

Data Descriptor

Not peer-reviewed version

A Dataset of Raw Fabric Grayscale Images for Defect Detection

[Rubén Pérez-Llorens](#), [Teresa Albero-Albero](#)^{*}, [Javier Silvestre-Blanes](#)

Posted Date: 20 April 2026

doi: 10.20944/preprints202604.1335.v1

Keywords: fabric defect detection; textile manufacturing; textile industry; image analysis; artificial vision; quality inspection; computer vision



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

A Dataset of Raw Fabric Grayscale Images for Defect Detection

Rubén Pérez-Llorens ^{1,2} , Teresa Albero-Albero ^{1,2,*}  and Javier Silvestre-Blanes ^{1,2} 

¹ Departamento de Informática de Sistemas y Computadores (DISCA), Universitat Politècnica de València (UPV). Plaza Ferrándiz y Carbonell s/n, Alcoy, 03801, Valencian Community, Spain

² Instituto Tecnológico de Informática (ITI). Camino de Vera, s/n, Valencia, 46022, Valencian Community, Spain

* Correspondence: talbero@disca.upv.es

Abstract

This article presents RAW-FABRID (RAW FABric Image Dataset), a publicly available annotated dataset for raw fabric defect detection using computer vision techniques. It addresses a major limitation in textile inspection, where reliance on private datasets hinders objective methodological comparisons. RAW-FABRID was acquired using a custom-built inspection machine equipped with controlled LED illumination and a line-scan camera. The dataset includes grayscale fabric images collected from several manufacturers to ensure variability in textures and patterns. It comprises 709 high-resolution images (1792×1024 pixels), including both defect-free and defective samples. To maximize reusability, data are provided in two complementary formats: high-resolution images (cropped to remove peripheral acquisition artifacts) for global analysis, and a patch-based organization following the widely adopted MVTec Anomaly Detection benchmark structure. The latter divides images into 256×256 pixel patches for direct machine learning integration. Crucially, the dataset is accompanied by comprehensive metadata (CSV) and precise COCO-formatted annotations (JSON) for both subsets, ensuring full traceability and supporting object detection and semantic segmentation. The dataset is publicly available through Mendeley Data, enabling reproducible research and objective benchmarking of defect detection algorithms.

Dataset: 10.17632/db6g85xsg.1

Dataset License: License under which the data set is made available (CC BY 4.0)

Keywords: fabric defect detection; textile manufacturing; textile industry; image analysis; artificial vision; quality inspection; computer vision

1. Summary

The primary motivation for compiling the RAW-FABRID dataset [1] was the need for a suitable public, comprehensively annotated benchmark in the field of fabric inspection and defect detection using computer vision and machine learning techniques. Automation in textile inspection has been recognized as a critical factor since the mid-1990s for reducing operational costs and improving product quality and customer satisfaction [2]. However, effective comparison of results across different studies has remained difficult, largely due to the widespread use of private datasets [3], heterogeneous imaging systems, and inconsistent acquisition parameters, which significantly limits the objectivity and reproducibility of reported methods.

The methodological context of the data generation involved image acquisition using a Basler raL2048-48gm line-scan camera, which features a 2048-pixel wide high-speed sensor. While this wide sensor allowed for comprehensive scanning, the final dataset images were systematically cropped to a width of 1792 pixels to eliminate non-informative dark borders and ensure uniform illumination across the samples. The system was designed with adjustable mechanical supports for both the camera

and LED illumination units, allowing control of the angle of incidence and working distance in order to obtain optimal spatial resolutions. This research was funded through the project “System for Defect Detection and Classification Using Artificial Vision Based on Deep Learning” [4]. This data article is not associated with a previously published research article.

