

Instance Segmentation for the Quantification of Microplastic Fiber Images

Viktor Wegmayr *

Aytunc Sahin *

Björn Sæmundsson

Joachim M. Buhmann

Department of Computer Science, ETH Zurich, Switzerland

{vwegmayr, aytunc.sahin, jbuhmann}@inf.ethz.ch

bjoernsm@student.ethz.ch

Abstract

Microplastics pollution has been recognized as a serious environmental concern, with research efforts underway to determine primary causes. Experiments typically generate bright-field images of microplastic fibers that are filtered from water. Environmental decision making in process engineering critically relies on accurate quantification of microplastic fibers in these images. To satisfy the required standards, images are often analyzed manually, resulting in a highly tedious process, with thousands of fiber instances per image. While the shape of individual fibers is relatively simple, it is difficult to separate them in highly crowded scenes with significant overlap. We propose a fiber instance detection pipeline, which decomposes the fiber detection and segmentation into manageable sub-problems. Well separated instances are identified with robust image processing techniques, such as adaptive thresholding, and morphological skeleton analysis, while tangled fibers are separated by an algorithm based on deep pixel embeddings. Moreover, we present a modified Intersection-over-Union metric as a more appropriate similarity metric for elongated shapes. Our approach improves significantly on out-of-sample data, in particular for difficult cases of intersecting fibers.

1. Introduction

Microplastics: Plastics products are an ubiquitous commodity in our society. Pollution with macroscopic plastics has caused significant stress on our environment, as they are highly resistant to natural degradation. Much less attention has been devoted to microscopic plastic particles in water and its accumulation in animal and human organisms; concerns have been raised only recently [5].

To develop effective solutions to reduce microplastic pollution, researchers have investigated its origins, such as the everyday washing of synthetic textiles [9, 10]. Risk-assessment, as in [2], depends on reliable statistical analysis of samples, which are typically microscopy images of



Figure 1: Example crop of a filter image containing microplastic fibers.

microplastic fibers filtered from water samples, as shown in Figure 1. While the fiber mass density can be efficiently weighed with a scale, the quantification of other key figures, such as fiber count density, and fiber length distribution, requires tedious identification and measurement of thousands of individual fiber instances.

Curvilinear Objects: In fact, statistical quantification of object instances from image data poses a common problem in the life sciences. While there exists a multitude of objects of interest, we focus on thin, and elongated microplastic fibers. However, curvilinear objects are an important subclass in general.

For instance, the *C. elegans* worm has served as an essential model organism for research in developmental biology [1]. Another example are amyloid fibrils of prion proteins, which play an important role in the understanding of neurodegenerative diseases [3].

We also like to mention short-fiber-reinforced polymers, which are basically a three-dimensional collection of densely packed, but non-intersecting, straight sticks. Their analysis helps material scientists to quantify the effects of stick length and orientation on their tensile strength [6].

*Equal contribution

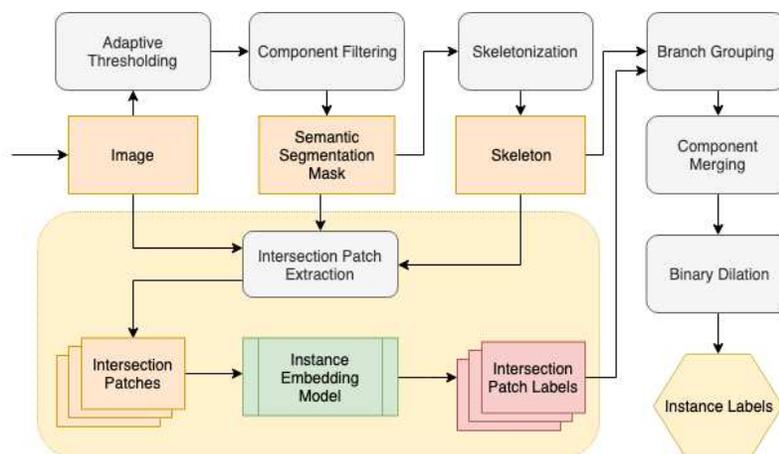


Figure 2: Flowchart of the proposed fiber instance segmentation pipeline. The Deep Pixel Embedding module is highlighted by the yellow square box.

Instance Segmentation: In general, automatic instance detection and segmentation in images define image processing tasks for computer vision. Recently, it received increased attention due to the progress of deep learning in semantic segmentation, and object detection.

