From Photo to Fabrication (Photo2Fab)

A Design-to-Fabrication Workflow using Irregular Tree Parts, Machine Learning, and Mixed Reality

Vernelle A. A. Noel[1], Karthick Bangaru Giridharan[2] and Shahzadi Padda[3]
[1,2,3] Carnegie Mellon University
[1]vnoel@andrew.cmu.edu, 0000-0001-6374-8217
[2]kbangaru@andrew.cmu.edu, 0009-0000-8791-7222
[3]spadda@andrew.cmu.edu, 0009-0008-7258-497X

Abstract. As concerns about climate change, sustainability, and resource circularity grow, there is a need to consider the role of waste materials. In this project, we explore the development of a computational design workflow for design and fabrication with non-standardized tree parts of varied shapes and sizes. We describe and compare three machine learning workflows: (1) Generative Adversarial Networks (GANs), (2) Denoising Diffusion Probabilistic Models (DDPM), and (3) Line Segment Detection (LSD) models for feature detection by testing their efficiency and the quality of their results for converting photos of irregular tree parts into usable digital geometry. How can we leverage tree waste, machine learning, and mixed reality to create novel design-to-fabrication workflows in architecture and construction? We then use the best performing workflow to build a digital materials library, digitally model and design with the resulting geometry, and use mixed reality tools and software to build structures. Our results show that DDPMs outperform GANs in detecting and generalizing skeletons of real tree branches from photos. While our dataset of Y-branches were predominantly planar in the Z-axis (depth), our project demonstrates that Photo2Fab is a bottom-up, ecologically responsive approach to material reuse in design, architecture, and construction.

Keywords. Sustainability, Tree Material, Machine Learning (ML), Computer Vision, Mixed Reality (MR), Photo2Fab, DDPM, LSD

1. Introduction

As concerns about climate change, sustainability, and resource circularity grow in the architecture and construction industries, there is a need to consider the role of waste materials. Design and construction with irregular materials of unpredictable shapes and sizes present a challenge as construction typically involves the processing and use of

HUMANISTIC COMPUTATION AND INTELLIGENCE, Proceedings of the 31st International Conference of the Association for Computer-Aided Architectural Design Research in Asia (CAADRIA) 2026, Volume 2, 103-112.

standardized materials with known dimensions and properties (Zboinska & Göbel, 2025). In this project, we explore the development of a novel computational design workflow for design and fabrication with non-standardized tree parts. Irregularly shaped tree parts and offcuts are often discarded during landscaping and construction. Additionally, increasingly extreme and unexpected weather events have resulted in the breaking of entire trees, branches, and limbs. Emerging computational tools and methods, such as computer vision and extended reality (XR), offer opportunities for new methodologies and a rethinking of architecture's relationships with nature and the built environment. This research project addresses these challenges by using these tools to create a novel design-to-fabrication pipeline we call, Photo2Fab. Our project utilizes computer vision to convert photos of irregular tree parts into usable digital geometry for design. It also employs mixed reality (MR) to assemble and fabricate structures with the tree parts. By developing a computational design-to-production pipeline, the project demonstrates a bottom-up, ecologically responsive approach to material reuse in architecture and construction.

Three objectives of this study are to: (1) reduce tree waste and upcycle non-standard parts to minimize environmental impact; (2) convert photos of geometrically irregular tree parts into design and fabrication-ready digital models; (3) develop a novel photo-to-fabrication workflow to design and build with tree parts. This paper is organized as follows: Background, situates our project within existing scholarship on sustainability and ecological fabrication, digital inventory, and XR in design and fabrication. In Methods, we describe three methods for converting tree geometry into digital geometry, and our digital modeling, design, and fabrication processes. Results, show outputs from the machine learning models, and our culminating structures, while Discussion and Conclusion evaluate the results, consider implications of the work, and outline possible future directions.

2. Background

2.1. SUSTAINABILITY AND ECOLOGICAL FABRICATION PRACTICES

Sustainability promotes resource efficiency, use of renewable materials, and low-impact construction methods. A circular economy advances sustainability in architecture by replacing the traditional start-end linearity with closed-loop systems, where materials are reused, repaired, or recycled (Kibert et al., 2000). Ecological fabrication integrates natural systems, material behavior, and computational design to create fabrication processes that align with ecological dynamics (Guy and Farmer 2001). Existing scholarship combining natural systems and computational design, includes the development of digital and robotic fabrication workflows for non-standardized tree parts (Haghnazar et al., 2024; Wei, 2023); 3D printing with biomaterials (Bell et al., 2025), wood (Dayyem Khan, 2023), and mycelium composites (Dessi-Olive, 2023).

