgery, and biosensing Wang et al. (2021). Based on their mode of actuation and target application, microrobots range between ~ 1 -1000 μm in size, and their average speeds vary vastly from ~ 1 $\mu\text{m/s}$ to 800 $\mu\text{m/s}$, with instantaneous speeds upwards of 800 $\mu\text{m/s}$ Li et al. (2020); Jiang et al. (2022). These properties make microrobots very effective in reaching currently inaccessible areas of the human body but incredibly difficult to visualize and track. Traditionally, microrobots have been studied in aqueous environments (Fig. 1A-B). The growing focus shift in the microrobotic field from system development to biomedical application-oriented implementations necessitates the operation and control of such systems in physiologically relevant environments, and aqueous media do not always represent the conditions and interactions experienced *in vivo*. Furthermore, advancing the current understanding of physical underpinnings




Figure 1: Self-propelled micromotors, such as bacteria, are easy to track in high-contrast fluorescent images (A); however, they are more difficult to detect in high acquisition rate bright-field images in liquid (aqueous) medium (B). The complexity further increases in dense environments such as collagen with feature shapes and sizes similar to those of the microscale objects of interest (C) The yellow arrows represent bacteria locations. (D, E, F, and G) are sample trajectories of bacteria in different motility sub-populations. All scale bars are 20 μm .

of microrobot behaviors *in vivo* requires the study of these systems in physiologically relevant *in vitro*, or *ex vivo*, or *in vivo* environments Soto et al. (2020).

The fast speeds of microrobotic systems (*i.e.*, 10s-100s of body lengths per second) necessitate high frame rate image acquisition, primarily attainable using bright-field imaging. The resulting grayscale images (Fig. 1B) have significantly lower contrast than fluorescent images (Fig. 1A), usually used for automated localization and tracking. Also, microrobots ($\sim 1 \mu\text{m}$ in size) often swim in the three-dimensional (3D) space, which translates to intermittent movement of the objects of interest moving in and out of the focal plane, adding another layer of complexity. Moreover, self-propelled micro-motors (*e.g.*, bacteria or catalytic motors) exhibit random walk, making it difficult to track them consistently in every frame. Operation in physiologically relevant environments with dense backgrounds and feature sizes and shapes comparable to those of microrobots (Fig. 1C) further exacerbated this problem. There has been significant progress in real-time tracking of microrobots using clinical imaging modalities that produce high-contrast images Aziz et al. (2019); Botros et al. (2023); Tiryaki et al. (2022); Pane et al. (2021); Aziz et al. (2020); Bae et al. (2020); Wang & Zhang (2020); however, automated tracking of micro/nano-scale objects in bright-field image sequences within dense backgrounds remains a largely unsolved challenge.

A few commonly used techniques exist for manual or semi-automated tracking of microscale objects. Each technique has different features and capabilities, such as particle identification and tracking or morphology analysis. The most widely used tool for manual tracking is ImageJ (Fiji), an open-source software for processing and analysis of scientific images Schneider et al. (2012). Fiji is equipped with several plugins for manual and semi-automated tracking, including Trackmate Tinevez et al. (2017); Ershov et al. (2021), MtrackJ Meijering et al. (2012), CellProfiler McQuin et al. (2018); Stirling et al. (2021), MosaicSuite - Particle Tracker Sbalzarini & Koumoutsakos (2005), FARSIGHT Bjornsson et al. (2008), BioImageXD Kankaanpää et al. (2012), and Icy De Chaumont et al. (2012). The semi-automated methods require user inputs, such as filters or thresholds and manual intervention to prune false positive data. None of these methods were developed for object detection in dense backgrounds, where the visual ambiguity between the dense background and microscale objects makes tracking error-prone, tedious, and time-consuming. Apart from these tools,

there exist several other methods for particle detection and tracking in 2D and 3D spaces Balomenos et al. (2017); Wang et al. (2010); Klein et al. (2012); Cornwell et al. (2020); Rust et al. (2011); Feng et al. (2011); Van Der Schaar et al. (2008); Stylianidou et al. (2016); Gardini et al. (2015). However, these methods primarily rely on the high contrast between the object of interest and the background that is unique to fluorescent images. The longer exposure time of fluorescence imaging (~ 100 ms compared to ~ 10 ms for bright-field imaging) results in a reduced acquisition rate and loss of important temporal data, making it unfeasible for many microrobotic applications. Other models have also been developed for tracking objects in aqueous environments using bright-field images. These models consider changes in cells’ appearance and overlaps during colony proliferation in time-lapse videos Vallotton et al. (2017); Spahn et al. (2022). They can also track single cells in aqueous media in bright-field but are still not fully automated and require manual interventions Schwanbeck et al. (2020); Xie et al. (2008); Gardini et al. (2015); Mehta et al. (2016); Newby et al. (2018); Hernandez et al. (2018); Pottash et al. (2017). To the authors’ best knowledge, automated tracking of micro-motors in dense environments, with features and dimensions similar to those of objects of interest, has not yet been realized.

In this work, we report the development of a deep learning-based approach for the detection and tracking of microrobots in dense environments. Existing deep learning-based multi-object tracking works have primarily focused on detecting and tracking objects that are easily distinguished from the background (*e.g.*, pedestrians and cars) Milan et al. (2016). In this work, we present a Motion Enhanced Multi-level Tracker (MEMTrack) for tracking microscale objects in dense environments, where the object of interest is almost indistinguishable from the background features (Fig. 2C). We train and test MEMTrack using bacterial micro-biomotors, one of the most commonly used biomotors in biohybrid microrobotic systems Webster-Wood et al. (2023). To demonstrate the wide utility of the MEMTrack, we demonstrate its performance in collagen, the most abundant extracellular matrix (ECM) protein in the body, with feature sizes and shapes comparable to those of bacteria (Fig. 1C) and in aqueous media (Fig. 1B). Our results demonstrate that our pipeline can accurately predict and track both visually identifiable and hard-to-detect bacteria. Our proposed pipeline represents a significant contribution to microrobot image analysis in dense environments using computer vision. Moreover, it opens the potential of applying deep learning-based methods for vision-based control of microrobots in dense and low-contrast settings for various applications, including disease diagnosis and treatment.

