AI-Driven Automated Pipeline for Tendon Fibril Segmentation and Morphological Analysis in TEM Images

Giorgio William Badrous
Department of Theoretical and Applied Sciences
University of Insubria
Varese, Lombardia, Italy
gwbadrous@studenti.uninsubria.it

Federico Ligas
Department of Theoretical and Applied Sciences
University of Insubria
Varese, Lombardia, Italy
fligas@studenti.uninsubria.it

Andrea Brambilla
Department of Medicine and Technological Innovation
University of Insubria
Varese, Lombardia, Italy
abrambilla2@studenti.uninsubria.it

Crescenzo Edoardo Mauriello
Department of Theoretical and Applied Sciences
University of Insubria
Varese, Lombardia, Italy
cemauriello@uninsubria.it

Davide Tosi
Department of Human Sciences and Innovation for the Territory
University of Insubria
Varese, Lombardia, Italy
davide.tosi@uninsubria.it

***Abstract*—The evolution of software engineering is increasingly driven by the integration of artificial intelligence and digital paradigms. In critical domains such as medical and biological imaging, modern software systems are transitioning from static, rule-based workflows to dynamic, AI-powered environments. This paper presents the design, implementation, and preliminary validation of an advanced, fully automated software pipeline for Transmission Electron Microscopy image analysis to address software engineering and analytical challenges in tendon fibril segmentation. The proposed architecture replaces a previous hybrid segmentation approach with a deep learning model based on the ViT-UNet architecture, integrating targeted preprocessing modules, advanced segmentation capable of capturing both local details and global spatial relationships, and an automated feature extraction and clustering engine. Initial experimental results show an improvement in clustering accuracy from 57.14% to 64.71%. While these findings suggest the potential of Transformer-based architectures to enhance medical image analysis, they also highlight the need for further evaluation on larger datasets to establish broader statistical significance and validate their overall effectiveness.**

***Index Terms*—deep learning, medical imaging, segmentation, vision transformers, tendon analysis**

I. Introduction

The evolution of software engineering is increasingly driven by the integration of artificial intelligence (AI) and digital paradigms. In critical domains such as medical and biological imaging, modern software systems are transitioning from static, rule-based workflows to dynamic, AI-powered environments. The development of these systems presents unique software engineering challenges: architects must ensure the reliability, robustness, and trustworthiness of AI-driven tools while managing the complexity of integrating advanced deep learning models into end-to-end analytical pipelines. Adapting traditional software engineering practices to co-evolve with these AI paradigms is essential for building scalable and verifiable solutions.

A prominent example of this challenge is the automated analysis of Transmission Electron Microscopy (TEM) images of biological tissues, specifically tendon collagen fibrils. The morphological analysis and spatial distribution of these fibrils are crucial for understanding tissue biomechanics, aging, and degenerative pathologies. Traditionally, this process relied on manual annotation, which is time-consuming and subject to human bias. Over time, semi-automated pipelines emerged to achieve better results faster, leaving professionals mainly to the task of fine-tuning parameters and applying minor corrections. However, developing a software solution to automate this task requires handling high-resolution images affected by significant noise (such as mixed Poisson-Gaussian noise), variable contrast, motion artifacts, and complex, densely packed structures.

Early attempts to automate this process within our research group led to the development of an initial image processing pipeline that combined classical computer vision techniques, such as the Circular Hough Transform (CHT), with foundational AI models like MicroSAM [23]. While this semi-automated approach demonstrated the feasibility of extracting

morphological metrics such as fibril diameter, density, and Nearest-Neighbor Distance (NND), and applying unsupervised clustering algorithms (k-means) to differentiate tendon types, the software architecture exhibited limitations. Its reliance on heuristic post-processing and parallel execution paths for different fibril sizes limited its full automation, scalability, metrics calculation, and overall segmentation accuracy, reaching a baseline clustering accuracy of 57%.

Recognizing the limitations of the baseline hybrid and semi-automated approach, and encouraged by recent trends in transmission electron microscopy that highlight the growing shift toward data-driven solutions [1], it was decided to switch to a leaner, AI-based end-to-end approach. This paper presents the design, implementation, and validation of an advanced, fully automated software pipeline for TEM image analysis to address these software engineering and analytical challenges.