2.2. IRREGULAR MATERIALS AND DIGITAL INVENTORY

Work in documenting and inventorying tree parts includes 3D scanning (Amtsberg et al., 2020); digital modeling (Allner et al., 2020); extracting scanned mesh centerlines

(Cousin et al., 2023); and photogrammetry with smartphones (Chai et al., 2024). These methods are sometimes expensive and require substantial computational power for processing. While replacing 3D scanning with computer vision algorithms or deep learning models presents opportunities, it also presents challenges like capturing the complex geometries of branches, noise and occlusion, and the uncontrolled nature of environments (Kim & Kantor, 2023; Y. Yang et al., 2024). Current research in machine learning methods for detecting irregular tree branches includes Pix2Pix GAN models (Chen et al., 2021; Tong et al., 2022); lightweight convolutional neural network (CNN) models (Wan et al., 2023); central skeleton extraction with cross sections (C. H. Yang et al., 2020); and CNN for probability prediction (Silva et al., 2022).

2.3. EXTENDED REALITY IN DESIGN, FABRICATION & ASSEMBLY

Augmented Reality (AR) and Mixed Reality (MR) are increasingly explored as these technologies become more ubiquitous, affordable, and mobile. Digital software environments and plugins often include Rhino 3D and Grasshopper 3D (digital modeling software by Robert McNeel & Associates), Microsoft HoloLens headsets, and the AR plugin for Rhino 3D, Fologram. Current work includes using AR to manually assemble complex structures from standardized dimensional lumber (Hanke et al., 2023), dowels (Mostafavi et al., n.d.), and friction-fit planes (Atanasova et al., 2023). Manual XR fabrication studies include the fabrication of low-cost, reusable molds (Azambuja Varela et al., 2022), fabrication and assembly of a steam-bent timber structure (Jahn, 2019), and a curved art installation made from bamboo splits (Goepel & Crolla, 2020). Work in manual fabrication and assembly with AR/ MR and irregular tree parts includes a tree-log structure (Lok, 2021), and a gridshell structure (Cousin et al., 2023).

3. Methods

3.1. COMPUTER VISION MODELS TO CAPTURE TREE GEOMETRY

3.1.1. Line Segment Detection (LSD) Model

Image segmentation is a computer vision technique that groups pixels of similar appearance while Line Segment Detection (LSD) models extract line segments from images (Lin et al., 2020; Szeliski, 2022). LSD models extract branch geometry and estimates a straight line representing the branch instead of point clouds and meshes. In our project, we photographed Y-branches on white sheets of paper to reduce environment noise. Photographs were then cropped and resized to 256 x 256 pixels then processed by the LSD model. The model predicted branch features to match the image, but with outer boundary lines rather than center line branch skeletons (Figure 1). Center lines were drawn in Rhino 3D, then piped to match branch diameters.

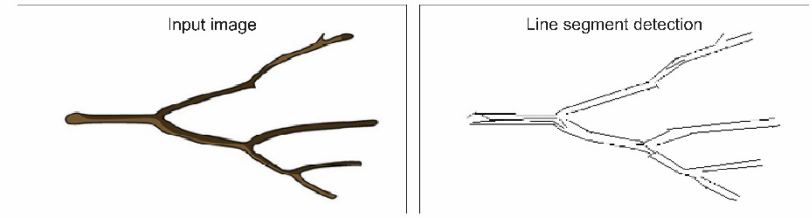

Figure 1. Photo of branch as input (left) and resulting predicted branch features from LSD model (right)

3.1.2. Pix2Pix GAN and Line Segment Detection (LSD) Models

Semantic segmentation models output outer boundaries of objects and cannot detect the central axis lines (Chen et al., 2021). To address this, we developed a Pix2Pix GAN model trained on a custom dataset of 5000 paired synthetic images of branches and their center line skeletons to predict RGB images representative of the branches. To create this synthetic dataset, we used L-systems in Grasshopper 3D to generate rendered branches of varying thickness, materials, and backgrounds in Rhino 3D (Figure 2). To test it, we converted RGB files of 256 x 256 pixels to black and white, then fed them to the Pix2Pix GAN model, and applied the line segment detection (LSD) model to get outer boundaries and vectors. The resulting vector line segments were plotted with the Python library Matplotlib and imported into Rhino 3D with a custom Python plug-in to make the geometry visible. These lines were then scaled up to match the real branch dimensions. The Pix2Pix GAN model was trained on an A100 GPU for 80,000 steps with binary cross entropy and L1 loss.