2 METHODS

2.1 MOTION ENHANCED MULTI-LEVEL TRACKER (MEMTRACK)

The proposed pipeline for MEMTrack is shown in Fig. 2. MEMTrack consists of four modules—Motion Enhancer, Multi-level Object Detector, False Positive Pruner, and Interpolated Tracker. Before describing each of the modules, we define the notations used throughout the paper. We define the input video with T frames as $I^{1..T} = [I^1, I^2, \dots, I^T]$, where $I^t \in \mathbb{R}^{C \times H \times W}$ denotes the t -th frame and C , H and W are the number of channels, height and width of the image, respectively. We utilized a tracking-by-detection approach, wherein tracking is done on top of predictions from the detection module.

2.1.1 MOTION ENHANCER MODULE

Object detectors designed for single-frame detection overlook the object’s position in preceding or subsequent frames, which limits their tracking effectiveness. Incorporating the concept of motion into the object detection model is crucial to enable accurate detection of microscale objects in dense environments with intermixing backgrounds and foregrounds (Fig. 1C). Therefore, we implemented feature engineering techniques to capture two types of motion features, optical flow features and median deviation features, which are then augmented or stacked with the image features to enhance detection accuracy, as shown in Fig. 2A and described below.

Optical Flow Features. Optical flow Horn & Schunck (1981) is a technique that estimates the motion of objects in an image sequence by analyzing the changes in pixel intensities between consecutive frames. We used the Lucas-Kanade Lucas & Kanade (1981) method for optical flow computation,




Figure 2: Overview of our proposed pipeline for bacteria tracking consists of four modules: (A) Motion Enhancer, which adds motion features to the input frames, (B) Multi-level Object Detector, which detects objects of varying motility levels using the RetinaNet model, (C) False Positive Pruner, which filters the predicted object occurrences to reduce false positives, and (D) Interpolated Tracker, which tracks objects over time using the SORT algorithm with interpolation.

which can be expressed mathematically as:

$$\partial I_x u + \partial I_y v + \partial I_t = 0, \quad (1)$$

where $u = \frac{dx}{dt}$ and $v = \frac{dy}{dt}$ represent the x and y components of the optical flow vector for the t -th frame, $\partial I_x = \frac{\partial I}{\partial x}$, $\partial I_y = \frac{\partial I}{\partial y}$, and $\partial I_t = \frac{\partial I}{\partial t}$ are the image gradients in the x , y , and time (t) dimensions, respectively. Solving this equation yields the optical flow vector $O = [u, v]$ for the t -th frame in the video. The length of this optical flow vector, which corresponds to the magnitude of motion at each pixel between consecutive frames, is then considered as an additional feature channel along with the grayscale image.

Median Deviation Features. While optical flow captures changes across consecutive frames, we also need features that capture slower motion trends with respect to a static background. To this effect, we used the median deviation as another feature channel, which is the pixel-wise difference between the intensity at a pixel and the median intensity at the pixel across all frames in the video. Note that the median represents the background at every pixel that remains static for the majority of frames. Mathematically, we defined median deviation $\Delta I = |I^t - \text{median}(I^{1..T})|$, where $\text{median}(I^{1..T})$ is the pixel-wise median for the video across frames 1 through T .

Finally, the input image features I^t , concatenated with the optical flow features O and the median deviation ΔI , are used as inputs to the Multi-level Object Detector.

2.1.2 MULTI-LEVEL OBJECT DETECTOR MODULE

Object detection is a fundamental task in computer vision and plays a crucial role in various applications such as automation and decision-making. The two main approaches used are two-stage detectors and one-stage detectors. Two-stage detectors such as Region-based Convolutional Neural Network (R-CNN) Girshick et al. (2016) and Fast R-CNN Girshick (2015) offer high accuracy but suffer from high computational complexity, while one-stage detectors such as You Only Look Once (YOLO) Redmon et al. (2016); Bochkovskiy et al. (2020); Redmon & Farhadi (2017); Wang et al. (2022) provide faster inference but may struggle with detection of small objects in dense environments.

In our pipeline, we used the one-stage deep learning detector RetinaNetLin et al. (2017) as the base object detector. RetinaNet introduces the novel “focal loss” function, which assigns higher weights to frequently misclassified objects during training, improving accuracy without sacrificing speed. RetinaNet strikes a balance between accuracy and efficiency, making it suitable for real-time

tracking in complex backgrounds, as required for microrobot tracking in dense environments. As is the case with most microrobots, bacterial biohybrid microrobots exhibit various motion patterns and speeds (Fig. 1D-G, Video S1) with different detection requirements, making a single model training for all the observed behaviors ineffective. We propose a Multi-level Object Detection model, where we train a different detector model for each motility category of low, medium, and high (Fig. 2B). The detection models for each category learn specific features and parameters tailored to their unique characteristics. This approach enhances the accuracy of the object detection system for each category and improves the overall performance of the system, as discussed in Section 3.

During the training phase, the object detector receives input in the form of annotated microscopy videos with bounding boxes measuring 30×30 pixels centered around bacteria centroids. When performing inference, the object detectors predict coordinates (x, y) for bounding boxes to indicate the presence of objects as well as their width and height. These predictions are accompanied by a confidence score that reflects the model’s level of certainty regarding the detection.

2.1.3 FALSE POSITIVE PRUNER MODULE

The Multi-level Object Detector module can produce a large number of false positives as well as duplicate predictions from the three different detector models. The False Positive Pruner was implemented to remove the false positives without losing the true positive predictions. To this end, first, we combined all the detections from the three motility models and then pruned the detections based on three exclusionary criteria, as shown in Fig. 2C: (1) Bounding Box Filter for removal of predictions that are greater than a prescribed area threshold, (2) Confidence Score Filter for removal of predictions lower than a prescribed confidence threshold, and (3) Non-Maximum Suppression (NMS) Bolme et al. (2010) Filter for elimination of redundant object detections by selecting the ones with the highest confidence score and discarding the other overlapping ones. NMS operates by evaluating the Intersection over Union ($\text{IoU} = \text{AreaOfIntersection} / \text{AreaOfUnion}$) between detected bounding boxes. The threshold on IoU serves as a basis to determine whether detected boxes correspond to the same object in the NMS Filter. Beginning with the most confident detected box, NMS eliminates overlapping boxes with lower confidence scores, resulting in accurate object selection and reduced redundancy of detections. Section 2.5 describes our process for selecting the thresholds for the three filters.