The proposed modernized architecture replaces the previous hybrid segmentation approach with a robust deep learning model based on the ViT-UNet (Vision Transformer combined with U-Net) architecture. This AI-powered pipeline integrates targeted preprocessing modules (e.g., CLAHE and Anscombe transform for denoising), advanced segmentation capable of capturing both local details and global spatial relationships, and an automated feature extraction and clustering engine. By comparing this novel architecture with the baseline system, we evaluate how traditional software engineering methodologies for pipeline design can adapt to Transformer-based AI models to improve system reliability, precision, and the interpretability of morphological metrics.

The remainder of this paper is organized as follows. Section II discusses related work in AI-augmented medical imaging and software pipeline design. Section III details the software architecture and methodology of the proposed ViT-UNet-based pipeline. Section IV presents the experimental results, including a comparative validation of the extracted metrics and clustering performance against the baseline system. Finally, Section V concludes the paper and outlines future directions for the software lifecycle management of this analytical tool.

II. Related Work

The automated analysis of medical and biological images, particularly the segmentation of dense cellular or fibrillar structures, has been approached through various methodologies ranging from classical computer vision to advanced deep learning paradigms.

A. Traditional Image Processing and Unsupervised Techniques

Historically, the segmentation of filament-like structures in biomedical imaging relied heavily on multiscale filters and morphological analysis. Techniques such as the Frangi filter leverage Hessian matrix information to identify regions with geometries compatible with tubular or elongated structures [2]. For collagen fibril segmentation in Transmission Electron Microscopy (TEM), semi-automatic processes using thresholding algorithms followed by watershed segmentation techniques have been proposed [3], [4]. Similarly, automated image analysis of muscle fibers has utilized multiscale ridge detection combined with Gradient Vector Flow (GVF) Snake algorithms based on deformable models [5]. While interpretable and computationally lightweight, these classical filter-based approaches show significant limitations when dealing with high structural noise, overlapping fibrils, and spatially varying contrast—conditions that are typical in TEM imaging.

Furthermore, because supervised learning is not always applicable due to the scarcity of required labeled datasets in the medical domain [6], unsupervised and semi-supervised techniques remain a popular choice. For partition clustering tasks, algorithms like k-means are widely used [7], while Unsupervised Domain Adaptation (UDA) and transfer learning have proven highly beneficial in contexts where labeled data is minimal [8], [9]. Notably, research indicates that increasing data quantity or model complexity does not always lead to improved diagnostic accuracy; in some cases, simpler models maintain equal or superior performance [10].

B. Deep Learning and Convolutional Architectures

The introduction of Convolutional Neural Networks (CNNs) represented a substantial advancement in biomedical segmentation. Stacked Autoencoders and Convolutional Autoencoders have been utilized for cell segmentation and nuclei detection [7]. Specifically, the U-Net architecture has become a reference standard due to its encoder-decoder structure and skip connections, which effectively combine high-level semantic information with high-resolution spatial details [11]. Variations such as Attention U-Net introduce spatial attention mechanisms that allow the model to focus on the most relevant regions, improving segmentation in areas with dense or complex structures [12]. Other specialized architectures, such as CircleSnake, leverage Graph Convolutional Networks (GCNs) with circle representation to segment circular medical objects like renal glomeruli [13]. However, traditional CNNs tend to rely predominantly on local information, struggling to capture global spatial relationships in high-density fibrillar images.

C. Vision Transformers and Hybrid Models

To overcome the limitations of strictly convolutional approaches, models based on Transformers have been adapted for computer vision. The Vision Transformer (ViT) divides images into patches and uses self-attention mechanisms to learn global relationships across different regions [14]. This capability is highly advantageous for TEM images, where the global organization of fibrils provides critical context. Advanced iterations like the Swin Transformer introduce a hierarchical structure with shifted windows, improving scalability for high-resolution images [15].

Currently, the state-of-the-art is shifting toward hybrid CNN-Transformer architectures, such as ViT-UNet, which integrate a Transformer-based encoder with a convolutional U-Net decoder. This hybrid schema leverages the Transformer's

ability to extract global contextual information and long-distance relationships, while the convolutional decoder ensures accurate reconstruction of fine, local structures. This compromise makes hybrid models exceptionally suited for resolving overlapping structures and analyzing the complex structural organization of tendon fibrils in TEM imagery.

III. Materials and Methods

The development of the automated analytical pipeline was structured into sequential modular phases: data preprocessing, AI-driven instance segmentation, post-processing for feature extraction, and unsupervised clustering. This architecture allows for the systematic comparison between a baseline hybrid segmentation model and an advanced Transformer-based approach.