The main contributions and key features of the RAW-FABRID dataset are summarized as follows:

- The RAW-FABRID dataset provides a publicly available, comprehensively annotated benchmark for raw fabric defect detection. This addresses a major limitation in textile inspection research, where most existing studies rely on private or poorly documented datasets, which restrict objective comparison and reproducibility of results [5].
- The dataset enables researchers to evaluate and compare computer vision and machine learning algorithms under consistent and well-documented acquisition conditions, facilitating fair benchmarking of defect detection and anomaly detection methods.
- The inclusion of high-resolution images with precise binary ground truth masks, further enriched by COCO-formatted bounding boxes and polygon annotations, allows for detailed spatial analysis of fabric defects. This supports the development of both traditional segmentation-based methods and advanced high-precision object detection algorithms.
- Reusability is drastically enhanced through a dual data organization strategy and full data traceability. The dataset is provided both as high-resolution 2D grayscale images (cropped to 1792×1024 pixels to remove peripheral acquisition artifacts) with corresponding masks, and as 256×256 pixel patches organized according to the widely adopted MVTec Anomaly Detection benchmark structure [6]. This dual format, accompanied by detailed CSV metadata, seamlessly supports both custom high-resolution processing workflows and standardized deep learning pipelines.

2. Data Description

This dataset contains a collection of high-resolution images designed to develop, test, and benchmark raw fabric defect detection algorithms. To ensure the robustness and generalizability of such methods, the fabrics used for image acquisition were collected from several different manufacturers, thereby introducing necessary variability in key characteristics such as texture, yarn thickness, thread density, and weave pattern.


To ensure a highly controlled experimental baseline, all acquisitions were performed on a standardized commodity fabric sourced from diverse manufacturers. Specifically, the base material consists of 100% cotton with a plain weave structure, featuring typical baseline parameters of approximately 9 threads/cm (warp) and 20 picks/cm (weft) using Ne 24/2 yarn. Sourcing this baseline specification from multiple suppliers introduces critical real-world intra-class variability, such as subtle loom tension differences and yarn batch variations, without confounding the anomaly detection algorithms with entirely different material architectures.

Furthermore, in the textile manufacturing pipeline, raw fabrics are systematically inspected prior to dyeing, printing, or finishing processes. Identifying structural anomalies at this early stage is critical to prevent the propagation of defects and avoid unnecessary chemical, environmental, and economic costs. Consequently, this dataset intentionally focuses on single-channel grayscale images to capture uncolored, structural defects, strictly reflecting this primary stage of industrial quality control.


All data are provided as 8-bit grayscale PNG images. The high-resolution images are stored in the *images* directory, having been cropped exclusively to remove non-fabric dark borders. This directory contains a total of 709 images, encompassing both defective and defect-free samples. The *masks* folder contains the 204 corresponding pixel-level ground truth binary masks for the defective images, sharing the exact same filenames as their associated samples. All high-resolution images have a final cropped resolution of 1792×1024 pixels, originally captured with a spatial resolution of 4 pixels per millimeter. The binary masks use two intensity values: 0 for non-defective background areas and 255 for defective

regions. These masks focus on binary defect segmentation (defect vs. defect-free) to support general anomaly detection tasks, without distinguishing between specific defect subclasses.


Representative examples of defect-free fabrics, defective samples, and their corresponding pixel-level ground truth masks are presented in Figure 1. These examples illustrate the variability in fabric textures as well as the pixel-level annotation quality provided in the dataset.



(a) Defect-free fabric image



(b) Fabric image with defect



(c) Corresponding binary ground truth mask

Figure 1. Representative high-resolution images from the RAW-FABRID dataset: defect-free fabric, defective fabric sample, and corresponding pixel-level ground truth mask.

In addition to the high-resolution images, the dataset is also provided in a patch-based organization within the *MVTec* directory, following the standard folder structure adopted in the MVTEC Anomaly Detection benchmark. The images are cropped into 256×256 pixel patches, enabling straightforward integration into existing anomaly detection frameworks and facilitating objective benchmarking. The *train* folder includes a *good* subfolder (14,196) containing only non-defective patches. The *test* folder contains two subfolders: *good* (4,969) with non-defective patches, and *defect* (687) with defective patches. Additionally, the *ground_truth* folder includes a *defect* subfolder (687) storing the corresponding binary masks for the defective test images, maintaining identical filenames with the suffix *_mask* added before the *.png* extension.