In particular, generic methods such as the celebrated Mask-RCNN [8] provide a solution for a wide range of objects. Even though fiber instances appear to be considerably less complex than objects commonly encountered in computer vision tasks, they present challenges due to complex instance configurations. A typical fiber scene contains hundreds to thousands of instances, whereas natural scenes usually contain a few dozens of instances. While this complication can be reduced by separation into independent parts, the remaining clusters still contain a sizable number of entangled instances that are highly crowded and overlap significantly. This situation is frequently recognized as a cause of failure for generic instance segmentation methods [11]. In particular, it aggravates the performance of methods based on bounding boxes, because each flat box will cover many thin instances.

Lastly, the imaging mechanism of bright-field microscopy introduces additional ambiguity at instance intersections, since objects are translucent, and this property renders it e.g. difficult to determine occlusions and depth ordering.

Contributions: We study the problem of instance segmentation in microscopy images containing many overlapping, fiber-like instances. A robust image processing pipeline similar to [17], which decomposes the instance segmentation problem into a sequence of easier subproblems as shown in Figure 2, is proposed. The main components are locally adaptive thresholding [18], and morphological

skeleton analysis [7] combined with branch grouping based on fiber continuity.

While this simple baseline works well for well-separated fibers, it has difficulty to resolve tangled fibers. Thus, we specifically address instance crossings with deep pixel embeddings (DPE), which is a proposal-free instance segmentation method, previously developed in the context of autonomous driving [4], and applied in polymers analysis [12].

Precision and recall of fiber detection shows that the combination of image processing with deep pixel embeddings improves the separation of difficult out-of-sample fiber clusters. In addition, we introduce a modified Intersection-over-Union score which is a more faithful measure for the similarity between elongated shapes.

2. Related Work

Little image processing work exists that specifically addresses the visual segmentation of microplastic fibers. The work of [13] uses adaptive thresholding, and connected component analysis to identify instances. However, their instances appear very well separated, which avoids many of the complications that we face in our image samples (Figure 1).

Related work in developmental biology has been devoted to the automatic analysis of *C. elegans*. The work of [17] first segments all worm pixels via thresholding. This binary mask is transformed into a morphological skeleton, which is used to identify endpoints, branches, and intersections. Crossings are resolved by connecting branches with the lowest mutual angle relative to the junction point.

The work of [15] is also based on image skeletonization. However, they apply a score-guided graph search method to group skeleton branches. Different branch groupings

are scored greedily using a probabilistic shape model for worms.

Complementary to skeleton-based approaches, the work of [14] uses active contours to detect worms. This approach minimizes contour energies that contain data and shape terms, but its initialization depends on ground-truth endpoints.

The automated quantification of amyloid fibrils was shortly addressed by [20], who applies a similar approach as [17]. As to short-fiber-reinforced polymers, the recent work of [12] uses deep pixel embeddings [4] to separate fibers.

Mask-RCNN: As we use Mask-RCNN as a baseline, we shortly summarize the method originally presented by [8]. Mask-RCNN is a complex, multi-modular deep learning model that has been very successfully applied to instance segmentation on popular datasets such as Pascal VOC, and COCO.

After extraction of image features with convolutional layers, a Region Proposal Network (RPN) [16] determines regions, which potentially contain an object. These regions are transformed to a fixed-size representation by so-called RoI Align layers. Each region feature vector is then passed to two branches. The first branch computes the confidence for each possible object class to be presented in the region, and a corresponding bounding box. The second branch computes a binary segmentation for each object bounding box.

In the end, Mask-RCNN returns the bounding box, and segmentation, of the object class with the highest confidence. In our case, there is only one object class, as we deal only with fibers of the same type.

Deep Pixel Embeddings: Deep pixel embeddings (DPE) is applied to the resolution of fiber crossings; therefore we briefly sketch the method originally presented by [4].

In contrast to Mask-RCNN, DPE does not employ region proposals or bounding boxes for object detection. Moreover, DPE does not depend on any particular network architecture, but is rather characterized by its loss function. Given an input image with shape $h \times w$, each pixel is mapped to a d -dimensional embedding vector, typically by some CNN. Let us assume that a number of k instances is present in the image. The resulting feature map with shape $h \times w \times d$, and a one-hot $h \times w \times k$ label map, are then subject to three different terms of the DPE loss.

The DPE loss function is constructed such that the embedding vectors of the pixels of one instance should be similar, while embedding vectors corresponding to different instances should be dissimilar [4]. Thus, the loss function has two competing terms to achieve this objective. First, a variance term forces the embedding vectors of one instance's

pixels to be close to their common mean to make them more similar to each other. Second, a repulsion term tries to push the embedding means of different instances apart from each other. Third, a regularization term penalizes the norm of each instance mean to bound their magnitude. After training, inference is performed by clustering the predicted pixel embedding vectors, e.g. by mean shift clustering.