2.1.4 INTERPOLATED TRACKER MODULE

Several approaches have been proposed for object tracking, including correlation filter-based Bolme et al. (2010), Kalman filter-based Bewley et al. (2016), and deep learning-based Valmadre et al. (2017); Wojke et al. (2017) methods. Simple Online Real-time Tracker (SORT) Bewley et al. (2016) is one of the most widely used tracking algorithms. It uses a combination of Kalman filtering Kalman (1961) and the Hungarian algorithm Kuhn (1955); Frank (2005) to assign detected objects to existing tracks. The Kalman filter in SORT works by recursively updating estimates of the current system state (in our case, positions of bacteria) based on the previous configuration of the state and the current measurements (*i.e.*, the detected bacteria positions in the current frame) while also taking into account the uncertainty of those measurements. SORT is simple and efficient and performs well on various tracking tasks, such as pedestrian or vehicle tracking. However, the random motion of bacteria and intermittent missing detections resulting from their 3D motion may limit the performance of this algorithm.

In our proposed method, we modify the SORT algorithm by incorporating interpolations for missed object detections. As shown in Fig. 2D, we applied the SORT algorithm to track the detected bacteria and produce tracklets from independent frame-wise predictions. We interpolated the missing detections by keeping the Kalman filter-based unmatched predictions for a given number of frames termed as the maximum age parameter (determined in Section 2.5), and dropping the predictions post that threshold. The maximum age parameter within the SORT algorithm ensures the persistence of a tracklet for a specified number of frames subsequent to a missed detection event, thereby upholding tracking continuity. If a track is missed further than the maximum age, it is discarded. Finally, we introduce the Track Length Filter to remove tracks whose length, in terms of the number of frames, does not meet a specified threshold (determined in section 2.5). This filter helps in excluding excessively short tracks that are prone to false positives.

2.2 EXPERIMENTAL METHODS

In order to evaluate the performance of MEMTrack, we recorded bacteria (*i.e.*, the micromotor in bacteria-based biohybrid microrobots) swimming behavior in collagen, as a tissue surrogate, and in an aqueous environment, as the baseline Traore et al. (2014); Zhan et al. (2022).

2.2.1 BACTERIA CULTURE

Six engineered strains of *Salmonella* Typhimurium VNP20009cheY⁺ Broadway et al. (2017) bacteria with different motile behaviors were used. Each strain was grown on a 1.5 % lysogeny broth (LB; 1 % tryptone, 0.5 % yeast extract, and 1 % sodium chloride) agar plate overnight at 37 °C. For each experiment, a single colony of the desired strain was isolated and used to inoculate 10 mL of LB media in a 125 mL smooth-bottom flask. Bacteria were cultured overnight at 37 °C and 100 RPM before being harvested and resuspended in fresh LB to a final concentration of $\sim 1.3 \times 10^6$ CFU/mL.

2.2.2 SWIMMING ASSAY IN COLLAGEN AND IN AQUEOUS MEDIUM

Bacteria motility in collagen was evaluated using an experimental setup similar to the traditional swim plate assay, in which bacteria migrate outward from a central inoculation point due to a combination of chemotaxis and growth. Collagen gel was prepared from a stock solution of collagen type I that was neutralized with 0.25 N NaOH, diluted to 5 mg/mL in LB, and supplemented with 100 μ g/mL ampicillin on ice. The cold collagen solution was pipetted into wells of a room-temperature well plate. The well plate was then immediately placed in a 37 °C incubator for 45 minutes to allow the collagen to gel. A 1 μ L aliquot of bacterial suspension, prepared as described in section 2.2.1, was introduced at the center of each collagen well. All data acquisition was performed using a Zeiss AxioObserver.Z1 inverted microscope equipped with a 40 $\times$ objective and an hSM camera (Carl Zeiss AG, Oberkochen, Germany). Phase contrast microscopy images were collected at 60 frames per second (FPS) for 2 minutes at 37 °C.

We also evaluated the motility of each strain in aqueous medium. For swimming speed analysis, a 100 μ L aliquot of the overnight bacteria culture, described in 2.2.1, was used to inoculate a fresh LB culture. The culture was grown at 37 °C for ~ 2.5 hours or until $OD_{600} = 1.0$ was reached. Subsequently, the culture was diluted 20 $\times$ in fresh LB medium. A 10 μ L aliquot of the bacterial suspension was placed between two No. 1.5 coverslips separated by a thin ring of vacuum grease. Timelapse imaging was performed as described above at 60 FPS for 10 seconds. For MEMTrack comparison with baseline models, the aqueous swimming assays were performed according to our previously developed methods Zhan et al. (2022).

2.2.3 BACTERIA TRACKING AND ANNOTATION

In order to generate training, validation, and test datasets for MEMTrack, the microscopy videos acquired in collagen and in aqueous media were imported into ImageJ Schneider et al. (2012) software. MTrackJ plugin Meijering et al. (2012) was used to label all bacteria in each frame of each video manually, and their x and y coordinates were recorded.

2.2.4 DETERMINING TRACK LENGTH AND MOTION CHARACTERISTICS

In order to determine the tracking period threshold for capturing bacteria random walk in collagen, we first tracked the bacteria for 150 frames and evaluated the diffusivity of each bacterium according to

$$D(\tau) = \frac{\sum_{t=0}^{2.5} ((x(t+\tau) - x_0(t))^2 + (y(t+\tau) - y_0(t))^2)}{4\tau} \quad (2)$$

where $\tau = 0.016$ s is the lag time between consecutive frames, and x_0 and y_0 represent the initial position.

As shown in Fig. 6, the diffusivity values plateaued at or before 1 s, indicating that a track length of 1 s (60 frames) in collagen is sufficient for capturing the bacteria random walk. Next, we used peak diffusivity values to divide bacteria into four categories based on their motility patterns and the

associated diffusivity values—No motility ($D \leq 0.075 \mu\text{m}^2\text{s}^{-1}$), low motility ($D \leq 0.25 \mu\text{m}^2\text{s}^{-1}$), medium motility ($D \leq 1 \mu\text{m}^2\text{s}^{-1}$), and high motility ($D > 1 \mu\text{m}^2\text{s}^{-1}$).

Based on the well-known motile behavior of bacteria in aqueous environments Berg (1993), we selected 0.5 s (30 frames) for testing MEMTrack’s performance on aqueous media datasets.