To ensure structured data handling and facilitate reproducibility, four auxiliary files are included in the dataset root directory. Comprehensive metadata is provided in two separate comma-separated values (CSV) files: *RAW_FABRID_HighRes_Metadata.csv* for the original high-resolution images, and *RAW_FABRID_Patches_Metadata.csv* for the cropped MVTEC dataset. Both files contain detailed information for each image or patch, including the filename, anonymized fabric origin ID, binary class label (good/defect), defect area, and number of defects per image. Furthermore, object detection and semantic segmentation annotations are supplied in two JSON files: *RAW_FABRID_HighRes_COCO.json* and *RAW_FABRID_Patches_COCO.json*. These files correspond to the high-resolution and patch-based datasets, respectively, and are formatted strictly as COCO-compatible JSON, defining precise bounding boxes and polygons [7,8]. Crucially, both the CSV and JSON files dedicated to the patch-based dataset explicitly map each 256×256 patch back to its original high-resolution source image, ensuring full data traceability.

In Figure 2, the directory structure of the dataset is illustrated. Additionally, Table 1 provides a statistical summary of the image distribution across the high-resolution and patch-based subsets. Examples of cropped patches organized following the MVTEC structure are shown in Figure 3.

Table 1. Quantitative distribution of images in the RAW-FABRID dataset.

Category	Original Images (1792 × 1024)	Patches (256 × 256)		
		Train	Test	Total Patches
Defect-free (Good)	505	14196	4969	19165
Defective	204	-	687	687
Total Images	709	14196	5656	19852

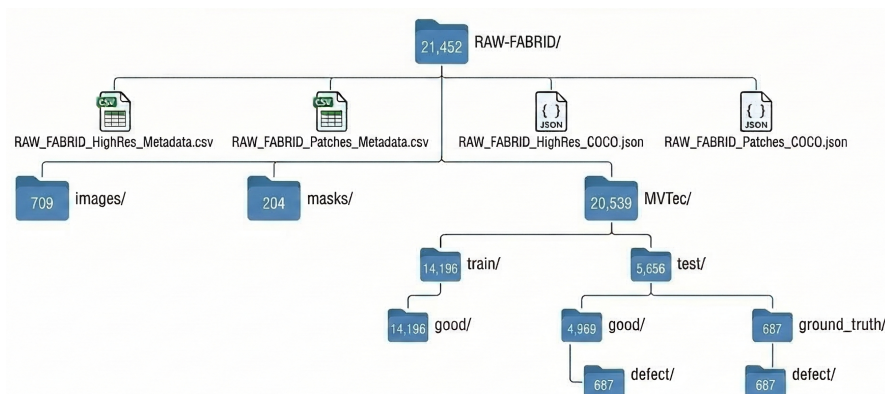


Figure 2. Overview of the file structure of the dataset.