A. Dataset and Image Characteristics

The dataset comprises 35 high-resolution Transmission Electron Microscopy (TEM) images of tendon cross-sections. The samples were extracted from 7 donors and belong to five distinct anatomical classes: Anterior Cruciate Ligament (ACL), Patellar Ligament, Quadriceps, Semitendinosus, and Supraspinatus. The analyzed Region of Interest (ROI) for each image is 1696×1696 pixels, with a spatial resolution of 1.78 nm per pixel.

From a software engineering perspective, processing these raw TEM images presents significant challenges due to acquisition artifacts. The images are highly affected by mixed Poisson-Gaussian noise, which depends on the exposure time during acquisition [16]. Furthermore, non-uniform illumination and motion artifacts (motion blur) obscure fibril boundaries, making automated detection highly complex.

B. Image Preprocessing Pipeline

To mitigate non-biological variability and enhance the signal-to-noise ratio before feeding the data into the neural networks, a robust preprocessing sequence was implemented. First, to address non-uniform illumination, Contrast-Limited Adaptive Histogram Equalization (CLAHE) was applied, adjusting local intensities to simulate a uniform distribution without amplifying background noise. To handle the mixed Poisson-Gaussian noise, the Generalized Anscombe Transform was utilized to stabilize the variance, effectively transforming the Poisson distributed noise into an approximate Gaussian distribution [17]. A Wiener filter was then applied in the spatial domain to remove the Gaussian noise, followed by the inverse Anscombe transform to restore the original dynamic range. Finally, high-pass filtering and unsharp masking techniques were executed to enhance local contrast and sharpen the edges of the densely packed fibrils.

C. AI-Augmented Segmentation Architectures

Two distinct segmentation strategies were developed and evaluated to isolate the collagen fibrils.

1) Baseline Hybrid Approach (MicroSAM + CHT): The initial software architecture utilized a hybrid approach combining classical computer vision and foundational deep learning models. The pipeline operated in two parallel branches:

- **Circular Hough Transform (CHT)** was employed to detect small, relatively circular fibrils. The detection parameters were tuned specifically for smaller objects, and overlapping circles were filtered using a distance-based heuristic.
- **MicroSAM**, a fine-tuned version of Meta's Segment Anything Model (SAM) optimized for microscopy (using a ViT-L encoder) [22], [23], was deployed to detect larger and more complex fibrillar structures. To accurately identify the high density of objects within the images, the tiled version of the automatic mask generation function was used. For the majority of the dataset, a tile shape of (212, 212) yielded the most optimal results. Contrary to the MicroSAM API documentation recommendations, a halo larger than half of the maximum radius of the objects was selected; this configuration improved computational efficiency, while the resulting false positives were straightforward to filter during subsequent post-processing. Furthermore, this approach enhanced the final detection accuracy, as the default parameter values failed to identify all objects due to the highly packed arrangement of the tendons. Finally, due to hardware constraints, the number of points per batch during inference was capped at 256 to prevent memory overflow.

Because overlapping masks frequently occurred between the two parallel processes, two heuristic overlap-resolution algorithms were implemented. The first, used for the AI model, filters object masks using a metric of Roundness as a threshold. The latter, related to the CHT, is based on the distance between centroids and the intersection area [18]. Once the masks are fused to create an initial segmentation, the last step consists in applying a watershed transformation to separate falsely merged adjacent fibrils.

2) Proposed Architecture (ViT-UNet): To overcome the limitations of the hybrid baseline, a modern deep learning architecture, the ViT-UNet, was implemented. This hybrid Convolutional-Transformer network leverages a Vision Transformer (ViT) encoder to capture global spatial relationships and long-distance dependencies across the tissue, while utilizing a convolutional U-Net decoder with skip connections to accurately reconstruct high-resolution local contours. The model outputs a continuous probability map that maintains uncertainty information in transitional regions, which is subsequently thresholded to generate a definitive binary mask.