3. Methods

Image Processing Pipeline

Fiber instance segmentation constitutes an essential part of the image processing pipeline to quantify fiber statistics. The core components of the pipeline are summarized by the flowchart shown in Figure 2.

The raw fiber images usually contain other elements besides the fiber instances themselves, such as a label and a ruler. Thus, the first step in all our processing requires to extract the fiber area. Images are converted to gray-scale and locally adaptive thresholding removes small image perturbations by using the average value of the 5×5 neighborhood of each pixel as a threshold. Then, small connected components of the mask are filtered out that correspond to dirt.

Next, an image morphology algorithm [19] computes the morphological skeleton of each connected component to determine their endpoints, junction points, and branches. At this point, we can already sort out skeletons of well separated instances by checking for the absence of junction points. The remaining skeletons, which contain at least one junction point, require additional processing to separate the tangled instances.

In the following, we discuss two approaches to group branches at junction points, which means the matching of fiber branches incident on a given junction. Or phrased in different words, we infer how to continue a fiber across a junction point.

Branch Grouping by Fiber Continuity: As a baseline approach for branch grouping, we implement the following intuitive strategy, which is also used in [17,20]: It computes the pairwise angles between skeleton branches incident on a junction point, and pairs them according to lowest mutual angle.

Branch Grouping with Deep Pixel Embeddings: The branch grouping strategy based on continuity is solely relying on skeleton information in the very vicinity of a junction. Image analysis of junction points would benefit from information on the image context around a junction. For this purpose we use Deep Pixel Embeddings.

This algorithm is the first step in our pipeline which involves a learning mechanism, i.e. model training on exam-

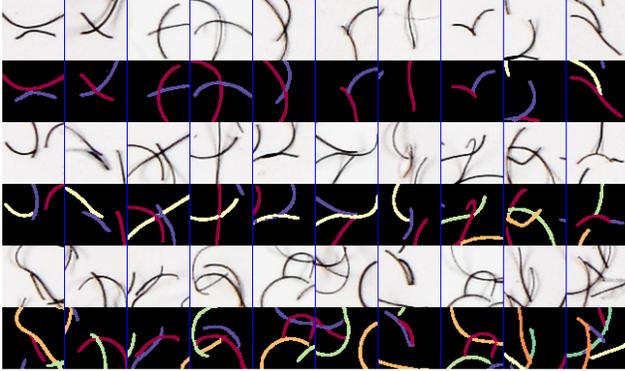


Figure 3: Instance segmentation results of Deep Pixel Embedding on intersection patches. The examples were selected randomly, and are ordered by increasing number of instances present in the input patch.

ples prior to inference. Specifically, we train the DPE model on intersection patches, and their corresponding instance labels. Intersection patches are fixed-size image crops centered on the junction, which are additionally masked with the previously obtained semantic mask.

After training, the DPE model can be used to predict instance labels on unseen intersection patches, as shown in Figure 3. These intersection patch labels are used to inform the skeleton branch grouping. This means, each skeleton branch is assigned to the majority label of the intersection area. Afterwards, pairs of branches with the same label are grouped together.

Component Merging

While the junction pixels are in principle shared between the overlapping fibers, they are not assigned to a particular instance in the case of grouping by continuity. In the case of grouping with DPE, the junction pixels are randomly assigned to one of the overlapping instances.

To join the grouped skeleton branches, we simply connect their ends with a straight line across the junction to form a connected fiber skeleton. All fiber skeletons are dilated again to their original width to obtain the final instance segmentations.

4. Experiments

Data

The image data and fiber labels used in this work were collected by researchers of the Environmental Risk Assessment and Management group at EMPA. They provided seven high resolution images (8100×5400 , $0.01\text{mm}/\text{pixel}$) of filters, which contain on average 1000 fiber instances (between 300 and 2400).

The corresponding label masks were obtained by manual

pixel annotation. Over all images, 66% of the fibers did not intersect with any other fiber, which we refer to as single fibers. Moreover, 15% intersect exactly with one other fiber, 6% intersect with two other fibers, and 3% intersect with three other fibers. We refer to fibers which intersect with one or more other fibers as “tangled” fibers.

Note, while a fiber rarely intersects with more than three other fibers, the size of fiber clusters can be much larger than that, because fibers are connected via other fibers.

