ICPA
17th INTERNATIONAL CONFERENCE ON PRECISION AGRICULTURE
ConBAP 2026
CONGRESSO BRASILEIRO DE AGRICULTURA DE PRECISÃO E DIGITAL
13-16 July 2026 | Porto Alegre, Brazil

SPARC-AI: Synthetic Procedural Agricultural Rendering and Annotation Framework for Crop Phenotyping and AI Applications

Rafael Volpe de Freitas[1], Vinicius Silveira Mello[1,2], Giusepe Disconzi Dallegrave[1,2], Eduardo Farinati Leite[1,2], João Francisco Valiati[1,2].

[1]CEDRA - Center for Embedded Devices and Research in Digital Agriculture, SENAI-RS;

[2]SENAI Innovation Institute for Sensing Systems (ISI-SIM), São Leopoldo, Brazil;

A paper from the Proceedings of the
17th International Conference on Precision Agriculture and 11th Brazilian Congress on Precision and Digital Agriculture
13-15 July 2026
Porto Alegre, Brazil

Abstract.

Between 20% and 40% of global agricultural production is lost annually to pests and diseases, generating economic damages estimated at over US$220 billion each year. This persistent challenge underscores the urgent need for scalable, precise, and cost-effective monitoring solutions. In this scenario, Artificial Intelligence (AI) based pathogen detection systems emerge as transformative tools, enabling high-resolution spatial and temporal monitoring of crop health. However, the performance and generalization capacity of AI models are fundamentally constrained by the availability of large, accurately annotated datasets, resources that are inherently difficult to obtain due to seasonality, phenological variability, and fluctuating disease severity in real-world conditions. To overcome this bottleneck, we introduce SPARC-AI, a fully synthetic, procedural data generation framework grounded in digitally reconstructed agricultural environments. Unlike conventional data augmentation strategies that still rely on real samples, our approach greatly reduces the dependency on real annotated datasets by generating entirely synthetic yet photorealistic training data. The system produces two levels of semantic annotation: infected leaf segmentation and segmentation of diseased regions within individual leaves, enabling multi-scale learning and precise pathology localization. The methodology begins with the laboratory acquisition of high-resolution photographic textures from real leaves, including translucency information extracted from the captured imagery. These textures are then mapped onto procedurally generated plant models, with the plant morphology and developmental progression being modeled entirely procedurally and allowing automated generation of distinct phenological stages. Developmental control is parameterized through stem length, which governs structural thickness, leaf density and scale, and the emergence of additional elements such as flowers, fruits, and grains. Disease patterns are likewise procedurally synthesized, and their spread is controlled through noise maps that, after thresholding, generate spatially coherent infected regions. Leaf translucency, a critical photorealistic attribute, is reproduced using a chromatic threshold in the white spectrum to *enhance light-tone saturation, combined with an alpha-channel mapping. This strategy enables realistic subsurface light diffusion using a single synthetic light source, modeled as an artificial sun, thereby optimizing rendering efficiency and enchasing visual fidelity. A validation rendering process conducted on a soybean crop model generated 100 unique scene compositions, each containing both photorealistic and semantically annotated outputs. The complete dataset was produced in 14 minutes and 11 seconds in an NVIDIA L40s GPU, averaging at approximately 2 seconds per image. Synthetic datasets were also successfully generated for grapevine, soybean, and corn crops, incorporating major diseases such as Esca (Phaeoacremonium aleophilum), Leaf Blight (Alternaria spp.), Black Rot (Guignardia bidwellii), Downy Mildew (Plasmopara viticola), and Asian Rust (Phakopsora pachyrhizi). The results directly address the core challenge of limited annotated data by demonstrating that our framework delivers a high-quality, efficient, and flexible fully synthetic pipeline. The methodology supports multiple crop types, with controllable infection stages, varying times of day, heterogeneous spatial disease distributions, and distinct phenological growth phases. By enabling scalable, precise, and adaptable dataset generation, SPARC-AI represents a powerful and practical advancement for AI-driven pest and disease detection in precision agriculture.*

Keywords.
synthetic data, plant disease recognition, procedural modeling, artificial intelligence, precision agriculture

Introduction

Agriculture plays a fundamental role in global food security and in the economic sustainability of numerous countries. However, a significant portion of global agricultural production is lost annually due to the incidence of pests and diseases affecting crops. According to estimates from a 2025 report from the Food and Agriculture Organization of the United Nations (FAO), approximately 40% of worldwide agricultural production is compromised by these factors, resulting in substantial economic impacts, reduced productivity, and risks to food supply stability.