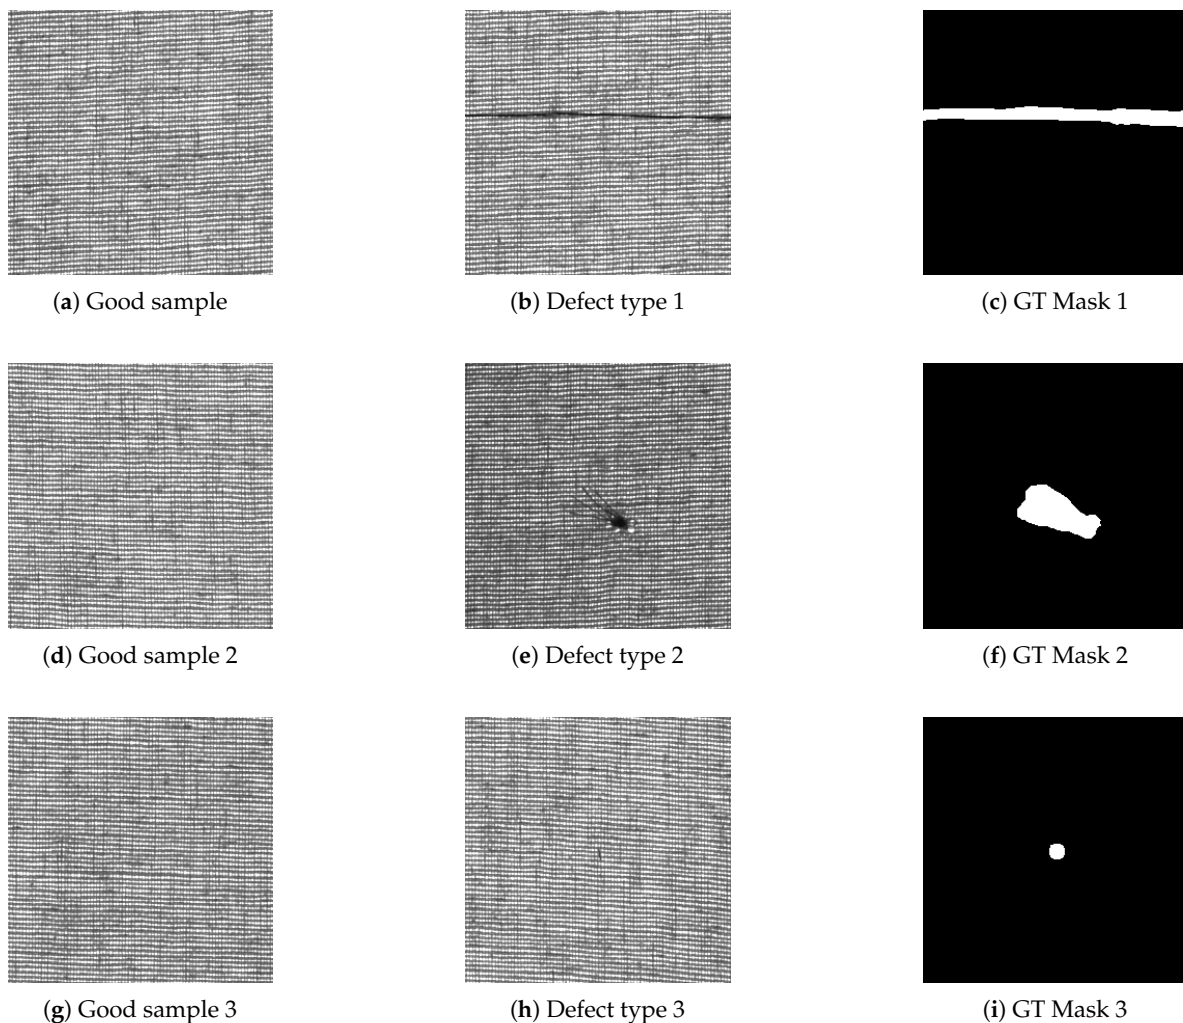


Figure 3. Representative examples of 256×256 cropped patches organized following the MVTec structure. The first column shows non-defective samples, the second column displays three different types of industrial defects, and the third column provides the corresponding ground truth masks.

3. Methods

The generation of the RAW-FABRID dataset followed a comprehensive experimental pipeline comprised of four main stages: data acquisition, manual expert annotation, data preprocessing and organization, and patch-based formatting for benchmarking. Initially, fabric rolls sourced from diverse manufacturers were digitized using the custom-built high-speed inspection machine described in the Background section. Following acquisition, the resulting high-resolution grayscale images underwent rigorous manual inspection to establish pixel-level ground truth masks, alongside precise bounding box and polygon annotations for all defective samples. Finally, to support the dual data organization strategy highlighted in Section 1, the annotated data was processed into two distinct formats: high-resolution 1792×1024 cropped images and standardized 256×256 patches organized according to the MVTec anomaly detection structure. The generation of comprehensive CSV metadata and COCO-compatible JSON files was integrated into this final stage to ensure full traceability. The specific hardware components, mechanical setup, and software procedures employed in each stage are detailed in the following subsections.

3.1. Data Acquisition System

The image acquisition system employed a Basler raL2048-48gm line-scan camera model with a spatial resolution of 2k. This resolution was selected to adequately cover the inspected textile area in the custom-built inspection machine while maintaining sufficient spatial detail for defect detection.

The sensor features a pixel size of $7\ \mu\text{m} \times 7\ \mu\text{m}$. Inspection of larger fabric widths can be achieved by employing higher-resolution sensors or multiple cameras. In the case of using a much larger sensor, for example, an 8k model, the pixel size is typically reduced to $3.5\ \mu\text{m} \times 3.5\ \mu\text{m}$, which implies lower light sensitivity, requiring significantly higher illumination intensity.

The camera operates at a maximum line rate of 51 kHz, which corresponds to fabric speeds exceeding 12 m/s at a spatial resolution of 4 pixels per millimeter, well above typical industrial textile inspection requirements. Image transmission is performed via a GigE interface, which provides sufficient bandwidth for real-time acquisition. The camera supports C-mount, F-mount, and M42 optics, allowing flexible lens selection depending on field of view and working distance requirements. Figure 4 shows the Basler line-scan camera used in the acquisition system.