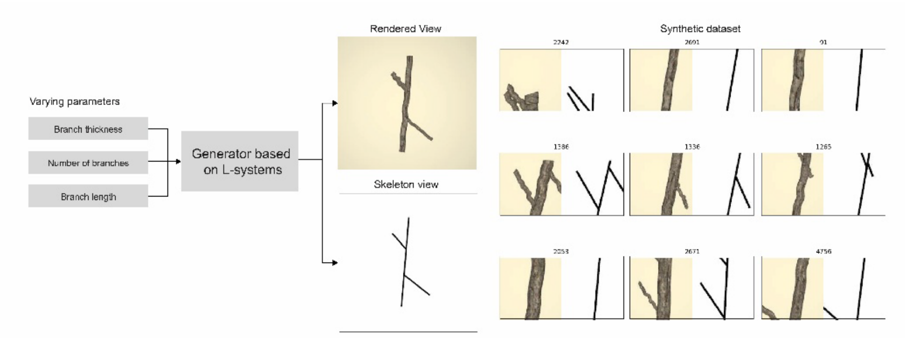

Figure 2. Parameters and synthetic dataset of rendered branches and their center line skeletons used to train Pix2Pix GAN model. This dataset also used to train DDPM model.

3.1.3. Denoising Diffusion Probabilistic Models (DDPM) and LSD

While our Pix2Pix GAN model produced samples closely matching the training images, its performance on real photos did not properly match. To address this, we utilized a Denoising Diffusion Probabilistic Model (DDPM), a machine learning technique for image generation. DDPMs generate high quality, realistic data from noisy images by reconstructing the data from noise (Ho et al., 2020; Zhao et al., 2025). DDPMs also perform better at data augmentation, thereby improving coverage and robustness. We built the DDPM model with the same synthetic dataset used for the Pix2Pix GAN model. To test it, we fed the DDPM our sized black and white images, plotted resulting vector line segments, and imported them into Rhino 3D (Figure 3). Our model was trained on A100 GPU for 10,000 steps with Mean squared error loss.

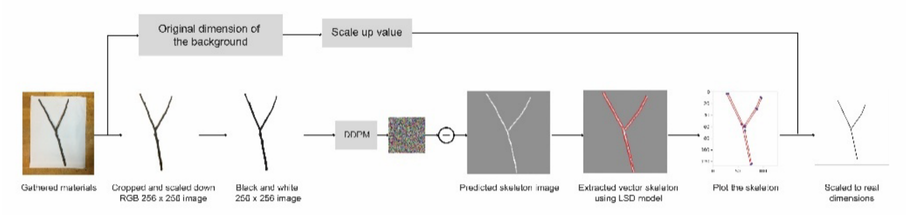

Figure 3. Image of irregular branches from the internet used as input (left), and line segment detection (LSD) model results (right)

3.2. DIGITAL DESIGN WITH IRREGULAR TREE PARTS

We collected Y-branches for the design and fabrication stage after finalizing our DDPM and LSD workflow for extracting geometry from photographs of irregular tree parts. Each branch was labelled before being photographed to improve sorting and recognition in the modeling and design phase (Figure 4).

Figure 4. Collection of labelled Y-branches

3.2.1. Digital Modeling and Sorting of Branches

We imported the vector outputs from our model into Rhino 3D, to digitally model, sort, design, and configure assembly logics and strategies. To digitally model the branches, we piped centerlines to match branch diameters, then labelled each branch in Rhino 3D for identification and assembly. We examined branch characteristics such as slant angles, height of branch crotch, narrow and wide fork distances, then grouped them based on these metrics. Since our system is currently unable to capture branch depth (Z-axis), we photographed our Y-branches to capture the more complex of the two sides, and privileged branches that did not have large deviations in the Z-axis. This was done to reduce the difference in depth dimension between 2D digital branches and 3D real branches. We designed and fabricated with 2D translations of the 3D branches.