To prevent these losses, the identification of diseases and pests in agricultural crops is traditionally carried out through human visual inspection. Artificial intelligence techniques, especially those based on computer vision and deep learning (Liu et al, 2020), have stood out as promising alternatives for automating plant recognition processes and agricultural disease detection. However, as stated by Alzubaidi (2023) and Najafabadi (2015), the effective training of these models depends directly on the availability of large volumes of annotated, diverse, and representative data from the different conditions found in the field. Obtaining these real datasets constitutes a significant challenge due to the high cost of collection, the need for specialists for manual annotation (Barbedo, 2016), and the dependence on external factors such as climate, seasonality, and the availability of crops throughout the year.

Despite its widespread adoption, this approach presents several limitations. Diagnostic accuracy often depends on the expertise of specialists and the subjective interpretation of symptoms, while different diseases may exhibit similar visual characteristics, increasing the likelihood of misclassification. In addition, complex backgrounds can hinder the accurate separation and annotation of target objects, making it susceptible to human errors and annotation inconsistencies (Beck, 2023). Furthermore, machine learning models developed using this approach require substantial amounts of accurately annotated data, the acquisition of which is costly and time-consuming (Schröder, 2020; Rechkemmer, 2023).

Faced with these challenges involving annotation and data scarcity, this work presents SPARC-AI, a pipeline aimed at the automated generation of synthetic datasets for artificial intelligence applications in the agricultural context. The proposal enables the simultaneous creation of realistic images and their respective annotations, significantly reducing the manual effort associated with data preparation. In addition, the pipeline mitigates potential human errors and inconsistencies, contributing to higher data quality and reliability. The pipeline was conceived as a generic and extensible framework for agricultural dataset generation, favoring the generation of diverse and representative scenarios. In this study, we investigate its application in selected crops for pest and disease detection, although its design readily supports adaptation to other crops, agricultural challenges, and future application domains.

Methodology

The developed pipeline combines traditional and procedural 3D techniques to generate realistic digital plant representations while maintaining flexibility and scalability. Initially, the subject is photographed under controlled conditions to acquire high-quality visual data. The resulting images are processed using image-editing software to extract and refine texture maps, capturing relevant surface characteristics such as color, patterns, and structural details. Subsequently, a 3D modeling software is employed to construct the geometric mesh, onto which the extracted texture maps are applied to produce visually accurate digital replicas. In addition to photorealistic rendering, the methodology was designed to generate a corresponding semantic representation for each synthesized image, enabling the simultaneous creation of visually realistic outputs and annotated datasets suitable for computational analysis, computer vision, and machine learning applications.

Image acquisition and texture conversion for leaves

The pipeline, as shown in Figure 1, begins with the acquisition of primary data intended for the generation of base textures. The material consists of photographic records of leaves, obtained under laboratory conditions. The images are captured against a neutral background, preferably white, under diffuse lighting conditions to minimize the formation of casted shadows, with an orthogonal framing (90° angle) relative to the leave plane, characterizing a top view (presented in Figure 2).

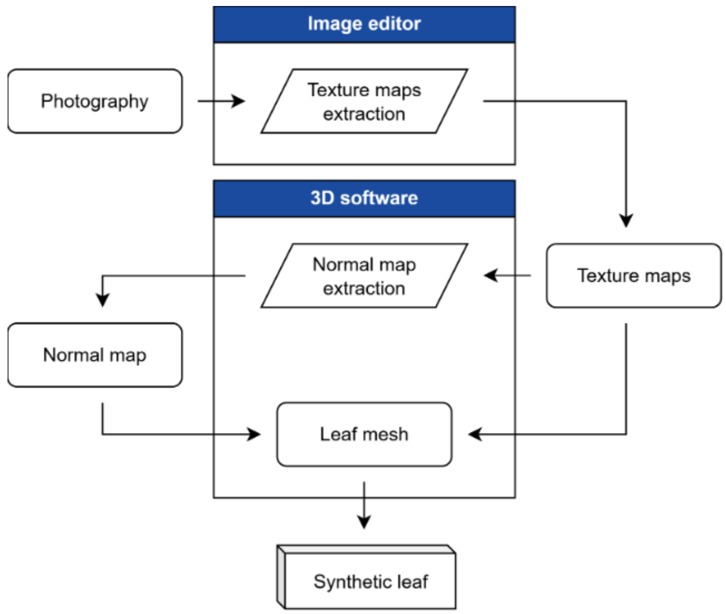

Fig 1. Diagram representing the texture work process for leaves.