2.3 TRAINING, VALIDATION, AND TESTING DATASETS

MEMTrack was trained and validated using the experimental datasets in collagen, while its performance was evaluated using both the collagen and the aqueous medium datasets. A subset of the collagen data, termed the training set, was used for learning the model parameters (*i.e.*, weights and biases of the deep learning models) using gradient descent algorithms. Simultaneously, another subset of the collagen data termed the validation set, was used during training for observing the performance of the model on data outside the training set and determining the configuration of hyper-parameters (*i.e.*, parameters of the model that are not directly learned using gradient descent such as the filter thresholds of the False Positive Pruner module) that yields best validation performance. Finally, a third subset of the collagen data termed the test set, which has no overlap with the other two sets, was used to report the model performance on “unseen” data. We used the entirety of the aqueous medium data for testing the collagen-trained model in that medium. To ensure consistency and provide comprehensive training data for different motion types and background scenarios across the six bacterial strains in the collagen, we allocated two videos per strain for training (a total of 12) and reserved one video per strain for validation (a total of 6). The collagen test set consisted of 16 videos representing all bacterial strains. The aqueous medium data set contained 5 videos that were all used for testing.

2.4 EVALUATION METRICS

We used precision (equation 3) and recall (equation 4) metrics to evaluate the performance of MEMTrack quantitatively. Precision is the fraction of the true positive (TP) predictions over the sum of all predictions (TPs and false positives (FP)). It is a measure of how precise or confident the model is in tracking real objects (*e.g.*, bacteria).

$$\text{Precision} = \frac{TP}{TP + FP} \quad (3)$$

Recall, on the other hand, signifies the fraction of ground truth (GT) or actual bacteria that the model is able to recover. It is defined as the fraction of the TP over the GT (TP and false negatives (FN)).

$$\text{Recall} = \frac{TP}{TP + FN} \quad (4)$$

Furthermore, we use the F1 score Tan et al. (2018) (equation 5), which is the harmonic mean of precision and recall, as another evaluation metric to select hyper-parameters in our model architecture.

$$\text{F1 Score} = 2 \frac{\text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}} \quad (5)$$

2.5 DETERMINING THE MODEL HYPER-PARAMETERS

Model hyper-parameters include the configurational parameters of a machine learning model that are not trained directly using gradient descent but rather need to be set manually before the training begins. The hyper-parameters of the MEMTrack platform are: Area threshold (Bounding Box Filter), Confidence thresholds (Confidence Score Filter), IoU threshold (NMS Filter), Maximum Age parameter (Tracking interpolation), and Track Length threshold (Track Length Filter).

The Area threshold in the Bounding Box Filter was set to a value of 35×35 pixels. Consequently, any detected object characterized by a bounding box area exceeding 35×35 pixels was omitted from our predictions. Note that this threshold exceeds the training bounding box size criterion of 30×30 pixels to be tolerant of predictions with larger bounding box size during inference compared to training.




Figure 3: Precision and recall of the three motility detectors (labeled over each plot) for varying confidence score thresholds. Confidence score thresholds of each model based on the maximum precision and maximum F-1 Score criteria are shown as a vertical dashed line and a vertical dotted line, respectively.

We experimented with two possible criteria for selecting the confidence score thresholds of the three-level object detector using the validation set—maximum precision criterion and maximum F1 score criterion. When the maximum precision criterion is used, the three confidence score thresholds (one for each motility level) are set to maximize the precision on the validation set for each motility level. This criterion ensures that the number of false positives is minimized in our predictions across all motility levels. However, by solely maximizing the precision, we cannot control for the number of false negatives and hence may obtain lower recall values. On the other hand, the maximum F1 score criterion ensures that the harmonic mean of precision and recall values are high for all three motility models, thus making a trade-off between reducing false positives and improving recall of the GT bacteria. The choice of maximum precision versus maximum F-1 score criteria depends on the user’s priorities. While the maximum F-1 score criterion would yield a higher recall than the maximum precision criterion, it would do so at the cost of a lower precision value (or, equivalently, more false positive tracks), requiring additional manual intervention to remove the false positive tracks.

Fig. 3 shows the precision-recall curves on the validation set of the three motility levels for the Multi-level Object Detector, where the threshold settings for the two criteria are indicated as vertical lines. For the remainder of this paper, all results from MEMTrack model were obtained using the max precision criterion unless otherwise indicated, since our focus was to obtain predictions of bacteria tracks with high precision without requiring additional manual post-processing steps to remove false positive tracks.

For the NMS Filter, we used an IoU threshold of 0.7, meaning that if two detected boxes overlap more than 70%, only the one with the highest confidence score is retained. Note that the threshold is greater than 0.5 since we wanted to retain GT objects that appear in close vicinity (with overlap between 0.5-0.7) as separate detections.

The maximum Age threshold for the SORT Tracking was selected to be 25 frames, indicating that interpolation for absent detections is conducted for up to 25 missed frames; subsequently, tracks beyond this threshold are terminated and discarded. Note that the IoU threshold and Maximum Age threshold settings were manually selected from a pool of other candidate settings to yield the best performance (precision and recall) on the validation set. However, we did not exhaustively search for the best setting of these thresholds using a comprehensive grid search, which we believe could further improve our results.

For the Track Length Filter, we used a threshold of 60 minimum frames to capture bacteria random walk in collagen, as described in 2.2.4. Given the significantly shorter randomization time of bacteria in aqueous environments Leaman et al. (2020), we used a minimum track length threshold of 30 frames. Figure 7 shows how varying the threshold for minimum track length impacts the precision and recall for our method and the baselines. Increasing this threshold will increase the precision while losing out on many bacteria that have shorter tracks due to them being missed in the detection phase, potentially due to very low-frequency motion.

3 RESULTS AND DISCUSSION

3.1 EVALUATION OF DETECTION BIAS

We first inquired if MEMTrack is able to detect and track bacteria from each of the four motility sub-populations equally well. Such performance capability is crucial to the accurate representation of population-scale behavior. To evaluate the presence of any biases, we segregated the GT bacteria and the TP tracked bacteria into 4 sub-populations based on their motility (refer to Section 2.2.4). Table 1 shows the number of GT and TP bacteria in each of the four categories. The % Detected column shows the fraction of the GT bacteria outputted from the MEMTrack. It can be seen that MEMTrack has a comparable detection rate for nearly all categories. The lower detection rate for the medium motility group may be attributed to the small number of bacteria in this sub-population, which leads to significant fluctuation in the calculation of % Detected.