3.2.2. Digital Design with Irregular Branches

We then designed three structures with branches in our digital material library (Figure 5). Structure A comprised eight Y-branches and measured 12” x 12” x 13” (W x L x H) in size. Tall Y-branches would be used for the sides of the arches with Y-branches alternating - one up and one down - while shorter branches rested horizontally across the tall branches. Structure B was dome shaped, comprised 19 Y-branches and measured 14” x 15” x 11”. Structure C was trapezoidal-shaped, comprise 14 Y-branches and measured 36” x 48” x 40”.

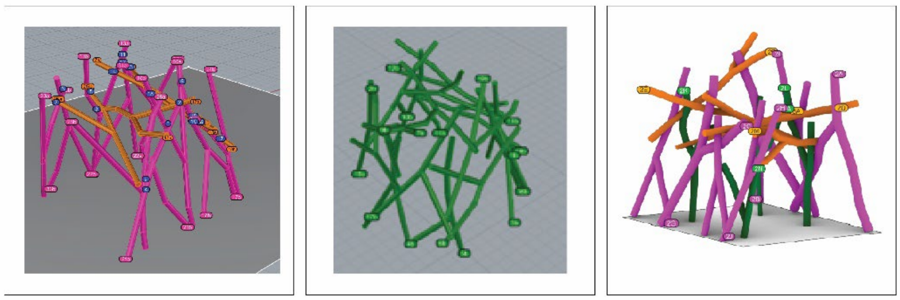

Figure 5. Digital models of designs. Structure A (left), Structure B (middle) and Structure C (right)

3.2.3. Mixed Reality Assembly and Manual Fabrication

The final phase of the workflow involved deploying mixed reality headsets (HoloLens 2) and AR software, Fologram, a plug-in for Grasshopper 3D, to superimpose the digital model from Rhino 3D into the physical workspace, thereby positioning and aligning real tree parts to match the digital model. For structures A and B, branches were secured with hot glue, while for the larger structure, C, we employed square lashing with waxed 9-ply 2mm twine of high tensile strength to rigidly secure the branches. Branch forks and connection points were labelled and marked to differentiate orientation and indicate connection location. Figure 6 shows mixed reality overlays in real space, the square lash, and assembly of the structure by designers wearing HoloLenses.

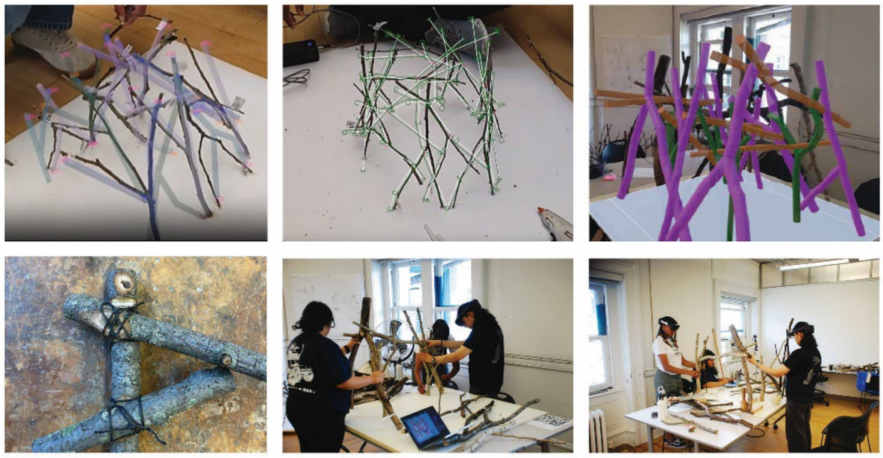

Figure 6. Top row: View through HoloLens showing overlay of digital model onto real space for structures A, B, and C (top left, middle, and right). Bottom row: Square lashing technique (left) and designers fabricating structures with the HoloLens headset on their heads.