Fig 2. Photograph of real leaf.

After image acquisition, the photographs are imported into image-editing software and subjected to a series of digital processing steps. Initially, the background and any residual shadows are removed, followed by repositioning the image within the frame to maximize the utilization of the available pixel area. The processed image is then used as the color map, which provides the texture's visual color information and surface appearance. Subsequently, the image is desaturated and its white and black levels are adjusted to generate roughness maps and height maps (Figure 3a and 3b). The roughness map defines the spatial variation in surface reflectance, controlling the degree of glossiness or matte appearance, whereas the height map encodes surface elevation information, enabling the simulation of fine-scale geometric relief and surface irregularities during rendering.

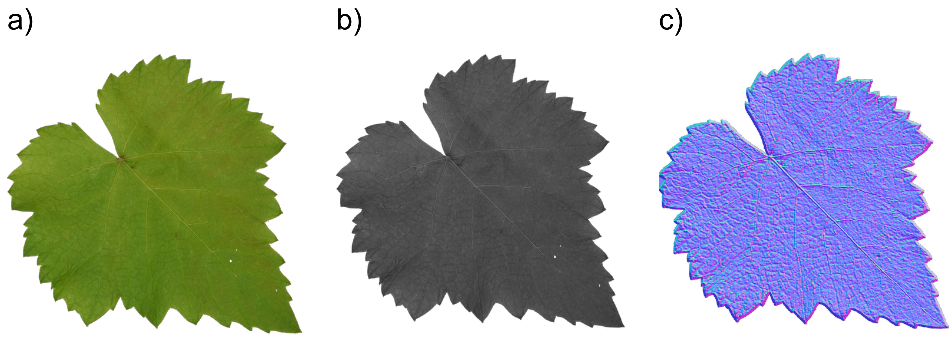

Fig 3. Texture maps extracted from the photograph.

The height map is then imported into the 3D modeling software, where it is used as a displacement map applied to a flat mesh with high polygon density. The resulting mesh acquires geometry corresponding to the original leaf. Next, the normals of this geometry are transferred to a simplified plane composed of a single polygon, allowing the generation of the normal map (Blinn, 1978) (as seen in Figure 3c).

Procedural texturing of synthetic diseases

Once the base leaf textures are prepared, the disease textures are generated and overlaid onto the primary leaf texture. These disease patterns are produced entirely through synthetic and procedural approaches using a node-based texture generation framework, enabling the creation of highly customizable materials and textures through mathematical parameterization, allowing substantial variability to be achieved by modifying only a limited set of input parameters.

The specifics of each synthetic disease vary from case to case, but the process consists of one or multiple noise maps, usually Perlin, as described by Taylor (2021), or Voronoi noise maps that are combined, have their values adjusted, and have levels of black and white filtered to create masks that resemble the patterns of the studied disease. After that, the grayscale values of the resulting pattern are converted into RGB colors to resemble the colors that the infected areas of the disease.

Plant morphology procedural modeling

The morphological structure of plants is modeled in a fully procedural manner, through a parametric node-based modeling system (Krecklau, 2012). The generation process starts from a single initial vertex, from which a hierarchical sequence of branches defined by Bézier curves (Khan, 2011) is developed. The specific branching sequence varies according to the type of plant being represented. After the complete curve structure is established, Perlin noise is applied to the control points in order to give greater organic quality to the generated forms. Next, a radial volume is distributed around the curves, which is then used as a parameter in a voxelization process responsible for generating the final mesh that composes the plant's branches.