D. Post-processing and Morphological Feature Extraction

Raw segmentation masks inherently contain spurious noise and artifacts that can skew quantitative analysis. A dedicated post-processing module was developed using morphological operations (erosion, dilation, and opening) to smooth fibril contours, followed by connected-component labeling (CCL)

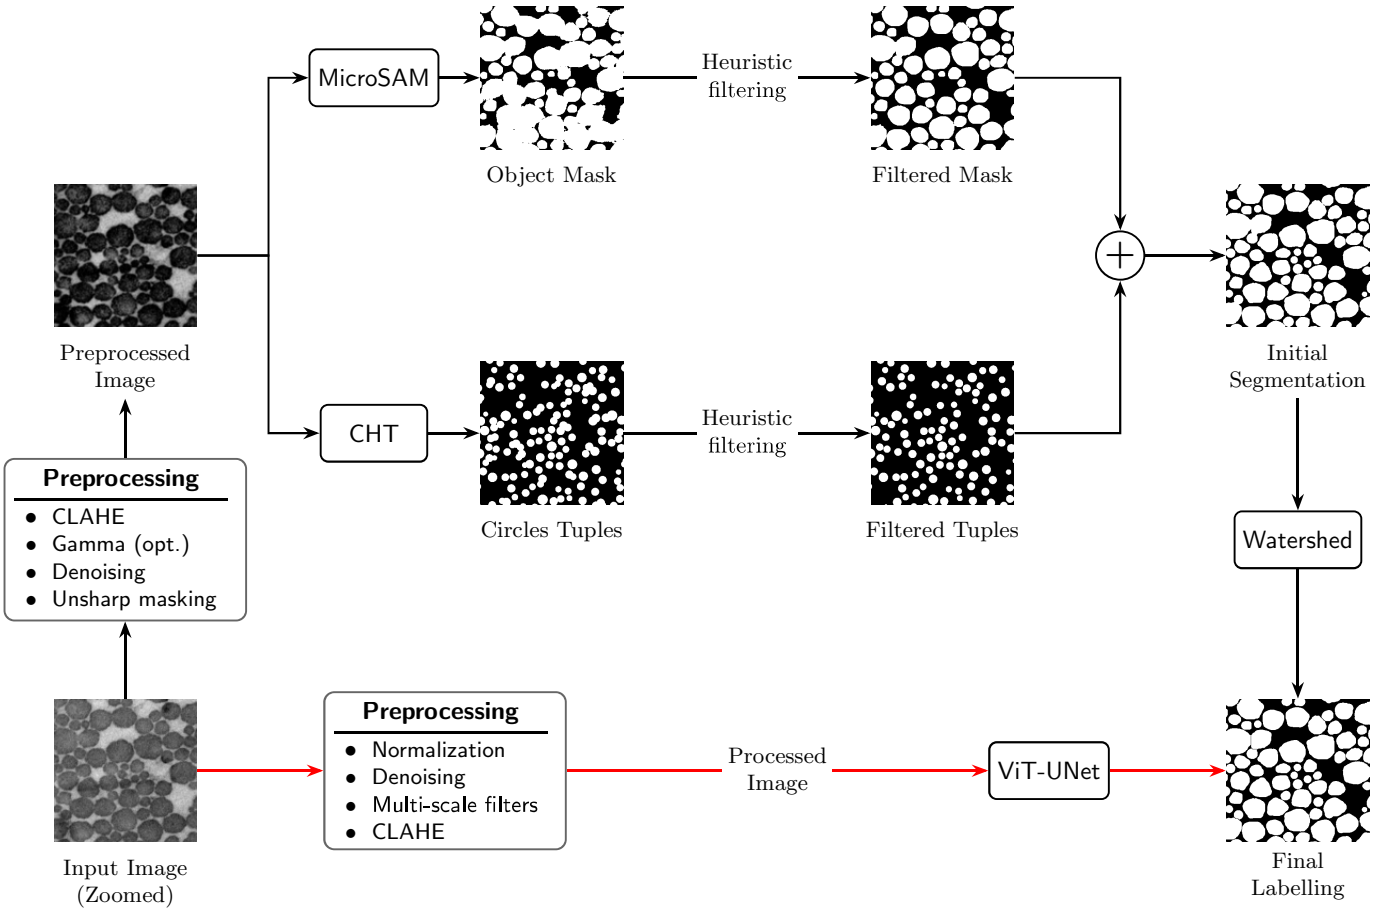

Fig. 1. Comparative architectural overview of the segmentation pipelines. The upper workflow represents the baseline hybrid approach: it integrates MicroSAM and Circular Hough Transform (CHT) branches which, after individual heuristic filtering, are combined through an **overlay** process to form the initial segmentation map. This combined output is then refined via watershed transformation for final labeling. The lower workflow (indicated by red arrows) illustrates the proposed architecture, which utilizes a unified ViT-UNet model to streamline the process by replacing the multi-branch overlay logic with a single, end-to-end deep learning stage.

to isolate individual structures. Skeletonization was also performed to extract the geometric center of each fibril, reducing the algorithm's dependency on local thickness variations.