4. Results

When it comes to the performance of the machine learning models, GAN achieves an overall accuracy of ~0.96, precision ~0.03, recall ~0.16, F1 score ~0.05, and IoU ~0.03. On the other hand, DDPM achieves ~1.00 accuracy, 0.66 precision, 0.76 recall, 0.71 F1, and 0.55 IoU. These scores are shown in Figure 7.

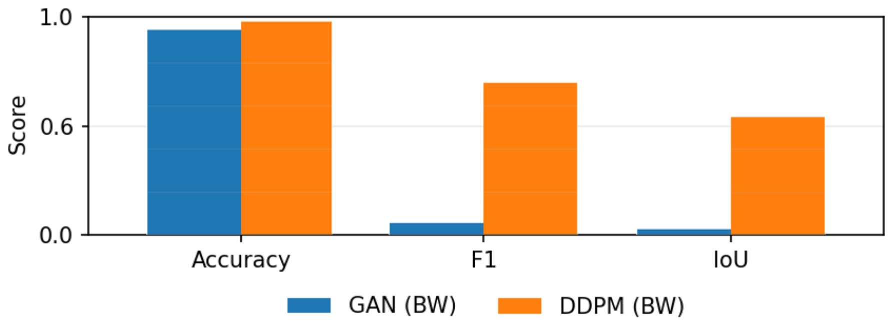

Figure 7. Comparative performance between GAN and DDPM models

Additionally, we present the three structures built with this computational design-to-production workflow comprising Y-branches, computer vision models, digital design and modelling, and augmented reality tools for assembly (Figure 8).

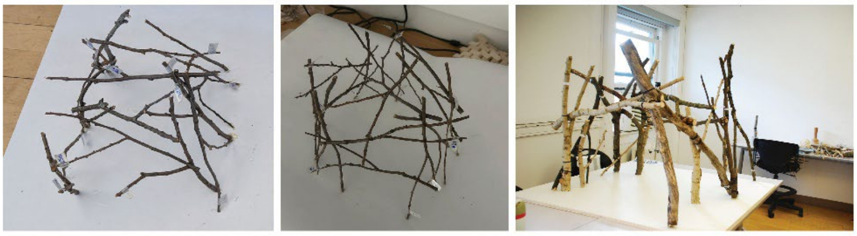

Figure 8. Fabricated structures built with augmented reality (AR) from irregular Y-branches. Structure A (left), structure B (middle) and structure C (right)

5. Discussion and Conclusion

This study aimed to leverage computer vision, digital design and modeling, and augmented reality — to transform otherwise discarded, irregular tree parts into usable building elements. Results show that our Photo2Fab workflow comprising a Denoising Diffusion Probabilistic Model (DDPM) performs better than GAN models in detecting and generalizing skeletons of photos of real tree branches. This project also demonstrates a bottom-up, ecologically responsive approach to material reuse in design, architecture, and construction.

Our finding that DDPM models outperform GAN models, is supported by our experiments and the models' performance, as shown in Figure 7. We evaluated performance based on precision (fraction of correctly predicted 'twig' pixels), recall (fraction of actual 'twig' pixels identified), and the F1-Score (the harmonic mean of precision and recall). In a dataset of approximately 5,000 paired images, we quantitatively compared GAN and DDPM-based predictors using pixel-wise segmentation metrics calibrated for class imbalance. Although the GAN achieves high overall accuracy, precision, recall, F1 score, and IoU, it yields few correct twig pixels. By contrast, DDPM achieves higher scores on all these metrics, reflecting better localization and coverage. These results stress that under severe class imbalance, accuracy is not a reliable diagnostic, and F1 and IoU are more informative, demonstrating better generalization of DDPM to target geometry.

One limitation of this study is that currently the resulting geometry from the model is two-dimensional (2D). This limitation arises from the high computational cost of converting 3D scans into lightweight 3D geometry. This work is still valid, however, since the deviation distance of branches from the mean in the third dimension did not affect design or fabrication at our current scale. Additional directions for future research include developing a ML model that is able to recognize branch depth (Z-axis) so that digital output is 3D, building a Generative Model that can take the geometry of irregular tree parts as input and create three-dimensional digital design options by analyzing and parsing the parts using metrics such as dimensions, shape, and other relevant characteristics.