The leaves are instantiated from flat objects to which specific texture maps are applied. These objects undergo manual deformations to reproduce morphological characteristics of real leaves and are subsequently procedurally instantiated at the ends of the stems of the synthetic plant (as shown in Figure 4). Additional elements, such as fruits and flowers, are also generated from manually modeled objects, whose distribution regions are adjusted according to the characteristics of the recreated plant species. A simplified diagram of the process can be seen in Figure 5.

Fig 4. Full synthetic grape tree.

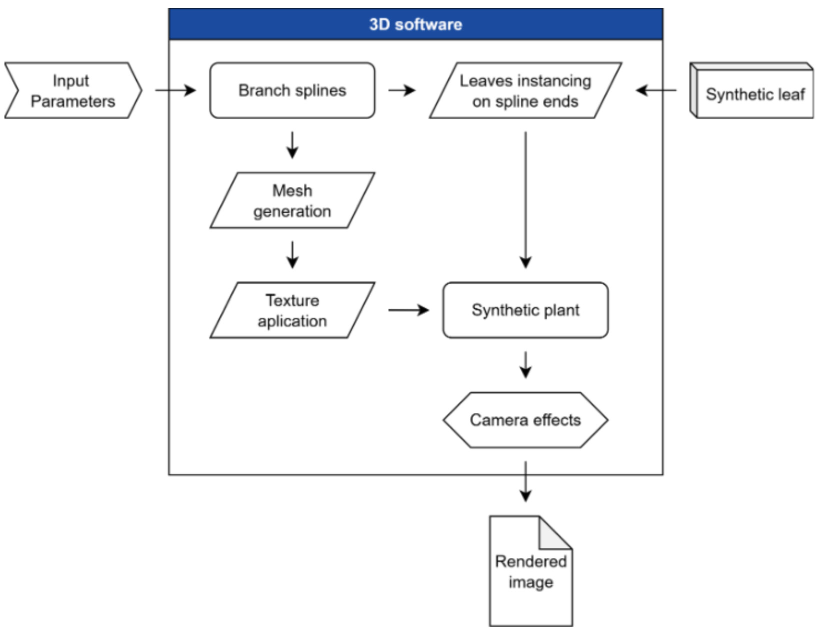

Fig 5. Diagram representing the synthetic plant pipeline.

A plant growth system was developed based on manipulating the extension of the vectors that compose the plant's structure. For this purpose, an input variable is used to define the plant's maturity stage, ranging from 0 (representing no growth) to 10 (representing the final developmental stage). This variable is used in the calculations that determine the degree of extension of the structural vectors, with the original previously defined size as the maximum limit. In addition, the system also controls the thickness of the branches, including both the overall thickness and the dimensional reduction gradient that occurs toward the plant's extremities.

Scene setup and lighting

The scene is generated using conventional 3D modeling and rendering techniques. The terrain is built using digital sculpting methods and is subsequently textured with traditional texture maps. Next, vegetation assets are procedurally distributed to compose the low-lying vegetation environment, using an input variable responsible for defining the distribution seed value. The result can be seen in Figure 6.

Fig 6. Grape tree scene with climate modified to be cloudy.

The lighting system acts on the turbidity variable, allowing the generation of atmospheres with higher particle density, resulting in visual effects associated with cloudy conditions. This system is based on the method proposed by García Liñán (2023), whose parameters are randomized in order to simulate different periods of day. Additionally, a climate control system based on the lighting model described by Preetham (1999) was implemented.

Annotated rendering

Regarding the annotated rendering, the texturing pipeline employs a distinct methodology. Rather than relying on the same texturing techniques seeing in the realistic rendering, the annotated visual aesthetic is achieved by assigning each object a monochromatic emissive material. This material type eliminates self-shadowing effects, thereby producing a uniformly unshaded appearance, as seen in Figure 7..

In contrast, the ground plane and surrounding scenery are assigned a solid black material with the same emissive properties, which visually isolates the target object in the final rendering. Furthermore, because all scene elements utilize emissive materials, the emitted light does not interact with or reflect upon other surfaces. Consequently, the resulting render exhibits a flat, solid-color appearance devoid of conventional lighting interactions or shading effects.