Preliminary analysis

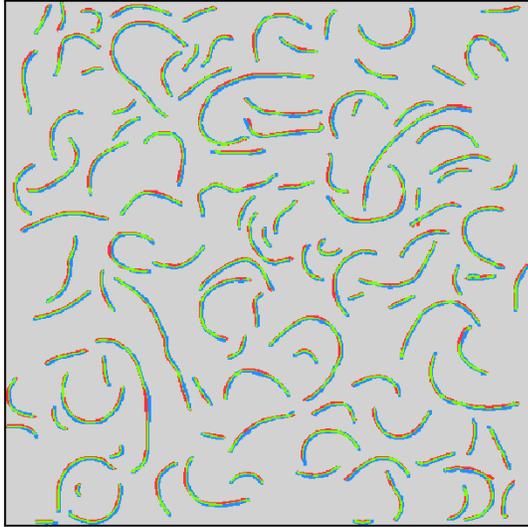
As discussed in the introduction, the analysis of fiber instances can be broken down into a sequence of easier problems. In this section we discuss particularly simple approaches to estimate the key figures *fiber count* and *fiber length distribution*. However, we will conclude that more sophisticated solutions are indeed necessary, in particular for the estimation of the fiber length distribution.

Fiber Counting by Endpoint Counting: An intermediate result of the skeleton analysis is the detection of fiber endpoints. As we know that fiber instances always have two endpoints, it is straightforward to estimate the number of fibers from the number of endpoints. This approach would avoid the identification of entire fiber instances, and reduce it to the simpler problem of endpoint detection.

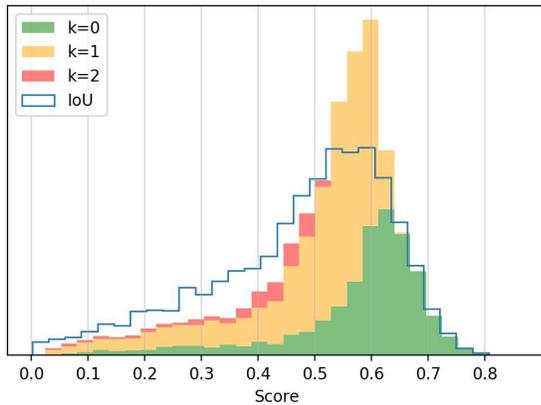
Averaged over all image samples, this approach achieves an endpoint detection precision of 0.90 ± 0.02 , and recall of 0.91 ± 0.01 . We use a distance of 10 pixels as cutoff to determine false positive (FP), true positive (TP) and false negative (FN) instances (as a reference, the mean fiber length is 76 pixels). We already observe that endpoints of tangled fibers are harder to detect than those of single fibers (FN fractions are 0.25, and 0.04, respectively). The average, relative fiber count difference was $10 \pm 2\%$. However this metric can be misleading, as it conceals the compensation of false negatives by false positives.

Length Distribution of Single Fibers: As mentioned in Section 4, two-thirds of all fibers do not intersect with any other fiber, which makes them much easier to segment. Due to the large count of fibers, we could hope to estimate the length distribution of all fibers by simply using the statistics of the single fibers. This approach would avoid the complicated resolution of tangled fibers.

However, our observations revealed the fact that tangled fibers are on average 63% longer than well-separated fibers, which invalidates this shortcut. We assume that longer fibers have a bias to get entangled more easily. Consequently, proper estimation of the length distribution needs to resolve the length of tangled fibers, too.



(a) Montage of Label/Prediction pairs



(b) IoU/mIoU histograms

Figure 4: (a) Montage of label/prediction pairs (blue/red) produced with the embedding method. Each example has $\text{IoU} < 0.40$ and $\text{mIoU} \in [0.50, 0.55]$. Overlap is indicated in green. (b) Histogram of scores over all label/prediction matches for the embedding method. Blue: IoU score. Green/Orange/Red: *Stacked* histograms of mIoU scores colored according to the argmax number of dilations k .

Evaluation of Fiber Detection

To evaluate the precision and recall of instance segmentation, we need a way to determine whether a predicted fiber is a TP or FP. For this purpose, we need to match it with one of the labeled fibers, and decide whether the match surpasses a true positive threshold to be counted as TP. The match is then classified as “good enough” for a TP instance.

The matching is achieved simply by selecting all label instance segmentations that intersect with the predicted in-

stance segmentation. If any of the intersecting labels obtains a “good enough” score with the prediction, we register a TP detection, otherwise FP. Every label that is not matched “sufficiently well” to any of the predictions, is registered as FN.

It is apparent, that the definition of “good enough” is critical for the evaluation. Commonly, the Intersection over Union (IoU) score, or Jaccard index, between label and prediction masks is invoked for this decision. While the Jaccard index is a reasonable measure of fit for flat objects encountered in standard computer vision problems, it is less appropriate for thin and elongated objects like fibers, which are characterized by a low area to border ratio. More specifically, the IoU score is not robust to slight misalignments between prediction and label, which makes it an overly pessimistic similarity score. As a consequence, predictions that are *semantically* close to a label, might still receive a very low IoU score.