Figure 4. Basler raL2048-48gm line-scan camera used for dataset acquisition.

To ensure stable operation and flexible configuration, both the camera and illumination units were mounted on adjustable mechanical supports, as detailed in Figure 5. Specifically, Figure 5(a) shows the pan and tilt mechanisms that allow precise positioning of the camera relative to the fabric surface. A similar adjustable mechanism, presented in Figure 5(b), was implemented for the LED illumination bars, enabling control over both height and angle of incidence, which is critical for highlighting different types of surface defects. For the acquisition of this specific dataset, a dual illumination strategy was employed, activating both back and front lights. The front illumination bars were configured with an incidence angle of 45° and positioned at a working distance of 5 cm from the camera's field of view.



(a) Mechanism for precise adjustment of camera pan and tilt



(b) Support system for LED illumination height and angle control

Figure 5. Adjustable mechanical integration of the vision system components. The figure details the specific mechanisms used for positioning relative to the fabric surface for both the camera (a) and the illumination bars (b).

The fabrics were transported through the inspection area in a controlled manner, while an encoder wheel synchronized the fabric motion with the line-scan camera acquisition to ensure a consistent spatial resolution across all captured images. This synchronization enabled precise line triggering and uniform sampling during image acquisition.

Following the acquisition phase, all captured high-resolution images underwent a rigorous manual quality control process supervised by domain experts. This step ensured that the dataset consists solely of samples with proper focus, adequate sharpness, and consistent illumination across the central region. Consequently, to standardize the data and systematically eliminate unavoidable peripheral shadowing or physical fraying at the roll edges, all approved images were systematically cropped to a final width of 1792 pixels. Any images exhibiting significant capture artifacts within this central region were entirely discarded prior to the annotation stage.

Figure 6 shows the complete custom-built inspection system used for generating the RAW-FABRID dataset.




Figure 6. Complete textile inspection machine used for dataset acquisition.

3.2. Defect Annotation Procedure

To establish reliable ground truth data for supervised learning stages, pixel-level semantic segmentation was performed on all images identified as containing defects. The annotation process was conducted manually by domain experts from the fabric supplier companies, ensuring highly accurate outlines of textile imperfections based on professional visual inspection.

In the defective regions, experts generated precise pixel-level segmentation masks. Consequently, only the defective samples in the *images* directory have corresponding files in the *masks* folder; defect-free images do not include mask files. Within these masks, pixels belonging to a defect region are strictly encoded with an intensity value of 255 (white), while non-defective background areas are encoded with a value of 0 (black). Furthermore, the generated annotation files (CSV and JSON) provide comprehensive information about the corresponding images, integrating precise bounding box and polygon coordinates for all defects.

3.3. Image Preprocessing and Formatting

Photometrically, the high-resolution grayscale images stored in the *images* directory remain in an unprocessed state, exactly as acquired by the camera sensor. No illumination correction, intensity adjustments, or histogram clipping were applied. Spatially, the original 2048-pixel wide captures were exclusively cropped to a width of 1792 pixels to remove uninformative dark borders. Preserving this photometrically raw format allows future researchers to develop and benchmark their own preprocessing or normalization algorithms.

To facilitate compatibility with standard defect detection frameworks, specifically the MVTEC Anomaly Detection benchmark structure, the high-resolution 1792×1024 images were divided into a grid of 256×256 pixel patches. This specific patch size (256×256) was selected to preserve the original spatial resolution (4 px/mm) while generating input dimensions compatible with the receptive fields and memory constraints of standard Convolutional Neural Networks (CNNs) and Vision Transformers

(ViTs). Resizing an entire high-resolution image to standard CNN input dimensions would cause severe interpolation artifacts and the potential loss of micro-defects. Therefore, this patch-based approach ensures seamless integration into modern anomaly detection architectures (e.g., PatchCore, PaDiM) without sacrificing critical pixel-level details.

The grid division yields a baseline of $7 \times 4 = 28$ patches per original image. To prevent defects from being split across patch boundaries, a dynamic cropping strategy was applied: patches containing defects near the grid edges were recentered, introducing deliberate spatial overlap to ensure the entire anomaly is captured within a single patch and to improve detection robustness at the borders.