Table 1: Fraction of bacteria detected by MEMTrack as a function of motility sub-population

Sub-population	Ground Truth (GT)	True Positive (TP)	% Detected
No motility	17	8	47
Low motility	140	64	46
Medium motility	16	5	31
High motility	51	25	49

3.2 MEMTRACK PERFORMANCE IN COLLAGEN

We next evaluated the MEMTrack performance in the detection and tracking of bacteria in collagen. (Figure 4, Video S2) shows the precision and recall values attained at each step of the pipeline for the test datasets in collagen. The precision significantly increases after each of the filters of the False Positive Pruner module with a trade-off in the recall. The Interpolated Tracker module further boosts the recall, although at the cost of adding some FPs or a slight reduction in precision. Informed by the physics of bacterial motion in collagen (Section 2.2.4), the Track Length Filter removes tracks that are shorter than 60 frames and achieves a reasonable balance between precision and recall.

To determine MEMTracks ability to accurately describe the average speed of the bacterial population, we compared the average speed of the GT bacteria $8.4 \pm 0.4 \mu\text{m/s}$ with the tracked bacteria $9.7 \pm 0.6 \mu\text{m/s}$, as shown in Fig. 5A. A t-test analysis did not indicate any statistically significant difference between the two groups; thus, MEMTrack is able to successfully track bacteria and ascertain population-scale motility speed in the dense collagen environment.




Figure 4: Precision and recall values at each step of the MEMTrack pipeline for the collagen test data.




Figure 5: (A) The average speed of the GT bacteria tracked by a human (blue) and the trajectories produced by MEMTrack (pink), irrespective of their match to the GT in collagen (N = 224 for GT, N = 153 for Tracked). (B) The average speed of the GT bacteria tracked by a human (blue) and the trajectories produced by MEMTrack (pink), irrespective of their match to the GT in liquid medium (N = 65 for GT, N = 38 for Tracked).

3.3 MEMTRACK PERFORMANCE IN AQUEOUS ENVIRONMENT

Next, we set out to determine the applicability of our model, trained on collagen data, to other media. We evaluated the performance of MEMTrack on bacteria motility data that was collected in an aqueous medium (Video S3).

To determine MEMTrack’s ability to accurately describe the average speed of the bacterial population, we compared the average speed of the GT bacteria, $17.6 \pm 1.7 \mu\text{m/s}$, with that of the tracked bacteria, $22.9 \pm 2.4 \mu\text{m/s}$, as shown in Fig. 5B. A t-test analysis did not indicate any statistically significant difference between the two groups; thus, MEMTrack is able to track bacteria successfully and ascertain their population-scale motility speed in the aqueous environment as well.

3.4 COMPARING MEMTRACK WITH BASELINE MODELS

Lastly, we aimed to compare our approach with other existing methods as baselines. Due to the lack of existing models for tracking micro/nano-scale objects in dense environments, we opted for the two most commonly used tracking methods: (1) Trackmate Ershov et al. (2021), and (2) MosaicSuite - Particle Tracker Sbalzarini & Koumoutsakos (2005). We chose these models because they are user-friendly, open source, freely available on the widely used platform FIJI Schindelin et al. (2012), and commonly used for micro/nano-scale object tracking. These tools also have the option to use ML and semi-automated tracking features. Trackmate7 Ershov et al. (2021) uses a tracking algorithm based on the Linear Assignment Problem (LAP) tracker Fukai & Kawaguchi (2023), which assigns objects to tracks based on their similarity. MosaicSuite - Particle Tracker Sbalzarini & Koumoutsakos (2005) presents a feature point tracking algorithm that is self-initializing, discriminates spurious detections, and can handle temporary occlusion as well as particle appearance and disappearance from the image region.

We used all the predictions from the benchmark models and compared them against the GT data. We then filtered out tracklets that have a minimum length of 60 frames in collagen and 30 frames in

Table 2: Comparison of MEMTrack performance in collagen against two benchmark models.

Method	Tracking		Post Track Length Filter (60)	
	Precision	Recall	Precision	Recall
Trackmate	0.43	0.20	0.45	0.17
MosaicSuite	0.03	0.94	0.07	0.71
MEMTrack (this work)	0.74	0.55	0.77	0.48

Table 3: Comparison of MEMTrack performance in liquid media against two benchmark models.

Method	Tracking		Post Track Length Filter (30)	
	Precision	Recall	Precision	Recall
Trackmate	0.84	0.39	0.83	0.37
MosaicSuite	0.27	0.94	0.39	0.86
MEMTrack -Max Precision	0.94	0.20	0.94	0.15
MEMTrack -Max F-1 Score	0.88	0.40	0.94	0.35

liquid. Results for the comparison of the MEMTrack model and the benchmark methods in collagen and aqueous environments are reported in Tables 2 and 3, respectively.

Considering a Track Length Filter of 60 frames (or 1 sec), MEMTrack achieved a precision of 77% precision and a recall of 48% in collagen (Table 2). In contrast, Trackmate had significantly lower precision and recall values of 45% and 17%, respectively. MosaicSuite showed the highest recall at 71% but at the cost of a significantly lower precision of 7% (*i.e.*, only 7% of tracks are TPs), making its results unusable without extensive manual post-processing for filtering out the FPs. Overall, our approach can provide significantly better precision than the benchmark methods with reasonably high recall.

Table 3 compares the performance of the pre-trained (in collagen) MEMTrack with the same Trackmate and MosaicSuite benchmark models in aqueous media. Consistent with the results in collagen, MEMTrack performed better than both benchmark models with a precision of 94% and a recall of 15% on the maximum precision setting and a precision of 94%, and a recall of 35% on the maximum F-1 score setting. Trackmate produced somewhat comparable results with a precision of 83% and a recall of 37%. MosaicSuite’s precision for liquid data was 39%, which was significantly improved compared to that of collagen data (7%); nonetheless, its significantly low precision necessitates manual intervention to remove false positives. As expected, the application of the Track Length Filter improved the precision in most cases and decreased the recall (Fig. 7). Altogether, these results demonstrate that our model can be used to effectively track micro/nano-scale objects in liquid media without additional training. MEMTrack performance in liquid media can potentially be further improved if we re-train models from scratch on liquid data instead of using the pre-trained Collagen models.