For the disease textures, the original grayscale noise map employed during the procedural texture generation process is reused. The same parameter values as those applied in the realistic counterpart are maintained to preserve pattern consistency between both representations. Subsequently, the grayscale values are converted into binary black-and-white regions. This resulting map is then utilized as a masking layer to determine the spatial distribution, specifying where the standard leaf emissive material is displayed and where the disease-specific variant is applied. Through this process, an annotated representation of the diseased leaf is generated.

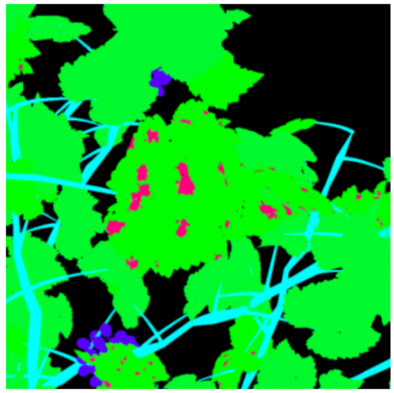

Fig 7. Piece of the grape tree rendered as annotated image.

The annotated image is rendered without the application of color conversion processes, thereby ensuring that the RGB values assigned to the annotated objects remain identical to the original material input values. Additionally, all anti-aliasing effects are disabled, resulting in serrated object boundaries. This approach prevents the introduction of transitional color tones that could negatively affect the consistency and reliability of subsequent AI training processes.

Tests and results

There were two separate preliminary tests. The first one generated synthetic datasets for grapevine, soybean, and corn crops, incorporating major plant diseases such as Esca (*Phaeoacremonium aleophilum*), Leaf Blight (*Alternaria spp.*), Black Rot (*Guignardia bidwellii*), Downy Mildew (*Plasmopara viticola*), and Asian Rust (*Phakopsora pachyrhizi*). The proposed approach could generate realistic renderings that convincingly reproduced both the morphology of the studied plants and the visual characteristics of disease infections. In addition, an accurately annotated version of each rendered image was simultaneously produced, enabling direct application in supervised machine learning tasks. A sample of the results showing soy and corn crops can be seen in Figure 8 rendered as realistic and annotated images.

a-1) a-2)

b-1) b-2)

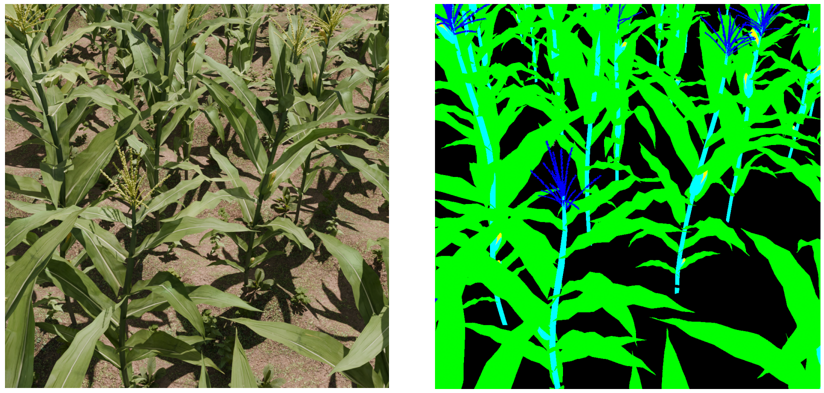

Fig 8. Soy (a) and corn (b) synthetic plantations rendered both realistic (1) and annotated (2).

The second preliminary test was conducted with the objective of generating a set of 200 images, consisting of 100 realistic images and their respective 100 annotated versions. Processing was carried out on an NVIDIA L40S GPU, with a total time of approximately 11 minutes. A synthetic vine was used as the test subject.

As a result, a large and highly diverse dataset was obtained, suitable for training artificial intelligence models aimed at detecting pests and diseases in the analyzed crops. Additionally, the annotated versions of the images were generated simultaneously with the realistic images, significantly reducing the time required for manual annotation.

Conclusion

It is concluded that the use of synthetic datasets represents a viable and promising approach for training artificial intelligence models aimed at agricultural disease detection. The results obtained demonstrate that the automated generation of realistic images and their respective annotations enables the construction of extensive, diverse datasets suitable for computer vision applications in the agricultural context.

Among the main advantages of this approach is the significant reduction in the time required to create and annotate datasets, since the labeling process can be carried out simultaneously with image generation. This factor considerably reduces the manual effort traditionally required in annotation tasks, as well as lowering operational costs related to data collection and preparation.