The resulting patches were automatically sorted based on their spatial overlap with the ground truth masks. Patches containing zero defective pixels were allocated to either the training (train/good) or testing (test/good) sets. Patches containing at least one defective pixel were placed in the testing set (test/defect). Simultaneously, the corresponding high-resolution ground truth masks were cropped using identical grid coordinates to generate precise 256×256 binary masks for every defective patch.

Finally, to facilitate automated data handling and guarantee reproducibility, this formatting stage concluded with the generation of the four comprehensive metadata and annotation files (CSV and JSON) detailed in Section 2, successfully mapping every patch back to its original source image and recording its specific data split allocation.

4. User Notes

The dataset consists of single-channel grayscale images. Therefore, it is specifically designed for the detection of structural and textural defects (holes, broken threads, etc.) and is not intended for analyzing color-based anomalies.

Author Contributions: Conceptualization, J.S.-B.; methodology, J.S.-B.; software, R.P.-L.; validation, R.P.-L. and T.A.-A.; formal analysis, R.P.-L. and T.A.-A.; investigation, R.P.-L. and T.A.-A.; resources, J.S.-B.; data curation, R.P.-L. and J.S.-B.; writing—original draft preparation, R.P.-L.; writing—review and editing, T.A.-A. and J.S.-B.; visualization, T.A.-A. and J.S.-B.; supervision, J.S.-B.; project administration, J.S.-B.; funding acquisition, J.S.-B. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Agència Valenciana de la Innovació (AVI) through the project “System for Defect Detection and Classification Using Artificial Vision Based on Deep Learning”, under the program Valorization and Transfer of Research Results to Companies, call 2022–2024, grant number INNVA1/2022/20. The project was co-financed by the Instituto Valenciano de Competitividad e Innovación (IVACE) and the European Union. The APC was funded by the aforementioned project INNVA1/2022/20.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The RAW-FABRID dataset presented in this study is openly available in Mendeley Data at <https://doi.org/10.17632/db6g85xsg.1> [1].

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

RAW-FABRID	RAW FABric Image Dataset
LED	Light-Emitting Diode
MVTec	Machine Vision Technologies
CSV	Comma-Separated Values
COCO	Common Objects in Context
JSON	JavaScript Object Notation
CNN	Convolutional Neural Network
ViT	Vision Transformer
GT	Ground Truth

References

1. Pérez-Llorens, R.; Albero-Albero, T.; Silvestre-Blanes, J. RAW-FABRID: RAW FABric Image Dataset for defect detection. Mendeley Data, V1, 2026. <https://doi.org/10.17632/db6g85xsg.1>.
2. Kumar, A. Computer-vision-based fabric defect detection: A survey. *IEEE Transactions on Industrial Electronics* **2008**, *55*, 348–363. <https://doi.org/10.1109/TIE.2007.906125>.
3. Hanbay, K.; Talu, M.F.; Özgüven, Ö.F. Fabric defect detection systems and methods: A systematic literature review. *Optik* **2016**, *127*, 11960–11973. <https://doi.org/10.1016/j.ijleo.2016.09.110>.
4. Silvestre-Blanes, J.; Pérez-Lloréns, R. System for Defect Detection and Classification Using Artificial Vision Based on Deep Learning [Research Project]. Available at: <https://www.upv.es/entidades/epsa/proyectos-de-investigacion-en-curso/tejidos-deep-learning/>, 2022. (Accessed: 20 March 2026).
5. Ngan, H.Y.T.; Pang, G.K.H.; Yung, N.H.C. Automated fabric defect detection—A review. *Image and Vision Computing* **2011**, *29*, 442–458. <https://doi.org/10.1016/j.imavis.2011.02.002>.
6. Bergmann, P.; Fauser, M.; Sattlegger, D.; Steger, C. MVTEC AD—A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 9592–9600. <https://doi.org/10.1109/CVPR.2019.00982>.
7. Lin, T.Y.; Maire, M.; Belongie, S.; Bourdev, L.; Girshick, R.; Hays, J.; Perona, P.; Ramanan, D.; Zitnick, C.L.; Dollár, P. Microsoft COCO: Common Objects in Context. In Proceedings of the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2014, pp. 740–748.
8. COCO Consortium. COCO Dataset Data Format. <https://cocodataset.org/#format-data>, 2026. Accessed: 2026-03-17.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.