From these refined masks, a comprehensive set of spatial and morphological metrics was extracted:

- **Quantity and Size**: Total fibril count, minimum, maximum, and mean cross-sectional area.
- **Spatial Occupancy**: Fibril area fraction (density), calculated as the ratio of the total fibril area to the image ROI area.
- **Statistical distribution of diameters**: The best-fit distribution of the fibril diameters was determined, as it can be directly related to the mechanical properties of the tissue [19] [20].
- **Spatial Pattern Analysis**: The organization of the tissue was quantified using the Nearest-Neighbor Distance (NND). To rigorously evaluate spatial randomness, the NND of the sample was compared against a theoretically generated Complete Spatial Randomness (CSR) Poisson distribution, yielding a standardized Z-score to determine if the pattern was clustered, random, or dispersed [21].

E. Unsupervised Clustering and Statistical Validation

To determine if the software pipeline could successfully differentiate the five anatomical tendon types purely based on the extracted morphology—without prior biological labeling—unsupervised machine learning was applied. All extracted features were standardized to ensure balanced weighting. Dimensionality reduction was performed using Principal Component Analysis (PCA) to project the high-dimensional feature space into a compact representation, minimizing redundant correlations.

A k-means clustering algorithm was configured with $k = 5$. Given that the dataset was inherently balanced (7 samples per anatomical class), a specialized balanced clustering approach was engineered. This was modeled as a combinatorial optimization task (the linear assignment problem) to strictly enforce equal cluster sizes. The Euclidean distances between samples and cluster centroids were computed to form a cost matrix, which was optimally solved using the linear sum assignment algorithm. The performance of the automated clus-

tering against the ground-truth anatomical labels was evaluated using accuracy metrics and the Adjusted Rand Index (ARI).

IV. Experimental Results

The experimental evaluation was designed to objectively compare the baseline hybrid pipeline against the proposed ViT-UNet architecture, assessing both segmentation quality and the discriminative power of the extracted morphological features through unsupervised clustering.

A. Segmentation Performance

Qualitative evaluation of the segmentation masks demonstrated that the ViT-UNet model significantly outperformed the baseline hybrid approach (MicroSAM combined with Circular Hough Transform). The baseline system struggled in regions with dense fibril packing and low local contrast, often requiring heuristic post-processing (such as watershed algorithms) that resulted in fragmented or merged masks. Conversely, the Transformer-based encoder of the ViT-UNet successfully leveraged global spatial context, allowing the model to accurately resolve individual fibril boundaries even in highly noisy and overlapping regions. This resulted in more consistent and stable masks, which directly translated to a reduction in non-biological variability during the feature extraction phase.

B. Morphological Feature Extraction

Following the segmentation and post-processing phases, a comprehensive set of morphological metrics was extracted for each TEM image (1696×1696 pixel ROI). The extraction process targeted both individual fibril geometries and their collective spatial organization. Key metrics included the total fibril count, minimum, maximum, and mean cross-sectional areas, and the fibrillar area fraction (density). Additionally, spatial pattern analysis was conducted using the Nearest-Neighbor Distance (NND) to evaluate tissue organization.

Notably, the **total fibril count**, the **maximum cross-sectional area**, and the **statistical distribution of diameters** (best-fit distributions) emerged as the most discriminative features for characterizing the structural differences between the distinct tendon types. These three parameters were selected to define the three-dimensional feature space utilized for the unsupervised clustering analysis and the subsequent performance evaluation, as visualized in Fig. 2.

C. Clustering and Classification Accuracy

To evaluate the descriptive capacity of the extracted metrics without introducing supervised bias, balanced k-means clustering ($k = 5$) was applied to the PCA-reduced feature space. The resulting clusters were evaluated against the ground-truth anatomical labels (ACL, Patellar Ligament, Quadriceps, Semitendinosus, and Supraspinatus) purely a posteriori.