Another relevant aspect is the independence from external factors, such as weather conditions, seasonality, and crop availability at different times of the year. Synthetic generation allows the reproduction of varied scenarios in a controlled manner, enabling the creation of data under different environmental conditions without the need to wait for specific growing periods or favorable weather conditions. This flexibility contributes to the development of more robust models, capable of operating in real environments subject to high variability.

In this way, synthetic datasets have high potential to accelerate research and applications in smart agriculture, making the development of artificial intelligence-based solutions more efficient, scalable, and accessible.

Future advancements

As a perspective for future work, the implementation and development of a system capable of providing greater climatic variability in the generation of synthetic scenarios is highlighted. Such an approach will allow the creation of datasets representative of different environmental conditions, extrapolating the predominantly sunny pattern currently used and encompassing situations such as cloudy, rainy, or low-light conditions.

Incorporating this variability tends to significantly enhance the robustness and generalization capacity of artificial intelligence models trained with this data, since the algorithms will be able to learn visual patterns under different lighting, contrast, and visibility conditions. Consequently, an improvement in the reliability and performance of pest and disease detection systems in real field applications is expected, where climatic and lighting conditions present high variability.

Funding

This work has been fully funded by the project FieldEye supported by the Center for Embedded Devices and Research in Digital Agriculture (CEDRA) of SENAI-RS, with financial resources from the PPI IoT/Manufatura 4.0 / PPI HardwareBR of the MCTI, grant number 053/2023, signed with EMBRAPII.

References

Alzubaidi, L., Bai, J., Al-Sabaawi, A., Santamaría, J., Albahri, A. S., Al-Dabbagh, B. S. N., ... & Gu, Y. (2023). A survey on deep learning tools dealing with data scarcity: definitions, challenges, solutions, tips, and applications. Journal of Big Data, 10(1), 46.

Barbedo J. G. A. (2016). A review on the main challenges in automatic plant disease identification based on visible range images

Beck, J. (2023). Quality aspects of annotated data: A research synthesis. AStA Wirtschafts-und Sozialstatistisches Archiv, 17(3), 331-353.

Blinn, J. F. (1978). Simulation of wrinkled surfaces. ACM SIGGRAPH Computer Graphics, 12(3), 286-292.

García Liñán, F., & Muñoz Orbañanos, A. (2023). Physically-Based Sky and Atmosphere Rendering in Real-Time

Khan, K., Lobiyal, D. K., & Kilicman, A. (2015). Bézier curves and surfaces based on modified Bernstein polynomials. arXiv preprint arXiv:1511.06594.

Kholghi, M., Sitbon, L., Zuccon, G., & Nguyen, A. (2017). Active learning reduces annotation time for clinical concept extraction. International journal of medical informatics, 106, 25-31.

Krecklau, L., & Kobbelt, L. (2012). Interactive modeling by procedural high-level primitives. *Computers & Graphics*, *36*(5), 376-386.

Liu, L., Ouyang, W., Wang, X., Fieguth, P., Chen, J., Liu, X., & Pietikäinen, M. (2020). Deep learning for generic object detection: A survey. International journal of computer vision, 128(2), 261-318.

Najafabadi, M. M., Villanustre, F., Khoshgoftaar, T. M., Seliya, N., Wald, R., & Muharemagic, E. (2015). Deep learning applications and challenges in big data analytics. Journal of big data, 2(1), 1.

Preetham, A., Shirley, P., & Smits, B. (1999). A practical analytic model for daylight. In Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (pp. 91–100). ACM Press/Addison-Wesley Publishing Co..

Rechkemmer, A., Williams, A. C., Lease, M., & Li, L. E. (2023). Characterizing Time Spent in Video Object Tracking Annotation Tasks: A Study of Task Complexity in Vehicle Tracking. In Proceedings of the AAAI Conference on Human Computation and Crowdsourcing (Vol. 11, pp. 140-151).

Schröder, C., & Niekler, A. (2020). A survey of active learning for text classification using deep neural networks. arXiv preprint arXiv:2008.07267.

Taylor, S., Sharpe, O., & Peethambaran, J. (2021). Prime gradient noise. Computational Visual Media, 7(3), 349-362.