4 CONCLUSION

In this work, we present MEMTrack, an automated pipeline for detecting and tracking microrobots in dense and low-contrast environments, such as collagen. This is a particularly challenging problem given the lack of visual features distinguishing the foreground objects from the background. Our approach leverages synthetic motion features, the RetinaNet object detection model and a modified SORT tracking algorithm with interpolation to achieve robust results against different background media and with high precision. Our test results demonstrate that the proposed pipeline can robustly substitute the tedious task of manually tracking microrobots for the prediction of population-scale speed values. Moreover, we show the broader applicability of our method by using it to track bacteria in liquid media. Our pipeline also has limitations that can be improved in future research. For instance, instead of using the SORT tracking algorithm, which only uses the positions of detected bacteria to perform tracking, deep learning-based tracking methods such as DeepSORT Wojke et al. (2017) can be used in future extensions to take into account the visual features of detected bacteria in every frame to perform tracking. Learning low-frequency motion patterns of bacteria from longer periods can also be explored to enhance accuracy in tracking no and low-motility bacteria categories. Finally, the tracking algorithm can be improved by including physics-based knowledge of bacteria motion in varying environments to guide the Interpolated Tracker module.

5 ACKNOWLEDGEMENTS

The authors acknowledge Behkam Lab member, Ying Zhan, for sharing previously published bacteria aqueous swimming data to test MEMTrack performance in speed predictions. This research was supported in part by NSF grants CBET-2133739 and CBET-1454226 to BB, 4-VA grant to BB, and NSF grant IIS-2107332 to AK. Access to computing resources was provided by the Advanced Research Computing (ARC) Center at Virginia Tech.

REFERENCES

- Azaam Aziz, Mariana Medina-Sánchez, Jing Claussen, and Oliver G Schmidt. Real-time optoacoustic tracking of single moving micro-objects in deep phantom and ex vivo tissues. *Nano letters*, 19(9):6612–6620, 2019.
- Azaam Aziz, Stefano Pane, Veronica Iacovacci, Nektarios Koukourakis, Jurgen Czarske, Arianna Menciassi, Mariana Medina-Sánchez, and Oliver G Schmidt. Medical imaging of microrobots: Toward in vivo applications. *ACS nano*, 14(9):10865–10893, 2020.
- Jang Pyo Bae, Siyeop Yoon, Malinda Vania, and Deukhee Lee. Three dimensional microrobot tracking using learning-based system. *International Journal of Control, Automation and Systems*, 18:21–28, 2020.
- Athanasios D Balomenos, Panagiotis Tsakanikas, Zafiro Aspidou, Anastasia P Tampakaki, Konstantinos P Koutsoumanis, and Elias S Manolakos. Image analysis driven single-cell analytics for systems microbiology. *BMC systems biology*, 11:1–21, 2017.
- Howard C Berg. *Random walks in biology*. Princeton University Press, 1993.
- Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Uptcroft. Simple online and realtime tracking. In *2016 IEEE International Conference on Image Processing (ICIP)*, pp. 3464–3468. IEEE, 2016.
- Christopher S Bjornsson, Gang Lin, Yousef Al-Kofahi, Arunachalam Narayanaswamy, Karen L Smith, William Shain, and Badrinath Roysam. Associative image analysis: a method for automated quantification of 3d multi-parameter images of brain tissue. *Journal of neuroscience methods*, 170(1):165–178, 2008.
- Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*, 2020.
- David S. Bolme, J. Ross Beveridge, Bruce A. Draper, and Yui Man Lui. Visual object tracking using adaptive correlation filters. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 2544–2550, 2010. doi: 10.1109/CVPR.2010.5539960.
- Karim Botros, Mohammad Alkhatib, David Folio, and Antoine Ferreira. Usmicromagset: Using deep learning analysis to benchmark the performance of microrobots in ultrasound images. *IEEE Robotics and Automation Letters*, 8:1–8, 2023.
- Katherine M Broadway, Seungbeum Suh, Bahareh Behkam, and Birgit E Scharf. Optimizing the restored chemotactic behavior of anticancer agent salmonella enterica serovar typhimurium vnp20009. *Journal of biotechnology*, 251:76–83, 2017.
- J.A. Cornwell, J. Li, S. Mahadevan, J.S. Draper, G.L. Joun, H. Zoellner, N.S. Asli, R.P. Harvey, and R.E. Nordon. Trackpad: Software for semi-automated single-cell tracking and lineage annotation. *SoftwareX*, 11:100440, 2020. ISSN 2352-7110. doi: <https://doi.org/10.1016/j.softx.2020.100440>. URL <https://www.sciencedirect.com/science/article/pii/S2352711019302390>.
- Fabrice De Chaumont, Stéphane Dallongeville, Nicolas Chenouard, Nicolas Hervé, Sorin Pop, Thomas Provoost, Vannary Meas-Yedid, Praveen Pankajakshan, Timothée Lecomte, Yoann Le Montagner, et al. Icy: an open bioimage informatics platform for extended reproducible research. *Nature methods*, 9(7):690–696, 2012.