To turn the common IoU into a more meaningful measure for the similarity of elongated objects, we propose to use a modified IoU:

$$\text{mIoU}(P, L) = \max_{k \in \mathbb{N}_0} S_k(P, L) \quad \text{with}$$

$$S_k(P, L) = \frac{|(P_k \cap L) \cup (P \cap L_k)|}{|P_k \cup L_k|} \quad (1)$$

where P_k , L_k denote k -time binary dilations of the prediction and label masks, respectively. Intuitively, the mIoU score is high, when a large intersection can be achieved with little dilation. Conversely, if intersection is small even with large dilation, the mIoU diminishes. This makes it tolerant against small differences, in particular shifts.

To demonstrate the difference between IoU and mIoU, we extract prediction/label pairs with IoU smaller than 0.40 and mIoU between 0.50 and 0.55. We show a *montage* of these fibers in Figure 4a. While these pairs represent perceptually good matches, their quality is underestimated by IoU. In contrast, mIoU measures their fit more faithfully with respect to the ground truth labeling.

We also report some simple formal properties of the mIoU. It is symmetric in its arguments L , and P and it is at most 1, if and only if $P = L$. Moreover, the modified IoU is always larger or equal to the usual IoU, and reduces to IoU when no dilation is performed, i.e. $S_0(P, L) = \text{IoU}(P, L)$. Lastly, the maximum operator in the definition of mIoU is well behaved for all finite-size masks, i.e. $\forall P, L : \text{argmax}_k S_k(P, L) < \infty$.

As shown in Figure 4b, the mIoU is practically often equal to the IoU, but in many cases the number of dilations is one and sometimes it exceeds two dilations.

Evaluation Results

Instance Detection: For these experiments we performed leave-one-out cross-validation on 7 high-resolution mi-