The baseline hybrid pipeline achieved an overall clustering accuracy of 57.14% (correctly associating 20 out of 35 samples). In contrast, the proposed ViT-UNet pipeline demonstrated a quantifiable improvement, achieving an overall accuracy of 64.71% (22 out of 34 valid samples). The accuracy increase was not uniform across all classes. For instance, the proposed pipeline achieved high classification accuracies for specific clusters, correctly associating 85.71% of the Supraspinatus samples and 71.43% of the Semitendinosus samples. However, classes such as the Patellar Ligament (50.0%) and Quadriceps (57.14%) remained challenging to isolate.

TABLE I
Baseline Approach Accuracy

Cluster	Correct	Incorrect	Acc. (%)	Err. (%)
Supraspinatus	4	3	57.14	42.86
Quadriceps	2	5	28.57	71.43
Patellar Lig	5	2	71.43	28.57
ACL	4	3	57.14	42.86
Semitendinosus	5	2	71.43	28.57
Total	**20**	**15**	**57.14**	**42.86**

TABLE II
Proposed Approach Accuracy

Cluster	Correct	Incorrect	Acc. (%)	Err. (%)
Supraspinatus	6	1	85.71	14.29
Quadriceps	4	3	57.14	42.86
Patellar Lig	3	3	50.00	50.00
ACL	4	3	57.14	42.86
Semitendinosus	5	2	71.43	28.57
Total	**22**	**12**	**64.71**	**35.29**

V. Discussion and Conclusion

This paper presented the evolution of a software engineering pipeline for the automated analysis of biological TEM images, transitioning from a classical hybrid architecture to an advanced, AI-augmented framework utilizing a ViT-UNet model. The primary objective was to determine whether modernizing the deep learning architecture could enhance segmentation robustness and the stability of downstream morphological metrics.

The experimental results confirmed that integrating Transformer-based architectures significantly improves the spatial coherence of fibril segmentation in dense, noisy images. Consequently, the morphological metrics derived from these refined masks proved more reliable, leading to an absolute improvement in unsupervised clustering accuracy from 57.14% to 64.71%. However, this result also highlights a fundamental limitation: while improved software pipelines and segmentation accuracy lead to better feature extraction, morphological and spatial features alone possess an intrinsic limit in their ability to perfectly discriminate between complex biological tissue types.

A. Limitations

Despite the improvements, the proposed system faces several limitations. First, the size of the manually annotated dataset is relatively small, which constrains the ability to train deeper network architectures without risking overfitting. Second, the inherent heterogeneity of TEM images—characterized

by variable contrast, motion artifacts, and noise—continues to pose challenges for absolute metric standardization. Finally, the inherent uncertainty and human bias in the ground-truth manual annotations introduce a degree of error in the model evaluation process.

B. Future Work

Future iterations of this software pipeline will focus on overcoming the current discriminatory limits. Planned developments include the integration of advanced texture descriptors and multi-scale feature analysis to supplement purely morphological metrics. Additionally, we aim to explore self-supervised learning techniques to leverage larger volumes of unannotated TEM images, thereby improving model generalization. Ultimately, the goal is to formalize an end-to-end, fully automated software lifecycle that seamlessly integrates image preprocessing, AI-driven analysis, and automated diagnostic reporting for digital healthcare environments.

REFERENCES

[1] S. R. Spurgeon et al., "Towards data-driven next-generation transmission electron microscopy," *Nature Materials*, vol. 20, no. 3, pp. 274–279, 2021.

[2] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever, "Multiscale vessel enhancement filtering," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 1998.

[3] B. V. Rego, D. Weiss, and J. D. Humphrey, "A fast, robust method for quantitative assessment of collagen fibril architecture from transmission electron micrographs," *Microscopy and Microanalysis*, vol. 29, no. 6, pp. 2099–2107, Dec. 2023.

[4] T. Starborg et al., "Using transmission electron microscopy and 3View to determine collagen fibril size and three-dimensional organization," *Nature Protocols*, vol. 8, pp. 1433–1448, July 2013.

[5] J. Mula et al., "Automated image analysis of skeletal muscle fiber cross-sectional area," *Journal of Applied Physiology*, vol. 114, no. 1, pp. 148–155, Jan. 2013.

[6] A. Barragan-Montero et al., "Artificial intelligence and machine learning for medical imaging: A technology review," *Physica Medica: European Journal of Medical Physics*, vol. 83, pp. 242–256, Mar. 2021.

[7] K. Raza and N. K. Singh, "A tour of unsupervised deep learning for medical image analysis," *Current Medical Imaging*, vol. 17, no. 9, pp. 1059–1077, 2021.