Acknowledgements

We would like to thank the Situated Computation + Design Lab (Sit.Co.De Lab) and Carnegie Mellon Architecture for funding the work. We thank the PJ Dick Innovation Fund for partial funding, CMU-Grounds for providing us with tree parts, CodeLab for sharing workspace, and CMU HURAY and SURA for research interns. We thank our interns: Tvisha Arora, Mahika Singh, Starr Wasler, and Brian Zhang for designing and fabricating large structures, and Jon Holmes for his advice and feedback.

References

Allner, L., Kroehnert, D., & Rossi, A. (2020). *Mediating Irregularity: Towards a Design Method for Spatial Structures Utilizing Naturally Grown Forked Branches.* In C. Gengnagel, O. Baverel, J. Burry, M. Ramsgaard Thomsen, & S. Weinzierl (Eds.), Impact: Design With All Senses (pp. 433–445). Springer International Publishing.

Atanasova, L., Saral, B., Krakovská, E., Schmuck, J., Dietrich, S., Furrer, F., Sandy, T., D'Acunto, P., & Dörfler, K. (2023). *Collective AR-Assisted Assembly of Interlocking Structures*. In C. Gengnagel, O. Baverel, G. Betti, M. Popescu, M. R. Thomsen, & J. Wurm (Eds.), Towards Radical Regeneration (pp. 175–187). Springer International Publishing. https://doi.org/10.1007/978-3-031-13249-0_15

Azambuja Varela, P., Lacroix, I., Güzelci, O. Z., & Sousa, J. P. (2022). *Democratizing Stereotomic Construction through AR Technologies—A reusable mold methodology to the production of customized voussoirs using HoloLens.* 225–232. https://doi.org/10.52842/conf.ecaade.2022.1.225

Bell, F., Friedman-Gerlicz, C., Urenda, L., & Buechley, L. (2025). *3D Printing Eggshells: Exploring Eco-Socio-Technical Relations through Biomaterial Design*. Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, 1–19.

Chai, H., Zhou, X., Gao, X., Yang, Q., Zhou, Y., & Yuan, P. F. (2024). *Integrated workflow for cooperative robotic fabrication of natural tree fork structures*. Automation in Construction, 165, 105524. https://doi.org/10.1016/j.autcon.2024.105524

Chen, Z., Ting, D., Newbury, R., & Chen, C. (2021). *Semantic segmentation for partially occluded apple trees based on deep learning. Computers and Electronics in Agriculture*, 181, 105952. https://doi.org/10.1016/j.compag.2020.105952

Cousin, T., Alkhayat, L., Pearl, N., Dewart, C. B., & Mueller, C. (2023). *Wild Wood Gridshells: Mixed-Reality Construction of Nonstandard Wood*. Technology|Architecture + Design, 7(2), 216–231. https://doi.org/10.1080/24751448.2023.2245725

Dayyem Khan, M. V. (2023). *BioMatters: The Robotic 3D-Printed Biodegradable Wood-Based Formwork for Cast-in-place Concrete Structures.* ACADIA 2023: Habits of the Anthropocene: Scarcity and Abundance in a Post-Material Economy [Volume 2: Denver. 26-28 October 2023. Edited by A. Crawford, N. Diniz, R. Beckett, J. Vanucchi, M. Swackhamer 174-183.

Dessi-Olive, J. O. (2023). *Computational Design and Craft of PHOENIX: A Hanging Myco-Structure*. ACADIA 2023: Habits of the Anthropocene: Scarcity and Abundance in a Post-Material Economy [Volume 2: Denver. 26-28 October 2023. Edited by A. Crawford, N. Diniz, R. Beckett, J. Vanucchi, M. Swackhamer 184-193.

Goepel, G., & Crolla, K. (2020). *Augmented Reality-based Collaboration—ARgan, a bamboo art installation case study.* 313–322. https://doi.org/10.52842/conf.caadria.2020.2.313

Haghnazar, R., Ashjazadeh, Y., Hauptman, J., & Nasir, V. (2024). *A computational design integrated digital fabrication framework for mass customization in industry 5.0 manufacturing with non-standard natural materials.* Results in Engineering, 23, 102400. https://doi.org/10.1016/j.rineng.2024.102400

Hanke, T., Kawasaki, J. Y., & Grunwald, G. (2023). *Manufacturing processes of complex shapes and structures using 3D printing and augmented reality*. Proceedings of the International Conference of Contemporary Affairs in Architecture and Urbanism-ICCAUA, 6(1), 62–66. https://doi.org/10.38027/iccaua2023en0134

Ho, J., Jain, A., & Abbeel, P. (2020). *Denoising Diffusion Probabilistic Models* https://doi.org/10.48550/arXiv.2006.11239

Jahn, G. W. (2019). *[BENT]*. ACADIA 19:UBIQUITY AND AUTONOMY (The University of Texas at Austin School of Architecture, Austin, Texas 21-26 October, 2019) P 438-47.