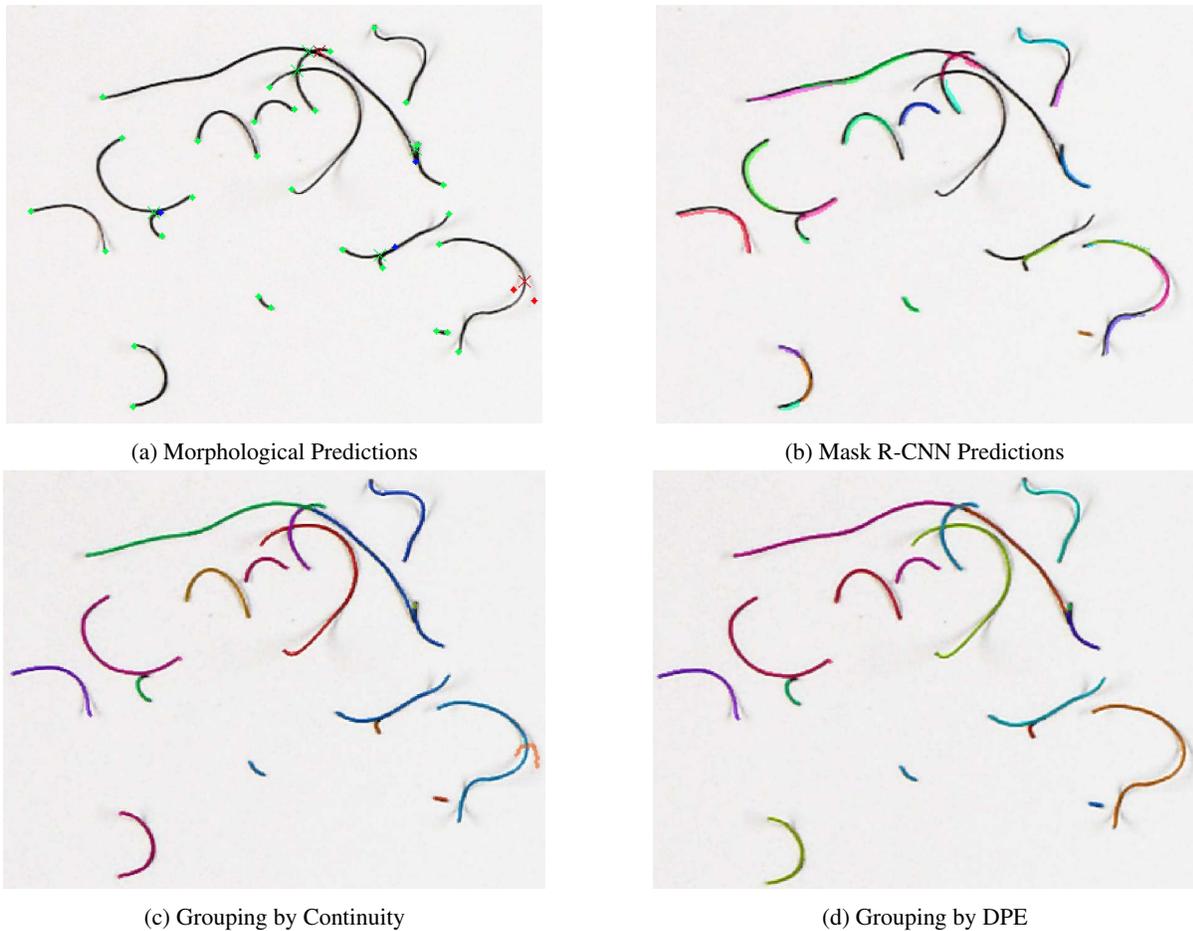


Figure 5: Comparison of instance predictions. (a) Endpoints (●) and crossings (×) detected by skeleton analysis. TP shown in green, FP in red, and FN in blue. (b,c,d) Exemplaric instance predictions of the three methods discussed in the main text. The input image is the same in each case.

crofiber images, i.e. for each fold a model is trained on six images, and tested on the remaining image. For the evaluation of instance detection, we choose an mIoU threshold of 0.50 to determine true positive (TP) predictions. The results

Table 1: Evaluation of fiber detection precision and recall on single and tangled fibers. The $mIoU_{>0.50}$ metric describes the fraction of instance predictions that achieve an mIoU score of more than 0.50.

		M-RCNN	Skeleton	Embedding
Precision	Single	0.57 ± 0.03	0.80 ± 0.03	0.83 ± 0.03
	Tangled	0.23 ± 0.07	0.30 ± 0.03	0.41 ± 0.04
Recall	Single	0.57 ± 0.02	0.80 ± 0.03	0.81 ± 0.03
	Tangled	0.23 ± 0.01	0.36 ± 0.03	0.43 ± 0.04
mIoU _{>0.50}	Single	0.57 ± 0.03	0.78 ± 0.02	0.81 ± 0.03
	Tangled	0.16 ± 0.02	0.29 ± 0.02	0.40 ± 0.04

in Table 1 clearly show the advantage of branch grouping with DPE over branch grouping by continuity. In particular, DPE significantly improves the detection quality of tangled fibers, which is also apparent in the instance prediction examples shown in Figure 5.

On single fibers, the advantage of DPE over continuity grouping is naturally small, because no junctions need to be resolved. However, in some cases, DPE can reject falsely detected junctions in single fibers, as for example shown in the bottom right corner of the examples in Figure 5.

When we use IoU_{>0.50} instead of mIoU_{>0.50} to evaluate precision and recall for the DPE approach, the results for single fiber precision/recall decrease to $0.64 \pm 0.02 / 0.63 \pm 0.02$, and tangled fiber precision/recall reduces to $0.30 \pm 0.04 / 0.32 \pm 0.04$. This decrease is expected, as IoU is naturally a rather pessimistic metric, which will reject many matches, even though they are reasonable. This undesirable property is also illustrated by the analysis in