[8] P. Kora et al., "Transfer learning techniques for medical image analysis: A review," *Biocybernetics and Biomedical Engineering*, vol. 42, no. 1, pp. 79–107, Mar. 2022.

[9] S. Kumari and P. Singh, "Deep learning for unsupervised domain adaptation in medical imaging: Recent advancements and future perspectives," *Computers in Biology and Medicine*, vol. 170, Mar. 2024.

[10] G. Varoquaux and V. Cheplygina, "Machine learning for medical imaging: Methodological failures and recommendations for the future," *npj Digital Medicine*, vol. 5, no. 48, Apr. 2022.

[11] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015.

[12] O. Oktay et al., "Attention U-Net: Learning where to look for the pancreas," *arXiv preprint arXiv:1804.03999*, 2018.

[13] E. H. Nguyen et al., "CircleSnake: Instance segmentation with circle representation," in *Proc. 13th International Workshop MLMI 2022: Machine Learning in Medical Imaging*, pp. 298–306, Sep. 2022.

[14] A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proc. International Conference on Learning Representations (ICLR)*, 2021.

[15] Z. Liu et al., "Swin Transformer: Hierarchical vision transformer using shifted windows," in *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.

[16] A. Suveer, "Methods for Processing and Analysis of Biomedical TEM Images," *Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology*, p.49, 2019.

[17] M. Mäkitalo, A. Foi, "Optimal Inversion of the Generalized Anscombe Transformation for Poisson-Gaussian Noise," *IEEE transactions on image processing*, vol. 22, pp.91–103, June 2012.

[18] J. Jiménez et al., "The Dynamic Identity Fusion Index," *Social Science Computer Review*, vol.34, Jan. 2015.

[19] K. L. Goh et al., "Age-related dataset on the mechanical properties and collagen fibril structure of tendons from a murine model," *Nature Scientific Data*, vol.5, pp.1–13, July 2018.

[20] D.A. Parry , G.R. Barnes , A.S. Craig, "A comparison of the size distribution of collagen fibrils in connective tissues as a function of age and a possible relation between fibril size distribution and mechanical properties," *Proceedings of the Royal Society of London. Series B, Biological Sciences*, vol. 203, Issue 1152 , pp. 305–321, Dec. 1978.

[21] Y. Yuan et al., "Point Pattern Analysis," *The Geographic Information Science & Technology Body of Knowledge*, 2020.

[22] A. Kirillov et al., "Segment anything," in *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4015–4026, 2023.

[23] A. Archit et al., "Segment anything for microscopy," *bioRxiv*, Aug. 2023.

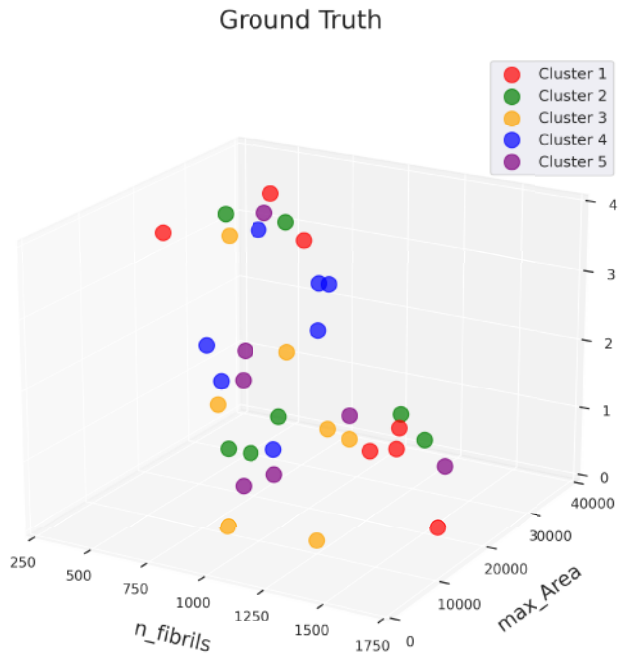

(a) Ground Truth (Dataset 1)