Kim, C. H., & Kantor, G. (2023). *Occlusion Reasoning for Skeleton Extraction of Self-Occluded Tree Canopies*. 2023 IEEE International Conference on Robotics and Automation (ICRA), 9580–9586. https://doi.org/10.1109/ICRA48891.2023.10160650

Lin, Y., Pintea, S. L., & Gemert, J. C. van. (2020). *Deep Hough-Transform Line Priors* https://doi.org/10.48550/arXiv.2007.09493

Lok, L. S. (2021). *Timber De-Standardized*. ACADIA 2021: Realignments: Toward Critical Computation. Online and Global. 3-6 November 2021. Edited by B. Bogosian, K. Dörfler, B. Farahi, J. Garcia Del Castillo y López, J. Grant, V. Noel, S. Parascho, and J. Scott. 222-231.

Mostafavi, S., Mehan, A., Montejano, E., Howell, C., & Stuckemeyer, J. (n.d.). *FabriCity-XR: A Phygital Lattice Structure Mapping Spatial Justice – Integrated Design to AR-Enabled Assembly Workflow.*

Silva, R., Junior, J. M., Almeida, L., Gonçalves, D., Zamboni, P., Fernandes, V., Silva, J., Matsubara, E., Batista, E., Ma, L., Li, J., & Gonçalves, W. (2022). *Line-based deep learning method for tree branch detection from digital images.* International Journal of Applied Earth Observation and Geoinformation, 110, 102759. https://doi.org/10.1016/j.jag.2022.102759

Szeliski, R. (2022). *Computer Vision: Algorithms and Applications*. Springer Nature.

Tong, S., Yue, Y., Li, W., Wang, Y., Kang, F., & Feng, C. (2022). Branch Identification and Junction Points Location for Apple Trees Based on Deep Learning. Remote Sensing, 14(18), 4495. https://doi.org/10.3390/rs14184495

Wan, H., Zeng, X., Fan, Z., Zhang, S., & Kang, M. (2023). *U2ESPNet—A lightweight and high-accuracy convolutional neural network for real-time semantic segmentation of visible branches.* Computers and Electronics in Agriculture, 204, 107542. https://doi.org/10.1016/j.compag.2022.107542

Wei, J. C. L. (2023). *Responsive Robotic Assembly System With Heterogeneous Materials: A Case Study with Unprocessed Wood.* ACADIA 2023: Habits of the Anthropocene: Scarcity and Abundance in a Post-Material Economy [Volume 2: Denver. 26-28 October 2023. Edited by A. Crawford, N. Diniz, R. Beckett, J. Vanucchi, M. Swackhamer 352-360.

Yang, C. H., Xiong, L. Y., Wang, Z., Wang, Y., Shi, G., Kuremot, T., Zhao, W. H., & Yang, Y. (2020). *Integrated detection of citrus fruits and branches using a convolutional neural network.* Computers and Electronics in Agriculture, 174, 105469. https://doi.org/10.1016/j.compag.2020.105469

Yang, Y., Lai, H., Chen, B., Huo, Y., Xia, K., & Huang, J. (2024). *Automatic Method for Extracting Tree Branching Structures from a Single RGB Image*. Forests, 15(9), 1659. https://doi.org/10.3390/f15091659

Zboinska, M. A., & Göbel, F. (2025). *Digital tool integrations for architectural reuse of salvaged building materials.* Automation in Construction, 170, 105947. https://doi.org/10.1016/j.autcon.2024.105947

Zhao, P., Zhang, W., Cao, X., & Li, X. (2025). *Denoising diffusion probabilistic model-enabled data augmentation method for intelligent machine fault diagnosis.* Engineering Applications of Artificial Intelligence, 139, 109520. https://doi.org/10.1016/j.engappai.2024.109520