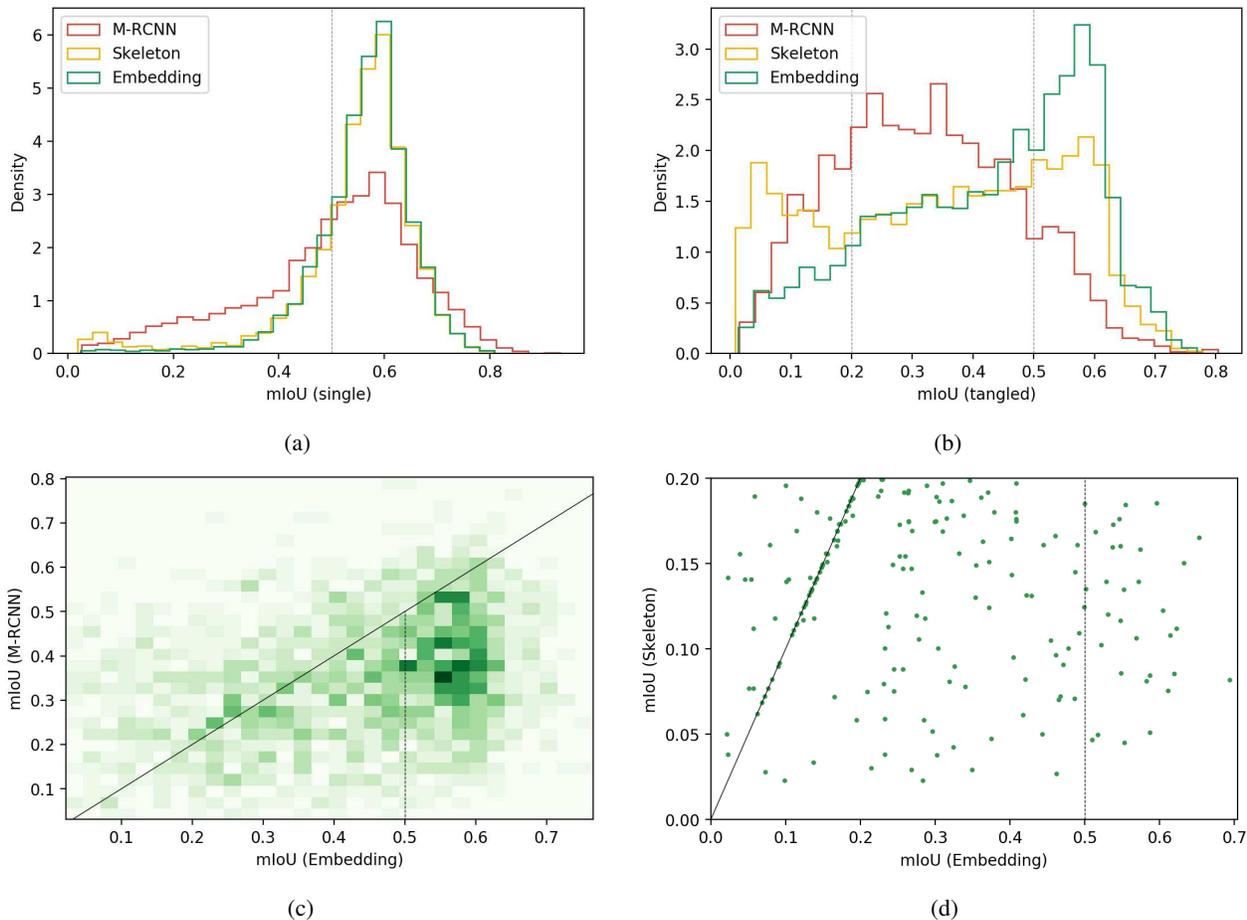


Figure 6: (a) mIoU distribution of predictions on single fibers; (b) mIoU distribution for predictions on tangled fibers; (c) joint distribution of mIoU scores between label-matched predictions of M-RCNN and Embedding; (d) Joint distribution of mIoU scores between label-matched predictions of Skeleton and Embedding. The diagonal line in (c,d) indicates the $y=x$ level.

Figure 4. In any case, using IoU would not change the ranking of the compared methods.

Instance Segmentation Quality: Moreover, we assess the quality of instance segmentation masks. As shown in the bottom row of Table 1, the approach using DPE achieves the highest fraction of instance predictions which have an mIoU larger than 0.50 with their matched label (denoted by $mIoU_{>0.50}$). Again, the difference is particularly pronounced for tangled fibers.

In addition to these bulk numbers, we also investigate the mIoU distribution of predictions in Figure 6. While the mIoU distributions for DPE and continuity grouping (*skeleton*) are virtually the same for single fibers (Figure 6a), the number of low-mIoU predictions of tangled fibers is clearly reduced for DPE (Figure 6b).

The joint distribution of the mIoU for fibers matched be-

tween Mask-RCNN and DPE (Figure 6c) shows that the average improvement is not the cumulative result of minor improvements for every case, but rather the result of better performance of DPE on difficult cases. This performance gain on difficult instances is also particularly apparent in Figure 6d, which shows that DPE significantly improves the failure cases for which continuity based grouping achieves only low mIoU.

Count and Length Results: We also report results for fiber count and length estimation, which are usually the relevant figures for practitioners, even though their diagnostic value for instance segmentation is questionable.

The mean relative count difference (averaged over images, taking into account every prediction, independent of its mIoU) is $7 \pm 2\%$ for the Mask-RCNN model, while it is $17 \pm 2\%$ for the skeleton baseline, and $7 \pm 2\%$ for the em-

bedding model. However, as mentioned already before, the count difference must be interpreted with caution, because FP predictions can compensate FN, and vice versa. For this reason, one should rather consider precision and recall as in Table 1.

The mean relative length difference (averaged over all predictions with $mIoU > 0.50$) for the Mask-RCNN model is $11 \pm 3\%$, while it is $9 \pm 3\%$ for the skeleton baseline, and $8 \pm 3\%$ for the embedding model. The difference of the length results between single and tangled fibers is minor.