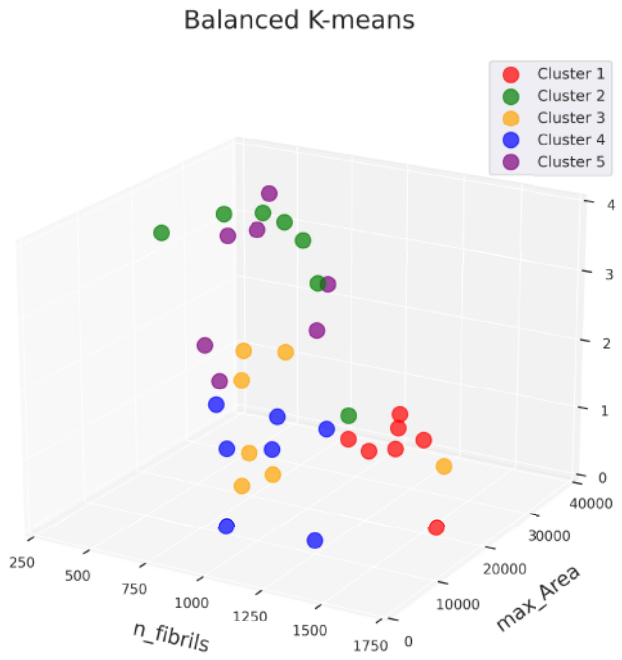

(b) Balanced K-means (Dataset 1)

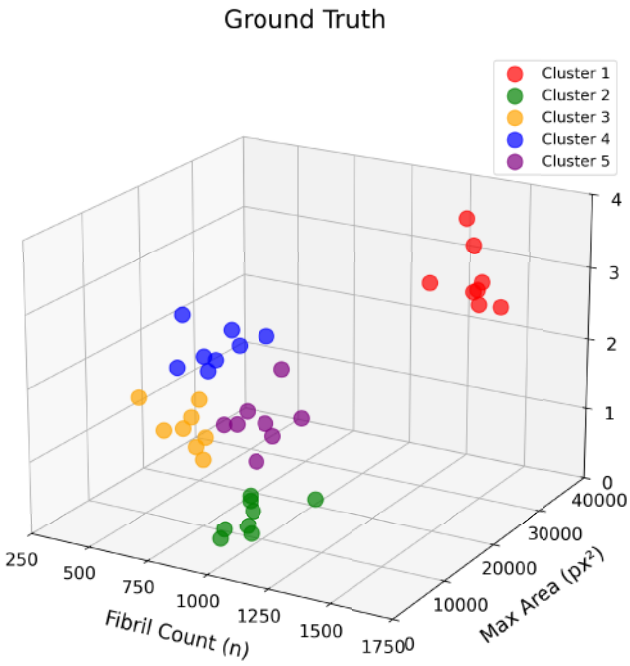

(c) Ground Truth (Dataset 2)

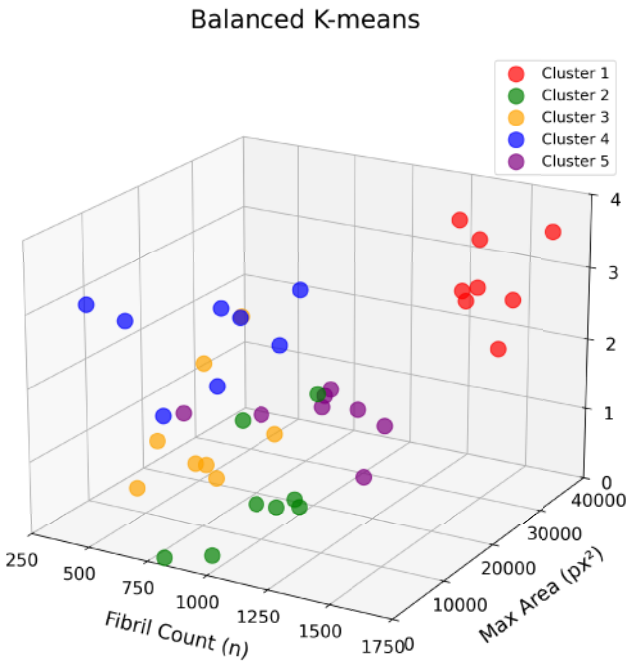

(d) Balanced K-means (Dataset 2)

Fig. 2. 3D feature space visualization and clustering performance. The plots illustrate the distribution of samples based on three key morphological descriptors: the total number of fibrils (X-axis), the maximum cross-sectional area (Y-axis), and the best-fit diameter distribution (Z-axis). The top row (a, b) displays the results for the first dataset, while the bottom row (c, d) refers to the second dataset. Comparing the ground truth labels with the Balanced K-means output highlights the algorithm's ability to accurately partition the descriptors into biologically consistent clusters.