- Dmitry Ershov, Minh-Son Phan, Joanna W Pylvänäinen, Stéphane U Rigaud, Laure Le Blanc, Arthur Charles-Orszag, James RW Conway, Romain F Laine, Nathan H Roy, Daria Bonazzi, et al. Bringing trackmate into the era of machine-learning and deep-learning. *BioRxiv*, pp. 2021–09, 2021.
- Linqing Feng, Yingke Xu, Yi Yang, and Xiaoxiang Zheng. Multiple dense particle tracking in fluorescence microscopy images based on multidimensional assignment. *Journal of structural biology*, 173(2):219–228, 2011.
- András Frank. On kuhn’s hungarian method—a tribute from hungary. *Naval Research Logistics (NRL)*, 52(1):2–5, 2005.
- Yohsuke T Fukai and Kyogo Kawaguchi. Laptrack: linear assignment particle tracking with tunable metrics. *Bioinformatics*, 39(1):btac799, 2023.
- Lucia Gardini, Marco Capitanio, and Francesco S Pavone. 3d tracking of single nanoparticles and quantum dots in living cells by out-of-focus imaging with diffraction pattern recognition. *Scientific Reports*, 5(1):1–10, 2015.
- Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1440–1448, 2015.
- Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Region-based convolutional networks for accurate object detection and segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(1):142–158, 2016. doi: 10.1109/TPAMI.2015.2437384.
- David E Hernandez, Steven W Chen, Elizabeth E Hunter, Edward B Steager, and Vijay Kumar. Cell tracking with deep learning and the viterbi algorithm. In *2018 International Conference on Manipulation, Automation and Robotics at Small Scales (MARSS)*, pp. 1–6. IEEE, 2018.
- Berthold Horn and Brian Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 08 1981. doi: 10.1016/0004-3702(81)90024-2.
- Jialin Jiang, Zhengxin Yang, Antoine Ferreira, and Li Zhang. Control and autonomy of microrobots: Recent progress and perspective. *Advanced Intelligent Systems*, 4(5):2100279, 2022.
- RE Kalman. A new approach to liner filtering and prediction problems, transaction of asme. *Journal of Basic Engineering*, 83(1):95–108, 1961.
- Pasi Kankaanpää, Lassi Paavolainen, Silja Tiitta, Mikko Karjalainen, Joacim Päivärinne, Jonna Nieminen, Varpu Marjomäki, Jyrki Heino, and Daniel J White. Bioimagexd: an open, general-purpose and high-throughput image-processing platform. *Nature methods*, 9(7):683–689, 2012.
- Johannes Klein, Stefan Leupold, Ilona Biegler, Rebekka Biedendieck, Richard Münch, and Dieter Jahn. Tlm-tracker: software for cell segmentation, tracking and lineage analysis in time-lapse microscopy movies. *Bioinformatics*, 28(17):2276–2277, 2012.
- H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955. doi: <https://doi.org/10.1002/nav.3800020109>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/nav.3800020109>.
- EJ Leaman, A Sahari, MA Traore, BQ Geuther, CM Morrow, and B Behkam. Data-driven statistical modeling of the emergent behavior of biohybrid microrobots. *APL Bioengineering*, 4(1):016104–016104, 2020.
- Dengfeng Li, Chao Liu, Yuanyuan Yang, Lidai Wang, and Yajing Shen. Micro-rocket robot with all-optic actuating and tracking in blood. *Light: Science & Applications*, 9(1):84, 2020.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pp. 2980–2988, 2017.
- Bruce D Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. In *IJCAI’81: 7th international joint conference on Artificial intelligence*, volume 2, pp. 674–679, 1981.

- Claire McQuin, Allen Goodman, Vasiliy Chernyshev, Lee Kamentsky, Beth A Cimini, Kyle W Karhohs, Minh Doan, Liya Ding, Susanne M Rafelski, Derek Thirstrup, et al. Cellprofiler 3.0: Next-generation image processing for biology. *PLoS biology*, 16(7):e2005970, 2018.
- Shalin B Mehta, Molly McQuilken, Patrick J La Riviere, Patricia Occhipinti, Amitabh Verma, Rudolf Oldenbourg, Amy S Gladfelter, and Tomomi Tani. Dissection of molecular assembly dynamics by tracking orientation and position of single molecules in live cells. *Proceedings of the National Academy of Sciences*, 113(42):E6352–E6361, 2016.
- Erik Meijering, Oleh Dzyubachyk, and Ihor Smal. Chapter nine - methods for cell and particle tracking. In P. Michael conn (ed.), *Imaging and Spectroscopic Analysis of Living Cells*, volume 504 of *Methods in Enzymology*, pp. 183–200. Academic Press, 2012. doi: <https://doi.org/10.1016/B978-0-12-391857-4.00009-4>. URL <https://www.sciencedirect.com/science/article/pii/B9780123918574000094>.
- 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.
- Jay M. Newby, Alison M. Schaefer, Phoebe T. Lee, M. Gregory Forest, and Samuel K. Lai. Convolutional neural networks automate detection for tracking of submicron-scale particles in 2d and 3d. *Proceedings of the National Academy of Sciences*, 115(36):9026–9031, 2018. doi: [10.1073/pnas.1804420115](https://doi.org/10.1073/pnas.1804420115). URL <https://www.pnas.org/doi/abs/10.1073/pnas.1804420115>.
- S Pane, V Iacovacci, E Sinibaldi, and A Menciassi. Real-time imaging and tracking of microrobots in tissues using ultrasound phase analysis. *Applied Physics Letters*, 118(1):014102, 2021.
- Eli Pottash, Ryan McKay, Chelsea Virgile, Hana Ueda, and William Bentley. Tumblescore: Run and tumble analysis for low frame-rate motility videos. *BioTechniques*, 62:31–36, 01 2017. doi: [10.2144/000114493](https://doi.org/10.2144/000114493).
- Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7263–7271, 2017.
- Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 779–788, 2016. doi: [10.1109/CVPR.2016.91](https://doi.org/10.1109/CVPR.2016.91).
- Michael J Rust, Melike Lakadamyali, Boerries Brandenburg, and Xiaowei Zhuang. Single-particle virus tracking. *Cold Spring Harbor Protocols*, 2011(9):1978–1987, 2011.
- I.F. Sbalzarini and P. Koumoutsakos. Feature point tracking and trajectory analysis for video imaging in cell biology. *Journal of Structural Biology*, 151(2):182–195, 2005. ISSN 1047-8477. doi: <https://doi.org/10.1016/j.jsb.2005.06.002>. URL <https://www.sciencedirect.com/science/article/pii/S1047847705001267>.
- Johannes Schindelin, Ignacio Arganda-Carreras, Erwin Frise, Verena Kaynig, Mark Longair, Tobias Pietzsch, Stephan Preibisch, Curtis Rueden, Stephan Saalfeld, Benjamin Schmid, et al. Fiji: an open-source platform for biological-image analysis. *Nature methods*, 9(7):676–682, 2012.
- Caroline A Schneider, Wayne S Rasband, and Kevin W Eliceiri. NIH image to ImageJ: 25 years of image analysis. *Nat. Methods*, 9(7):671–675, July 2012.
- Julian Schwanbeck, Ines Oehmig, Jérôme Dretzke, Andreas E Zautner, Uwe Groß, and Wolfgang Bohne. Ysmr: a video tracking and analysis program for bacterial motility. *BMC bioinformatics*, 21(1):1–8, 2020.
- Fernando Soto, Jie Wang, Rajib Ahmed, and Utkan Demirci. Medical micro/nanorobots in precision medicine. *Advanced Science*, 7(21):2002203, 2020.
- Christoph Spahn, Estibaliz Gómez-de Mariscal, Romain F Laine, Pedro M Pereira, Lucas von Chamier, Mia Conduit, Mariana G Pinho, Guillaume Jacquemet, Séamus Holden, Mike Heilemann, et al. Deepbacs for multi-task bacterial image analysis using open-source deep learning approaches. *Communications Biology*, 5(1):688, 2022.