To put these results into perspective, the variability between four human raters was measured for one image. The mean pairwise difference of the total count between raters was found to be $4 \pm 1\%$, while it was $6 \pm 2\%$ for the fiber length, averaged over all fibers. As expected, this result indicates that the disagreement between human raters is lower than the disagreement with the presented automatic methods. Moreover, the advantage of human annotators is more apparent for the fiber count, while it is barely significant when measuring fiber length.

5. Conclusion

We have presented a pipeline for instance segmentation of microplastic fibers in microscopy images. The image processing pipeline relies on a combination of robust image processing methods to deal with the bulk of single fibers, and Deep Pixel Embeddings to resolve the difficult tangled cases.

We demonstrate the advantage of using DPE for skeleton branch grouping over grouping strategies by continuity and perceptual organization. DPE improves the resolution of complex fiber junctions, because it considers the image context around the junction for disambiguation and visual semantics extraction.

Additionally, we propose a modified Intersection-over-Union score for elongated objects to make the assessment of their similarity more faithful. The proposed algorithmic pipeline for fiber detection and analysis serves as an example for flexible data analysis with robustness by adaptive design of the various processing steps.

Acknowledgement: The authors like to thank Yaping Cai and Bernd Nowack for fruitful discussions and for sharing the micro-fiber dataset including the image annotations.

References

- [1] A. A. Aboobaker and M. L. Blaxter. Medical significance of *caenorhabditis elegans*. *Annals of medicine*, 2000.
- [2] V. Adam, T. Yang, and B. Nowack. Toward an ecotoxicological risk assessment of microplastics: Comparison of available hazard and exposure data in freshwaters. *Environmental Toxicology and Chemistry*, 2019.
- [3] J. Adamcik, J.-M. Jung, J. Flakowski, P. D. L. Rios, G. Dietler, and R. Mezzenga. Understanding amyloid aggregation by statistical analysis of atomic force microscopy images. *Nature nanotechnology*, 2010.
- [4] B. D. Brabandere, D. Neven, and L. V. Gool. Semantic instance segmentation for autonomous driving. *CVPR Workshops*, 2017.
- [5] M. Cole, P. Lindeque, C. Halsband, and T. S. Galloway. Microplastics as contaminants in the marine environment: a review. *Marine pollution bulletin*, 2011.
- [6] S.-Y. Fu and B. Lauke. Effects of fiber length and fiber orientation distributions on the tensile strength of short-fiber-reinforced polymers. *Composites Science and Technology*, 1996.
- [7] J. Goutsias and D. Schonfeld. Morphological representation of discrete and binary images. *IEEE Transactions on Signal Processing*, 1991.
- [8] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. In *ICCV*, 2017.
- [9] E. Hernandez, B. Nowack, and D. Mitrano. Polyester textiles as a source of microplastics from households: A mechanistic study to understand microfiber release during washing. *Environmental Science and Technology*, 2017.
- [10] D. Kawecki and B. Nowack. Polymer-Specific Modeling of the Environmental Emissions of Seven Commodity Plastics As Macro- and Microplastics. *Environmental science & technology*, 2019.
- [11] A. Kirillov, E. Levinkov, B. Andres, B. Savchynskyy, and C. Rother. Instancecut: from edges to instances with multi-cut. In *CVPR*, 2017.
- [12] T. K. Konopczynski, T. Kröger, L. Zheng, and J. Hesser. Instance Segmentation of Fibers from Low Resolution CT Scans via 3D Deep Embedding Learning. In *BMVC*, 2018.
- [13] J. Lorenzo-Navarro, M. C. Santana, M. Gómez, A. Herrera, and P. A. Marín-Reyes. Automatic counting and classification of microplastic particles. In *ICPRAM*, 2018.
- [14] D. Ochoa, S. Gautama, and B. Vintimilla. Detection of individual specimens in populations using contour energies. In *International Conference on Advanced Concepts for Intelligent Vision Systems*. Springer, 2007.
- [15] T. R. Raviv, V. Ljosa, A. L. Conery, F. M. Ausubel, A. E. Carpenter, P. Golland, and C. Wählby. Morphology-guided graph search for untangling objects: *C. elegans* analysis. In *MICCAI*. Springer, 2010.
- [16] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.
- [17] N. B. Rizvandi, A. Pizurica, and W. Philips. Machine vision detection of isolated and overlapped nematode worms using skeleton analysis. In *ICIP*. IEEE, 2008.
- [18] J. J. Sauvola and M. Pietikäinen. Adaptive document image binarization. *Pattern Recognition*, 2000.
- [19] P. Soille. *Morphological image analysis*. Springer Verlag, 1999.
- [20] Y. Yin, S. Prigent, J. Torrent, H. Rezaei, D. Drasdo, and M. Doumic. Automated quantification of amyloid fibrils morphological features by image processing techniques. In *ISBI*, 2019.