- David R Stirling, Madison J Swain-Bowden, Alice M Lucas, Anne E Carpenter, Beth A Cimini, and Allen Goodman. Cellprofiler 4: improvements in speed, utility and usability. *BMC bioinformatics*, 22:1–11, 2021.
- Stella Stylianidou, Connor Brennan, Silas B Nissen, Nathan J Kuwada, and Paul A Wiggins. Supersegger: robust image segmentation, analysis and lineage tracking of bacterial cells. *Molecular microbiology*, 102(4):690–700, 2016.
- Pang-Ning Tan, M Steinbach, A Karpatne, and V Kumar. *Introduction to data mining (2nd edition)*. Pearson Education Limited, 2018.
- Jean-Yves Tinevez, Nick Perry, Johannes Schindelin, Genevieve M. Hoopes, Gregory D. Reynolds, Emmanuel Laplantine, Sebastian Y. Bednarek, Spencer L. Shorte, and Kevin W. Eliceiri. Trackmate: An open and extensible platform for single-particle tracking. *Methods*, 115:80–90, 2017. ISSN 1046-2023. doi: <https://doi.org/10.1016/j.ymeth.2016.09.016>. URL <https://www.sciencedirect.com/science/article/pii/S1046202316303346>. Image Processing for Biologists.
- Mehmet Efe Tiryaki, Sinan Ozgun Demir, and Metin Sitti. Deep learning-based 3d magnetic microrobot tracking using 2d mr images. *IEEE Robotics and Automation Letters*, 7(3):6982–6989, 2022.
- Mahama A Traore, Carmen M Damico, and Bahareh Behkam. Biomanufacturing and self-propulsion dynamics of nanoscale bacteria-enabled autonomous delivery systems. *Applied Physics Letters*, 105(17):173702, 2014.
- Pascal Vallotton, Antoine M Van Oijen, Cynthia B Whitchurch, Vladimir Gelfand, Leslie Yeo, Georgios Tsiavaliaris, Stephanie Heinrich, Elisa Dultz, Karsten Weis, and David Grünwald. Diatrack particle tracking software: Review of applications and performance evaluation. *Traffic*, 18(12):840–852, 2017.
- Jack Valmadre, Luca Bertinetto, Joao Henriques, Andrea Vedaldi, and Philip HS Torr. End-to-end representation learning for correlation filter based tracking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2805–2813, 2017.
- Hilde M Van Der Schaar, Michael J Rust, Chen Chen, Heidi van der Ende-Metselaar, Jan Wilschut, Xiaowei Zhuang, and Jolanda M Smit. Dissecting the cell entry pathway of dengue virus by single-particle tracking in living cells. *PLoS pathogens*, 4(12):e1000244, 2008.
- Ben Wang, Kostas Kostarelos, Bradley J Nelson, and Li Zhang. Trends in micro-/nanorobotics: materials development, actuation, localization, and system integration for biomedical applications. *Advanced Materials*, 33(4):2002047, 2021.
- CY Wang, A Bochkovskiy, and HYM Liao. Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. arxiv 2022. *arXiv preprint arXiv:2207.02696*, 2022.
- Qianqian Wang and Li Zhang. Ultrasound imaging and tracking of micro/nanorobots: From individual to collectives. *IEEE Open Journal of Nanotechnology*, 1:6–17, 2020.
- Quanli Wang, Jarad Niemi, Chee-Meng Tan, Lingchong You, and Mike West. Image segmentation and dynamic lineage analysis in single-cell fluorescence microscopy. *Cytometry Part A: The Journal of the International Society for Advancement of Cytometry*, 77(1):101–110, 2010.
- Victoria A Webster-Wood, Maria Guix, Nicole W Xu, Bahareh Behkam, Hirotaka Sato, Deblina Sarkar, Samuel Sanchez, Masahiro Shimizu, and Kevin Kit Parker. Biohybrid robots: recent progress, challenges, and perspectives. *Bioinspiration and Biomimetics*, 18(1):015001, 2023.
- Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *2017 IEEE International Conference on Image Processing (ICIP)*, pp. 3645–3649. IEEE, 2017.

Jun Xie, Shahid Khan, and Mubarak Shah. Automatic tracking of escherichia coli bacteria. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2008: 11th International Conference, New York, NY, USA, September 6-10, 2008, Proceedings, Part I 11*, pp. 824–832. Springer, 2008.

Ying Zhan, Austin Fergusson, Lacey R McNally, Richey M Davis, and Bahareh Behkam. Robust and repeatable biofabrication of bacteria-mediated drug delivery systems: Effect of conjugation chemistry, assembly process parameters, and nanoparticle size. *Advanced Intelligent Systems*, 4 (3):2100135, 2022.

A SUPPLEMENTARY VIDEOS

All supplementary video are publicly available at this link: (<https://doi.org/10.5281/zenodo.10001477>). This repository contains 3 videos as detailed below:

Video S1: Video shows the movement of representative bacterial biomotors from each of the four motility subpopulations in collagen. Video slowed down by 2x

Video S2: Videos show ground truth annotation (red) by manual tracking and MEMTrack-enabled automated detection and tracking of bacteria in collagen. All scale bars are 20 μm and timestamp indicates s:ms. Video slowed down by 2x

Video S3: Videos show ground truth annotation (red) by manual tracking and MEMTrack-enabled automated detection and tracking of bacteria in aqueous media. All scale bars are 20 μm and timestamp indicates s:ms. Video slowed down by 2x

B SUPPLEMENTARY FIGURES




Figure 6: Bacteria diffusivity in collagen plateaus after approximately 1s.




Figure 7: Effect of the Track Length Filter threshold on the precision and recall values for the liquid dataset. The top row shows the results with max-precision criteria, and the bottom row shows the results for max F1 criteria.